



HAL
open science

Perceiving the world under the strobe of attention: psychophysical and electroencephalographical investigations

Julien Dubois

► **To cite this version:**

Julien Dubois. Perceiving the world under the strobe of attention: psychophysical and electroencephalographical investigations. Life Sciences [q-bio]. Université Paul Sabatier - Toulouse III, 2011. English. NNT: . tel-00619137

HAL Id: tel-00619137

<https://theses.hal.science/tel-00619137v1>

Submitted on 5 Sep 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par *Université Toulouse 3 Paul Sabatier (UT3 Paul Sabatier)*

Discipline ou spécialité : *Neurosciences*

Présentée et soutenue par *Julien Dubois*

Le vendredi 2 septembre 2011

Titre : *Perceiving the world under the strobe of attention :
psychophysical and electroencephalographical investigations*

*(Percevoir le monde sous le stroboscope attentionnel :
études psychophysiques et électroencéphalographiques)*

JURY

Jean-Philippe Lachaux - INSERM, Bron

Ole Jensen - Donders Centre for Cognitive Neuroimaging, Nijmegen, Pays-Bas

Andreas Kleinschmidt - INSERM, Gif-sur-Yvette

Arnaud Delorme - CNRS, Université Paul Sabatier, Toulouse

Pier-Giorgio Zanone - Université Paul Sabatier, Toulouse

Rufin VanRullen - CNRS, Université Paul Sabatier, Toulouse

Rapporteur

Rapporteur

Examineur

Examineur

Examineur

Directeur de Thèse

Ecole doctorale : *CLESCO*

Unité de recherche : *Centre de Recherche CERveau et COgnition, UMR5549*

Directeur(s) de Thèse : *Rufin VanRullen*

Rapporteurs : *Ole Jensen et Jean-Philippe Lachaux*

A lifetime is roughly 20 billion moments.

J. M. Stroud (Stroud, 1967)

To my wife Christine

*who let me spend millions of our precious moments away
so I could perform this work.*

Acknowledgements

How does our material brain--the most complex physical system known--produce our immaterial but vital sense of awareness? [...] The key to finding an answer, Koch says, is to trace the activity of neurons--the "neural correlates"--of the simplest type of consciousness, which is the awareness of something we see. "Some of my colleagues think I'm naive," Koch remarks, "that this rather narrow focus won't reveal the workings. And they might be right. But as a scientist, I think this is the most likely way to solve this problem."

Above is an excerpt from the article *The Quest of Christof Koch* published in June 2005 in Scientific American Mind. In late 2005 I stumbled upon this article, and it is the main reason that I am here now, writing this thesis.

In 2002, after two years of classes préparatoires (an intensive syllabus of Math, Physics, Chemistry, Biology and Earth Sciences), I entered the École Normale Supérieure at rue d'Ulm in Paris where I majored in Earth and Planetary Science. The first year of my masters was highlighted by a six-month internship in the biogeochemistry lab of the late Pr. Harold Helgeson at UC Berkeley – where I incidentally met my wife Christine. As this year drew to a close, I became aware that I was not passionate enough for the field to pursue a PhD. I reoriented myself slightly, with a second masters year in biochemistry, focusing on the problem of the origins of Life. A 5-month internship in Rome, Italy, in the lab of Pr. Pier Luigi Luisi, working on the Minimal Cell project, left me even more perplexed... I thus decided to take a year off from my studies, travelling with Christine and waiting for something to happen. This “something” turned up in the form of the Scientific American Mind article I mention above.

After reading the article, I purchased Christof Koch's book, *The Quest for Consciousness*, and read it avidly. By the end of the book I knew I had found my way. I boldly contacted Christof by email, proposing that I come work for him for a year, starting in the fall of 2006. After meeting with him, he agreed to let me work with him, take classes at Caltech and learn about neuroscience. He also introduced me to my other role model: Rufin VanRullen, my thesis advisor. In the summer of 2006 I met Rufin, and knew right away that I wanted to work with him. I started this thesis work with Rufin in the fall of 2007, continuing to spend a lot of time in the Koch lab at Caltech, so I could live with Christine in California. We were married in the summer of 2010.

It has been a wonderful 5 years both professionally and personally since I read that Scientific American Mind article. I hope throughout this thesis that I can share some of this scientific and personal happiness. But first I want to thank a few people for making this thesis possible.

When a young scientist has made up his mind and wants to obtain a PhD degree, he is looking for a Doktorvater and presents himself as a doctoral candidate. He hopes to find the optimal support from this specific professor. The Doktorvater, on his side, decides to accept the young scientist because he trusts in him; he desires to pass his experience, knowledge and scientific ethics through to the next generation.

Ziko van Dijk, blog entry (28 Feb 2011)

Rufin – you have been a true Doktorvater to me. You are one of the brightest and most insightful scientists I have ever met, and it was an amazing opportunity to work with you and learn from your practice of science. It is quite stimulating to watch you think, and try to follow your racing mind! You are also the best adviser anyone could wish for, available and

motivating; at all times when I needed your support, you were there. And then there was wakeboarding, volleyball, meals at your house... I made a great friend (also with you, Leila), and hopefully a lifelong collaborator. I thank you for the trust you had in me, and I apologize for not being around as much as you probably would have wished, not bringing you as much in return as I got from you...

CerCo – thank you for welcoming me each time I made an unexpected appearance in Toulouse... I've always felt part of the family, even though I was a pretty distant relative. So, thank you to everyone in the lab. More especially... Marianne, Seb & Téry, James & Charlotte, Rufin & Leila, Jan, Maxime & Romain & Gladys & Flo : thank you for offering me a place to stay at times when I was "homeless". I realize I've lived in 12 different places over the last four years, including a Formula 1 hotel! Simon, Jean-Michel, Rufin, thank you for the lunches at the grown-up canteen, and the stimulating discussions which invariably happened. Laura, Rodi, thank you for asking almost every day if I wanted to have lunch at the student canteen, despite my almost invariably negative answer... Michèle, and Claire, thank you for always being available and helpful in all sorts of administrative matters, and making life at Cerco as easy as humanely possible. Maxime and Romain : thanks for taking care of my daily dose of sports in the last stages of writing this thesis manuscript!

Christine – thank you for bearing these long times we spent apart, especially the last 6 months, divided by the Atlantic pond and a few mountain ranges. I wish we could have done it differently... I'm coming home now, and you'll be stuck with me for the rest of your days. I'm looking forward to each and every one of them. I love you.

My families – Papa, Maman, Seb, Laurine and Tom, Margaret, Mary, Steve – thanks for loving me and taking care of me at all times.

June 1st 2011

Julien Dubois

Publications

Reuves internationales à comité de lecture

1. Dubois, J., Hamker, F. H., & VanRullen, R. (2009). *Attentional selection of noncontiguous locations: The spotlight is only transiently “split”*. **Journal of Vision**, 9(5):3
2. Busch, N. A., Dubois, J., & VanRullen, R. (2009). *The phase of ongoing EEG oscillations predicts visual perception*. **Journal of Neuroscience**, 29(24): 7869–7876
3. VanRullen, R., Busch, N. A., Drewes, J. & Dubois, J. (2011) *Ongoing EEG phase as a trial-by-trial predictor of perceptual and attentional variability*. **Frontiers in Psychology**, 2:60
4. Dubois, J., VanRullen, R. (2011). *Visual Trails: Do the Doors of Perception Open Periodically?* **PLoS Biology**, 9(5): e1001056
5. VanRullen, R., Dubois, J. (2011). *The Psychophysics of Brain Rhythms*. **Frontiers in Psychology**, 2:203

Posters et Présentations à des conférences internationales

1. Busch, N., Dubois, J. & VanRullen, R. (2009). *The phase of ongoing EEG oscillations predicts visual perception*, **Proceedings of the 9th Vision Sciences Society annual meeting @ Naples, FL, USA**, poster
2. Dubois, J., VanRullen, R. (2009). *Evaluating the contribution of discrete perceptual mechanisms to psychometric performance*. **Association for the Scientific Study of Consciousness meeting @ Berlin**, poster
3. Dubois, J., Macdonald, J. & VanRullen, R. (2010). *Broadband frequency tagging:Reevaluating the sustained division of the attentional spotlight at high temporal resolution*. **Proceedings of the 10th Vision Sciences Society annual meeting @Naples**, poster
4. Dubois, J., VanRullen, R. (2011). *Visual trails: When perceptual continuity breaks down*. **Proceedings of the 11th Vision Sciences Society annual meeting @Naples**, poster
5. Dubois, J., VanRullen, R. (2011). *Do the Doors of Perception open periodically?.* **European Conference on Visual Perception @Toulouse**, présentation orale

Abstract

Perceiving the world under the strobe of attention *Psychophysical and electroencephalographical investigations*

Our sensory experience of the world is smooth and continuous. Yet, it could rely on discrete sampling of incoming sensory information – this sampling being instantiated by attentional mechanisms. Under continuous lighting an observer attending to a rotating spoked wheel may experience illusory reversals, which have been interpreted as a temporal aliasing artefact, suggesting that attentional motion perception relies on position samples. We sought to use similar paradigms and better characterize the attentional rhythm, but practical difficulties arose. The purported periodicity of attention was put to the test in another context: we predicted that concurrently attended spatial locations should be sampled serially. With a novel analysis of an existing paradigm we indeed found evidence against a sustained division of the attentional spotlight. We also studied some pathological perceptual manifestations of motion perception which are compatible with underlying sampling mechanisms. Brain oscillations are likely to support the periodicities that we evidenced behaviorally. A fronto-central rhythm at about 7Hz could predict whether a faint visual stimulus at an expected location would be detected or completely missed. We also looked for a phasic influence of ongoing brain activity on the perception of simultaneity. Finally we sought to track the position of the attentional spotlight in “real time” in a paradigm enforcing attention to two concurrent locations to gain knowledge about its intrinsic rhythm. This work revealed periodicities in attention and perception; however, a complete theoretical understanding of how these rhythms truly shape perception remains ahead of us.

Keywords alpha, brain oscillations, discrete perception, phase, serial attention, temporal aliasing, temporal framing, visual trails

Percevoir le monde sous le stroboscope attentionnel *Études psychophysiques et électroencéphalographiques*

Notre expérience du monde est fluide, continue. Pourtant, elle pourrait reposer sur un échantillonnage discret, par l'attention, de l'information sensorielle entrante. Sous illumination continue un observateur portant attention à une roue à rayons en rotation peut percevoir des inversions illusoire; ceci semble être un artefact d'aliasing temporel, suggérant que la perception du mouvement par l'attention utilise des échantillons. Nous avons cherché des signes d'aliasing dans d'autres tâches, mais avons rencontré des obstacles pratiques. La périodicité présumée de l'attention nous a mené à prédire que des positions spatiales attendues simultanément devraient être échantillonnées tour à tour. Nous avons des résultats psychophysiques allant à l'encontre d'une division soutenue de l'attention spatiale. Nous avons également étudié certains désordres pathologiques de la perception du mouvement, compatibles avec des mécanismes perceptuels discrets. Certaines oscillations cérébrales sont certainement à l'origine des périodicités découvertes. Un rythme fronto-central à environ 7hz nous a permis de prédire si un sujet détecterait un stimulus visuel très faible apparaissant à une position attendue. Nous avons aussi cherché une influence phasique de l'activité spontanée du cerveau sur le jugement de simultanéité. Enfin, nous avons voulu suivre la position de l'attention spatiale en temps réel dans un paradigme où le sujet devait porter attention à deux endroits simultanément. Ce travail a révélé des périodicités de l'attention et de la perception; il reste à faire pour parvenir à une compréhension théorique complète de la façon dont ces rythmes forment notre expérience.

Mots-Clés alpha, oscillations cérébrales, perception discrète, phase, attention sérielle, aliasing temporel, cadrage temporel, traînées visuelles

Résumé substantiel

“Je ne crois que ce que je vois”. Cet adage populaire, attribué à St Thomas, traduit bien la confiance que nous humains portons à notre vision. Mais comment voit-on? en général cette question ne se pose pas, il s’agit de quelque chose qui se fait automatiquement en ouvrant les yeux, il n’y a rien de mystérieux. Mais pour celui qui étudie la vision, qui cherche à en comprendre les mécanismes fondamentaux, rapidement cette faculté semble relever du miracle. L’image du monde vu à travers la rétine est incroyablement déformée. La résolution est correcte autour du point de fixation, mais elle chute très vite dès que l’on s’en éloigne. Il y a un trou dans chaque rétine, pour laisser passer le nerf optique, car les neurones de la rétine sont du mauvais côté par rapport aux capteurs de lumière. Il n’y a pas de capteurs pour les longueurs d’ondes courtes au centre du champ de vision. Malgré ces déformations substantielles de l’information visuelle dès son acquisition, notre vision semble claire partout dans notre champ de vision, il n’y a pas de zones vides, et nous ne sommes pas conscients de percevoir moins de couleurs au point de fixation. Il faut dès lors se rendre à l’évidence: nous ne voyons pas vraiment ce que nos yeux voient, mais plutôt une construction mentale élaborée, qui doit utiliser un certain nombre d’astuces silencieuses pour nous donner l’illusion que notre vision reproduit parfaitement ce qui se passe au dehors. Les heuristiques utilisées par le cerveau sont peu à peu révélées par les chercheurs en sciences de la vision, par exemple au travers d’illusions optiques qui mettent ces heuristiques en défaut.

Notre vision du monde nous semble continue, sans interruption. En jouant au volleyball, par exemple au moment de recevoir un service, la position de la balle nous est connue à tout moment, elle semble être mise à jour à chaque milliseconde. Et s’il s’agissait là aussi d’une simple illusion? Si le cerveau nous donnait l’impression d’une perception continue, alors que celle-ci repose en réalité sur des mises à jour périodiques (avec une période d’au moins plusieurs dizaines de millisecondes)? C’est cette hypothèse, qui a été considérée à plusieurs reprises dans la littérature sans jamais convaincre la majorité des académiques, qui a guidé les travaux réalisés au cours de cette thèse.

Cette hypothèse n’est pas résolument farfelue. L’activité des neurones donne naturellement lieu à des oscillations, qui sont enregistrables à travers le scalp en électroencéphalographie (EEG) ou en magnétoencéphalographie (MEG). Ces oscillations représentent des modulations concertées du potentiel de membrane de millions de neurones, et ont des conséquences fonctionnelles sur l’activité de ces neurones – par exemple, favorisant la genèse de potentiels d’action dans des fenêtres temporelles restreintes. Ces conséquences, d’abord soupçonnées d’un point de vue théorique, ont été empiriquement démontrées *in vitro*, mais aussi *in vivo* chez l’animal. Si l’activité neuronale est ainsi contrainte par des rythmes spontanés, il est logique de proposer que ces contraintes devraient se répercuter sur nos fonctions cognitives, et notamment sur notre perception visuelle.

L’idée de mécanismes périodiques pour la perception fut mentionnée dès le dix-neuvième siècle par William James. Elle revit le jour avec l’invention de la cinématographie: cette technologie projette une séquence rapide de photographies (24 par seconde) sur un écran, donnant lieu à un percept apparemment continu. Se pourrait-il que le cerveau échantillonne le monde comme une caméra? La popularité de cette question atteint son maximum il y a quelque 50 ans, avec la théorie du “moment perceptuel” par Stroud (1956): selon lui, des échantillons sensoriels sont prélevés selon un processus périodique dont la période peut changer en fonction de la tâche à accomplir. Quels étaient les arguments avancés à l’époque en faveur de cette hypothèse? L’un des arguments forts reposait sur des études psychophysiques de la perception de l’ordre et de la simultanéité. Par exemple, Hirsh (1959) présenta successivement deux sons brefs, et observa qu’un ISI (intervalle entre les deux stimuli) de 100ms était nécessaire pour déterminer l’ordre de présentation à 95% correct. Ces résultats sont

également valables pour des stimuli visuels, et pour des stimuli inter-modaux. Ceci peut être interprété en termes de moments perceptuels : il faut que les deux stimuli soient présentés dans des moments perceptuels successifs pour être discriminables. Une autre expérience, réalisée par Lichtenstein (1961), se servit d'une présentation cyclique de 4 flashes lumineux (5ms), aux quatre coins d'un losange. Il montra qu'avec un cycle durant moins de 125ms, l'observateur percevait un clignotement synchronisé des 4 flashes – en faisant varier la séparation temporelle des flashes au sein d'un cycle, il montra également que l'ISI entre deux flashes n'avait aucune importance. Ceci semble indiquer que la perception repose sur des moments perceptuels d'environ 125ms chacun. White et ses collègues (1952) utilisèrent des clicks ou des flashes (ou des stimuli tactiles) à une fréquence de 10, 15 ou 30 stimuli par seconde, présentés pour une durée variable. Dans ces conditions, les observateurs semblent atteindre un plafond à 10-12 stimuli par seconde, dans toutes les modalités sensorielles, lorsqu'ils doivent rapporter le nombre de stimuli présentés dans chaque séquence. Ces études de simultanéité et de numérosité temporelle restent cependant interprétables en termes d'une simple période d'intégration temporelle, plutôt que de moments perceptuels. D'autres résultats sont plus difficiles à concilier avec une période d'intégration, comme la perception de causalité qu'étudia Shallice (1964) : il s'agissait de montrer un disque en mouvement touchant un autre disque, qui entrait en mouvement lui-même avec un certain délai. Shallice rapporta qu'un délai de moins de 70ms donnait lieu à la perception d'un lien causal direct; entre 70ms et 140ms le lien semblait indirect; enfin, pour plus de 140ms, les événements semblaient indépendants. Un autre résultat intéressant est la mise en évidence de périodicités dans des histogrammes de temps de réaction, avec des périodes de 25ms et de 100ms. Enfin, une expérience prometteuse fut réalisée par Latour (1967), qui consistait à présenter deux faibles flashes lumineux successivement, avec un intervalle temporel assez long entre les deux pour qu'ils soient facilement distinguables. Latour trouva que le seuil visuel pour la perception des deux flashes oscillait avec une période d'environ 25 à 30ms, correspondant à la période du phénomène d'échantillonnage hypothétique sous-jacent.

Le problème auquel se heurtent les scientifiques est qu'ils n'ont généralement pas accès à l'état du cerveau, notamment à la phase du rythme perceptuel supposé, au moment où ils présentent une stimulation. Il s'ensuit que toute méthode d'analyse consistant à moyenniser des dizaines d'essais pour en déduire la performance du sujet dans une condition particulière ne peut pas détecter les effets du rythme perceptuel. Par exemple, la détection d'un faible flash visuel a beau être dépendante d'une oscillation cérébrale, avec une probabilité de détection pouvant être exprimée comme $p_0(1+a.\sin(w.t))$, l'expérimentateur n'ayant pas accès à t ne pourra mesurer que p_0 , la probabilité de détection moyenne. La méthode astucieuse de Latour permet d'accéder à w , mais elle est difficile à appliquer en pratique.

Une illusion découverte récemment a ouvert une nouvelle fenêtre sur le débat de la perception discrète : l'illusion de la roue de chariot, sous illumination *continue*. Souvent au cinéma, une roue de voiture (ou une hélice d'avion) tourne dans le mauvais sens. Il s'agit alors d'un phénomène physique lié au fonctionnement discret de la caméra, connu sous le nom d'aliasing temporel. Par exemple, si une hélice monopale tourne dans le sens des aiguilles d'une montre, effectuant un tour complet toutes les 100ms, et qu'une caméra prend une photo de l'hélice toutes les 75ms, la pale tourne effectivement de 90 degrés dans le sens inverse des aiguilles d'une montre (270 degrés dans le sens des aiguilles d'une montre) d'une image à la suivante : lors de la projection du film elle sera perçue comme tournant dans le sens inverse des aiguilles d'une montre avec une période de 300ms, interprétation la plus probable de l'information enregistrée par la caméra. Aussi surprenant que cela puisse paraître, une variante de ce phénomène a lieu dans des conditions d'illumination continues (en plein jour). La version continue de l'illusion diffère de la version cinématographique sur plusieurs points. Notamment, le percept du mouvement erroné n'est pas stable : des périodes d'inversion alternent avec des périodes de perception du mouvement réel, avec une dynamique correspondant à beaucoup de phénomènes bistables en psychologie. La probabilité de percevoir des inversions (mesurée comme le pourcentage du temps total pendant lequel le sujet rapporte

une direction opposée à la vraie direction) est maximale pour une fréquence temporelle de 10Hz : si cette illusion est liée au traitement du flot perceptuel en une série d'instantanés, on peut en déduire que l'échantillonnage est réalisé à une fréquence d'environ 13.3Hz. L'illusion a été observée pour du mouvement de premier-ordre (défini par des modulations de luminance) aussi bien que pour du mouvement de second-ordre (défini par des modulations de contraste); ces deux types de mouvement étant traités différemment dans le cerveau, le phénomène est difficilement explicable en termes de processus de bas niveau. L'illusion requiert l'attention: si l'attention est dirigée sur une autre tâche (comme une présentation sérielle rapide de lettres), les observateurs sont moins affectés par l'illusion. La seule composante du spectre de l'électroencéphalogramme modulée par l'illusion est à ~13Hz, dans le lobe pariétal droit; les changements dans cette bande de fréquence peuvent prédire les inversions du percept au-dessus du niveau de la chance, deux secondes avant qu'elles soient rapportées par le sujet. Bien qu'étant essentiellement corrélacionnelle, cette observation semble peser en faveur de l'hypothèse d'échantillonnage attentionnel à environ 13Hz. En perturbant l'activité du lobe pariétal droit par stimulation magnétique transcranienne, il est possible de diminuer la probabilité des inversions illusoire – un effet qui n'est pas observé lorsque le lobe pariétal gauche est visé. Ce résultat incrimine le lobe pariétal droit de manière causale. Ces résultats ont mené Rufin VanRullen, mon directeur de thèse, à formuler l'explication suivante pour cette illusion : un système attentionnel de perception du mouvement fonctionne en capturant des échantillons périodiquement, et est en compétition avec un système automatique de perception du mouvement qui fonctionne en continu. Cette compétition, lorsque les interprétations des deux systèmes diffèrent, se traduit par des épisodes d'inversion de la direction du mouvement.

Un aspect de cette interprétation qui n'est pas souvent bien assimilé par la communauté est la proposition que l'attention est une ressource périodique. Une expérience d'attention divisée, utilisant jusqu'à quatre disques à surveiller simultanément pour détecter une cible difficile, a montré que l'attention spatiale se comporte en effet comme un "projecteur clignotant" – même lorsqu'il n'y a qu'un seul endroit où l'attention doit être portée, l'attention semble moduler le traitement de l'information périodiquement, à une fréquence de 7Hz environ. Lorsque plusieurs endroits doivent être attendus simultanément, l'attention passe d'un endroit à l'autre à cette même fréquence. Cette réalisation rend la question de la perception discrète intimement liée à celle de la dynamique de l'attention. C'est donc cette hypothèse qui a guidé tout le travail réalisé au cours de cette thèse : notre perception est-elle basée sur un échantillonnage périodique de l'information visuelle par l'attention?

Dans un premier temps, nous avons cherché à mettre en évidence des périodicités liées à cet échantillonnage sous-jacent, par des études psychophysiques et comportementales.

Inspirés par les résultats de l'illusion de la roue de chariot en lumière continue, nous avons essayé de voir si des artefacts d'aliasing temporel pourraient être observés dans d'autres modalités. Nous avons choisi l'audition, qui est souvent citée comme le sens ayant la meilleure résolution temporelle. Nous nous sommes aperçus pourtant que la résolution temporelle pour la perception de mouvement auditif spatial était sévèrement limitée – au delà de 3hz environ, il devenait impossible de juger la direction de mouvement d'une source sonore. Il se trouve que l'information spatiale en audition repose sur des calculs assez complexes de délais et de niveau sonore entre les deux oreilles, et est donc une information construite. L'information directement à disposition pour le système auditif est la fréquence d'une source sonore. Nous avons donc exploré la possibilité de créer du mouvement périodique dans le domaine fréquentiel. Cependant, les méthodes existantes pour créer ce genre de stimuli (échelle de Shepard ou glissando de Risset) sont également limitées par artefacts proéminents dès que le rythme de présentation dépasse 3 ou 4Hz. Notre échec pour la modalité auditive, ainsi que l'échec rapporté par un autre groupe pour la modalité tactile, nous ont poussé à revenir vers la vision. Nous avons exploré la

performance en fonction de la fréquence de présentation dans un paradigme de color-orientation binding, mais nous sommes heurtés à nouveau à une limite de 3-4Hz au-dessus de laquelle la tâche n'était plus faisable. Finalement, nous sommes revenus à l'étude de la perception du mouvement. Certains stimuli ne peuvent pas être traités par le système automatique de perception du mouvement: si seul le système attentionnel de perception du mouvement est utilisé, les effets de l'aliasing temporel devraient être très visibles. Nous avons choisi un paradigme de mouvement interoculaire, dans lequel chaque oeil ne reçoit pas d'information sur le mouvement mais l'intégration de l'information présentée aux deux yeux définit une direction non ambiguë. La performance moyenne dans cette expérience, dans laquelle 10 sujets furent inclus, tombe aux alentours de 7Hz, sans remonter par la suite comme nous le prévoyions. Le problème rencontré dans toutes ces expériences est que, si la performance tombe trop vite avec la fréquence de présentation, il est impossible d'observer de l'aliasing temporel (qui se manifestait maximale à 10Hz dans le cas de l'illusion de la roue de wagon, par exemple). Cependant, dans le cas du mouvement, nous pouvons utiliser une approche de modélisation pour interpréter les données obtenues : en effet, en supposant qu'il existe deux systèmes de perception du mouvement, l'un recevant l'information de manière continue et l'autre l'échantillonnant, et que ces systèmes contribuent à la performance du sujet de façon additive, nous pouvons dériver certains paramètres comme la contribution du système discret, sa fréquence, etc. Avec toutes les précautions qui doivent être prises avec ce genre d'approche basée sur des modèles, nos résultats sont compatibles avec la théorie formulée pour l'illusion de la roue de wagon – nous trouvons une contribution du système attentionnel de 38% en moyenne dans l'expérience interoculaire, avec une fréquence d'échantillonnage de 13Hz environ. Cette contribution est beaucoup plus importante que dans l'expérience contrôle binoculaire où nous trouvons une contribution de 16% en moyenne – qui correspond bien au fait que la probabilité des inversions est d'environ 15-20% dans l'illusion de la roue de wagon – toujours à une fréquence de 13Hz environ. Cette approche permet donc de mettre en évidence des mécanismes sous-jacents qui sont difficiles à détecter en moyenne. Mais au final, notre quête d'aliasing temporel a été peu fructueuse, pour des raisons pratiques.

Certaines études psychophysiques semblent avoir démontré qu'il était possible de diviser l'attention spatiale, c'est-à-dire de porter son attention simultanément à deux endroits disjoints tout en ignorant ce qui se passe entre ces deux endroits. Mais si l'attention oscillait entre les deux endroits attendus? Un comportement sériel de l'attention nous fournirait un argument fort en faveur de sa périodicité. Nous avons repris un paradigme existant et nous en avons analysé les résultats avec une méthode nous permettant de déterminer si l'attention peut réellement être divisée équitablement entre deux endroits attendus. Il s'agit de présenter 8 formes, dont 4 carrés et 4 cercles, arrangées de façon circulaire. Deux de ces formes sont rouges et les autres sont vertes. Ces formes restent à l'écran pour une durée variable (53.3ms, 106.6ms, 186.6ms ou 213.3ms) puis des lettres apparaissent dans chaque forme, avant d'être masquées au bout de 66ms. Le sujet doit, à la fin de chaque essai, décider si les deux formes rouges étaient les mêmes ou différentes, puis dans un deuxième temps rapporter les lettres qu'il/elle a vues et dont il/elle se souvient. La probabilité de rapporter les lettres est utilisée comme mesure de la quantité d'attention qui était allouée à chaque endroit. Notre approche originale fut de considérer la probabilité de détecter les deux lettres cibles, et la probabilité de n'en détecter aucune, pour pouvoir déterminer d'éventuels déséquilibres d'allocation attentionnelle entre les deux lettres cibles, via un formalisme d'équation de second degré. Nos résultats indiquent que, si l'attention semble en effet pouvoir être distribuée équitablement entre les deux cibles, même disjointes, de façon transiente, elle est biaisée vers l'une des deux cibles au bout d'environ 100ms. Ceci est un argument allant à l'encontre d'une division soutenue de l'attention spatiale – et qui a pour conséquence directe que dans une telle situation, l'attention spatiale devrait avoir un comportement sériel (périodique).

S'il est vrai que le cerveau utilise un échantillonnage attentionnel pour construire notre perception, il s'en cache très bien, et nos efforts pour dévoiler ces mécanismes reçoivent assez peu de succès. Qu'est-ce qui rend notre expérience si fluide?

La réponse est peut-être à chercher du côté de certains troubles de la perception, pour lesquels la perception semble être décomposée en une série d'instantanés. Le cas le plus célèbre est celui de la patiente L.M. qui, après une destruction bilatérale de ses aires corticales de perception du mouvement, perçoit le monde comme s'il était illuminé par un stroboscope. Elle dit ne pas pouvoir remplir une tasse de thé, car elle voit l'eau à un certain niveau puis, le moment suivant, la tasse a débordé sans qu'elle ne s'en aperçoive. Elle ne peut pas non plus traverser la rue, car des voitures au loin se retrouvent pratiquement à l'écraser l'instant d'après. Pourquoi sa perception est-elle ainsi constituée d'une succession d'instantanés, chacun clairement défini? Bien que la fréquence de rafraîchissement de sa perception n'ait pas été mesurée directement, ses troubles sont compatibles avec l'hypothèse d'une perception normalement discrétisée et rendue fluide par les mécanismes de perception du mouvement. Un autre désordre perceptuel, également lié au mouvement, atteint certaines personnes qui utilisent le LSD (diéthylamide de l'acide lysergique), mais aussi certains patients traités avec de hautes doses d'antidépresseurs (par exemple, néfazodone, trazodone, rispéridone, mirtazapine) ou de drogues antiépileptiques (topiramate), et peut-être certains patients migraineux. Il s'agit de la perception d'une série de répliques d'un objet en mouvement, qui suivent cet objet, un peu comme si de multiples photos étaient prises le long de la trajectoire et restaient chacune visible assez longtemps pour que plusieurs soient perceptibles simultanément. Ce phénomène peut être expliqué de plusieurs façons; une des hypothèses est celle d'un échantillonnage sous jacent qui ne serait plus masqué et deviendrait perceptible. Nous avons réalisé une enquête en ligne, visant une population ayant pris du LSD dans le passé. Cette étude semble indiquer que la période d'échantillonnage perceptuel serait dans les 75-125ms en moyenne.

Dans un deuxième temps, nous nous intéressons directement à l'activité oscillatoire enregistrable en EEG et cherchons les rythmes qui pourraient être à l'origine de l'échantillonnage attentionnel.

Nous revenons sur le problème de la détection d'un stimulus visuel si faible qu'il n'est perçu par le sujet qu'une fois sur deux en moyenne. Pourquoi y'a-t-il des essais dans lesquels on perçoit le flash, et d'autres essais dans lesquels on le rate complètement? Puisque le stimulus est exactement identique d'un essai à l'autre, il faut bien chercher la source de variabilité ailleurs. L'hypothèse que la perception est modulée périodiquement peut être testée directement si l'on a accès à l'activité cérébrale : il suffit de trouver un rythme cérébral spontané dont la phase au moment de la présentation du stimulus peut prédire si celui-ci sera perçu ou non. Nous avons exploré l'ensemble des rythmes cérébraux détectables en EEG et avons trouvé une influence de la phase d'une oscillation fronto-centrale, à 7hz, sur la détection du stimulus. Ceci est très compatible avec un rythme attentionnel, tant en termes de topographie (la source de ce rythme pourrait être au niveau des champs oculomoteurs frontaux) qu'en termes de fréquence (correspond à la fréquence du "projecteur clignotant" de l'attention mis en évidence précédemment). L'attention étant toujours dirigée à l'endroit où le flash devait apparaître, nous ne pouvions pas avec cette étude seule incriminer l'attention de manière certaine. Cependant, une étude réalisée par la suite, manipulant explicitement l'attention, a montré que l'influence de la phase de ce rythme fronto-central sur la détection d'un flash faible était maximale quand l'endroit où le flash devait apparaître était attendu, confirmant donc cette interprétation.

Revenons maintenant à la question de la perception de simultanéité, car il s'agit d'un phénomène très discuté dans le débat opposant la théorie des moments perceptuels à une simple intégration temporelle. Lorsque deux flashes visuels sont présentés successivement, il existe un délai où le sujet les perçoit comme étant simultanés dans la moitié des essais, et successifs dans l'autre moitié. La théorie des moments perceptuels part du principe qu'un rythme cérébral organise la perception en une série d'échantillons. Pour un délai donné entre les deux flashes, en fonction de la phase du rythme en question, les flashes peuvent tomber dans le même échantillon ou dans deux échantillons successifs, et c'est cela qui

déterminerait le percept. Peut-on donc, de la même façon que dans l'expérience décrite précédemment, trouver un rythme dont la phase détermine la perception? Il y a 30 ans, Francisco Varela avait trouvé que la phase de l'alpha occipital prédisait la perception de simultanéité, mais nos efforts ne permirent pas de confirmer ce résultat statistiquement. Nous avons observé quelques complications par rapport à l'hypothèse originale : par exemple, une présentation de deux flashes au sein d'un hémichamp visuel donne lieu à une perception plus fine de leur relation temporelle que si les deux flashes sont présentés dans des hémichamps visuels opposés, ce qui suggère que la perception de simultanéité n'est pas entièrement déterminée par des moments perceptuels globaux. Nous avons observé quelques tendances dans nos résultats qui vont dans le sens des observations de Varela, et peut-être une nouvelle expérience avec plus d'essais et quelques améliorations du protocole, par exemple en termes de clarté du percept, nous permettrons de redonner vie à son résultat – et du même coup, d'offrir une démonstration difficilement discutable du découpage temporel de la perception par des oscillations cérébrales.

Enfin, nous avons voulu montrer empiriquement que l'attention oscille entre plusieurs endroits attendus – suivant notre démonstration psychophysique que l'attention ne peut pas être divisée de manière soutenue. Bien que certains résultats d'électroencéphalographie aient prétendu avoir démontré que l'attention pouvait être divisée spatialement, la possibilité d'une alternance rapide entre les différents endroits attendus n'a jamais été falsifiée – elle a juste été rejetée sur la base d'arguments peu convaincants. Il est nécessaire de pouvoir trouver une signature cérébrale quasi-instantanée de la position de l'attention spatiale, pour pouvoir la suivre avec une bonne résolution temporelle et enfin savoir si, dans une situation d'attention divisée, elle est effectivement divisée ou se déplace périodiquement entre les différents endroits. C'est dans cette optique que nous avons développé une nouvelle technique, s'inspirant du marquage fréquentiel (frequency tagging). Le marquage fréquentiel consiste à présenter un patch lumineux clignotant à une certaine fréquence sur un écran. Cela entraîne une oscillation cérébrale, généralement recueillie au niveau des électrodes occipitales (cortex visuel), à la même fréquence. L'amplitude de l'oscillation cérébrale entraînée est modulée par l'attention spatiale. Pour avoir un rapport signal sur bruit suffisant, il est généralement nécessaire d'utiliser plusieurs cycles pour la détermination de l'amplitude de la réponse, ce qui ne permet pas d'étudier la dynamique temporelle fine de l'attention. Notre idée consistait donc à entraîner l'activité cérébrale avec un stimulus contenant de l'information dans plusieurs bandes de fréquence, et à utiliser à chaque instant la signature fréquentielle de la réponse cérébrale pour savoir où se trouvait l'attention – nous avons utilisé deux disques dont la luminance était modulée avec un spectre plat entre 0 et 80Hz, et bien sûr complètement indépendants l'un de l'autre, pour entraîner des oscillations cérébrales. Les résultats semblent indiquer que l'attention spatiale, lorsqu'elle est portée à deux endroits simultanément, échantillonne en fait les deux endroits chacun à son tour – à un rythme d'un échantillon toutes les 120-140ms, correspondant bien au rythme du projecteur clignotant précédemment estimé.

Ce travail de thèse a mené à la publication de trois articles de recherche, et de deux articles de revue de la littérature et de méthodologie. Nous y avons notamment démontré l'importance de s'intéresser à la phase des oscillations cérébrales, une approche plutôt absente de la littérature; la plupart des scientifiques s'intéressent surtout à l'amplitude des oscillations – qui est aussi une approche tout à fait justifiable et importante pour de nombreuses questions. Étudier la phase est une approche complémentaire qui peut élucider certains mécanismes avec une résolution temporelle plus fine.

Il semble que l'attention est bien une ressource périodique, avec une fréquence caractéristique à 7Hz, et que la perception pourrait être construite sur la base d'une série d'échantillons attentionnels. Nous proposons pour finir un modèle théorique de la façon dont ce rythme attentionnel s'intègre avec d'autres résultats de la littérature, comme les changements de distribution de l'amplitude du rythme alpha dans les lobes occipitaux avec l'attention, par exemple. Nous

insistons sur le fait qu'un comportement périodique de l'attention couverte semble tout à fait défendable d'un point de vue évolutif. En effet, de nombreux rythmes sont en place au niveau des capteurs d'information: les mouvements des vibrisses des rats lors de l'exploration tactile, les mouvements oculaires (saccades) lors de l'exploration visuelle, les reniflements répétitifs lors de l'exploration olfactive... l'attention spatiale couverte étant très intimement liée aux mouvements oculaires, avec des circuits neuronaux partagés, il n'est pas du tout surprenant d'y trouver une rythmicité également. Il est possible que le rythme soit plus rapide pour l'attention couverte que pour les mouvements oculaires (attention ouverte), puisque les mouvements de l'attention couverte sont moins coûteux métaboliquement. Notre modèle propose un rôle potentiel de ce rythme attentionnel.

Table of Contents

I. INTRODUCTION	1
A. FROM THE PROJECTION ON THE RETINA TO OUR PERCEPTION OF THE WORLD : A LEAP OF FAITH	3
1. <i>The image on the retina: a far cry from what we perceive</i>	3
a) Photoreceptors are unevenly distributed	3
b) The eyes keep moving	3
c) The blind spot	4
2. <i>The visual brain 101 : a hierarchy of visual processing areas organized in two main pathways</i>	5
a) A hierarchy	5
b) The main stages of visual processing	6
c) Two pathways	7
3. <i>Perception is a constructive act (a.k.a. a con job)</i>	8
a) Perception is fallible	8
b) Resolving ambiguity	9
B. THE “MIND’S EYE” : HOW ATTENTION FILTERS INCOMING VISUAL INFORMATION	11
1. <i>Too Much Information?</i>	11
2. <i>What is attention? come on, everybody knows</i>	12
a) Bottom-up and Top-down, a.k.a. exogenous and endogenous attention	13
b) Spatial, feature-based, object-based, what have you	14
3. <i>The neural basis of attention</i>	15
a) What areas of the brain guide attention allocation?	15
b) The effects of attention on neural activity	16
C. NEURAL OSCILLATIONS AND THEIR FUNCTIONAL SIGNIFICANCE	19
1. <i>Oscillations in the EEG</i>	19
2. <i>How do oscillations come about?</i>	21
3. <i>How can oscillations serve brain function?</i>	21
a) A reference signal for a spike timing neural code	21
b) Nested oscillation coding	22
c) Enhanced communication through phase synchrony	24
d) A communication protocol	24
D. SUMMARY AND AIM OF MY THESIS WORK	26
II. BEHAVIORAL EVIDENCE OF RHYTHMS IN ATTENTION AND PERCEPTION	27
A. TWO CLASSICAL DEBATES IN THE LITERATURE	28
1. <i>The discrete perception debate</i>	28
a) Studies of simultaneity and temporal rate : discrete epochs, or integration period?	28
b) Some studies undoubtedly point to periodicities in perceptual processes	30
2. <i>The sequential attention debate</i>	31

a)	Many classical paradigms fail to provide a definite answer	31
b)	The blinking spotlight of attention	33
B.	TEMPORAL ALIASING ARTEFACTS AS EVIDENCE THAT PERCEPTION RELIES ON PERIODIC PROCESSES	34
1.	<i>Aliasing in the continuous Wagon Wheel Illusion</i>	34
a)	A definition of temporal aliasing	34
b)	The continuous wagon wheel illusion (cWWI): a case of temporal aliasing?	35
c)	How to probe the brain for temporal aliasing artefacts	40
2.	<i>Seeking aliasing artifacts in other sensory modalities – the case of auditory motion perception</i>	40
a)	Auditory motion in the spatial domain	41
b)	Auditory motion in the frequency domain	42
3.	<i>Does attentive (third order) visual motion perception rely on a sampling process?</i>	44
a)	Interocular contrast-defined motion	45
b)	Other third order motion stimuli	46
4.	<i>A model fitting approach to detect aliasing in psychometric curves</i>	46
5.	<i>Summary and Discussion</i>	49
C.	PSYCHOPHYSICAL EVIDENCE AGAINST A SUSTAINED SPLIT OF THE ATTENTIONAL SPOTLIGHT	52
1.	<i>Can the spotlight of attention be split?</i>	52
a)	Existing experimental evidence is inconclusive	52
b)	The Bichot et al 1999 study	53
2.	<i>PAPER 1 : Dubois, Hamker and VanRullen (2009) – Journal of Vision</i>	56
3.	<i>Summary and Discussion</i>	68
D.	PATHOLOGICAL BEHAVIORAL EVIDENCE FOR RHYTHMS IN PERCEPTUAL PROCESSES	70
1.	<i>Akinetopsia : a rare and striking disturbance</i>	70
2.	<i>A more common, poorly studied disorder : Visual trails</i>	71
a)	PAPER 2 : Dubois, J. and VanRullen, R. (2011) – PLoS Biology	71
b)	An online survey among past LSD users	76
3.	<i>Summary and Discussion</i>	83
E.	PAPER 3 : VANRULLEN AND DUBOIS (2011) – FRONTIERS IN PSYCHOLOGY	85
III.	EEG EVIDENCE OF RHYTHMS IN ATTENTION AND PERCEPTION	97
A.	METHODS FOR LINKING ONGOING EEG PHASE TO TRIAL-BY-TRIAL VARIABILITY	99
B.	THE RHYTHM OF ATTENTIONAL SAMPLING FOR VISUAL DETECTION	103
1.	<i>Phase of ongoing oscillations and perception</i>	103
2.	<i>PAPER 4 : Busch, Dubois and VanRullen (2009) – The Journal of Neuroscience</i>	104
3.	<i>Summary and Discussion</i>	117
C.	A BRAIN RHYTHM FOR PERCEPTUAL FRAMING?	118
1.	<i>“Perceptual Framing and Cortical Alpha Rhythm”</i>	118
2.	<i>Replication attempt</i>	122
a)	Stimulus, procedure and EEG acquisition	122
b)	Behavioral results	123

c)	Correlation between simultaneity judgement and phase of alpha at electrode O1?	124
d)	Correlation between perception of simultaneity and phase of any rhythm at any electrode?	127
3.	<i>Summary and Discussion</i>	130
D.	DOES THE SPOTLIGHT OSCILLATE BETWEEN CONCURRENTLY ATTENDED LOCATIONS?	131
1.	<i>A sustained division of the attentional spotlight has not been convincingly shown yet</i>	131
2.	<i>Our methodological approach : from SSVEPs to broadband frequency tagging</i>	134
a)	How the brain responds to different stimulation frequencies.....	134
b)	Attentional modulation of SSVEPs	135
c)	Increasing the temporal resolution of the SSVEP technique.....	136
3.	<i>Where is the spotlight of attention now?</i>	137
a)	Stimuli, procedure and EEG acquisition	137
b)	Occipito-parietal alpha indexes attentional allocation	138
c)	The EEG phase locks to stimuli (mostly at lower frequencies).....	140
d)	Attention increases phase locking	141
e)	The phase locking information available on occipital electrodes can be used to decode attentional allocation (with modest performance).....	142
f)	What is the temporal resolution of our technique?	145
g)	The ~7Hz blinking spotlight of attention oscillates between simultaneously attended locations	145
4.	<i>Summary and discussion</i>	147
E.	PAPER 5 : VANRULLEN, BUSCH, DREWES AND DUBOIS (2011) – FRONTIERS IN PSYCHOLOGY.....	150
IV.	DISCUSSION	161
A.	WHY ARE OSCILLATORY MECHANISMS NOT OBVIOUS PERCEPTUALLY?.....	163
1.	<i>An absence of perception cannot be perceived</i>	163
2.	<i>Early areas buffer incoming visual information until the next sample</i>	164
3.	<i>Motion perception mechanisms are the glue that links discrete epochs</i>	164
4.	<i>Resetting the rhythm may be the only way to find a behavioral effect</i>	165
B.	THE RHYTHM OF COVERT ATTENTION	166
1.	<i>7hz? 10hz?13hz?</i>	166
2.	<i>Other estimates of the attentional rhythm in the literature</i>	168
a)	The speed of attentional shifts	168
b)	Serial search	168
3.	<i>Task dependence</i>	169
a)	When time slows down.....	169
b)	A rhythmic and a continuous mode	169
4.	<i>Entraining the rhythm of covert attention</i>	170
5.	<i>Covert rhythmic sampling and overt rhythmic sampling: the egg and the chicken?</i>	171
C.	PUTTING IT ALL TOGETHER : COULD PERCEPTION RELY ON A 7HZ ATTENTIONAL SAMPLING OF INFORMATION NESTED IN OCCIPITO-PARIETAL ALPHA OSCILLATIONS?	174
1.	<i>The pulsed inhibition account of attentional selection</i>	174

2. <i>Attentional selection through pulsed inhibition works if downstream areas integrate information over time</i>	176
3. <i>The 7Hz attentional rhythm as a readout rhythm</i>	177
EPILOGUE	179
REFERENCES	181

I. Introduction

In this thesis, I pry open the black box of visual perception with psychophysical and electrophysiological tools. The aim is to find evidence of a purported periodic updating mechanism – of which we are completely unaware in our daily lives.

As humans we take vision for granted. We open our eyes to experience brightness, color, texture, depth, and motion. It is often said that “seeing is believing”; for most of us, there is no reality that can supersede what we see. But for the scholar interested in vision science, understanding the details of how we see is a definite eye opener... In this introduction, I start by describing the data that the brain receives as input. Our eyes are very imperfect sensors; they do not faithfully encode all the information that is available in their field of view. I will give a few examples of what information is lost at this stage of the seeing process. The brain constantly needs to interpret the input, using heuristics and computations of which we are completely unaware. Does one know that information is actually segregated and processed along two pathways? Of course not, there is no conscious access to the underlying operations. The brain may not always be right... as vision scientists we take pleasure in tricking it with visual illusions. These illusions serve as tools to uncover some of the heuristics the brain implements. The intent of this rather generic “vision 101” section in the scope of my thesis is to plant an important idea in the reader’s mind: one cannot trust introspection in the study of perceptual processes, much to our philosopher friends’ dismay. It is important that doubt about the ways of the brain settle into the reader’s mind, and that curiosity about the underpinnings of perception be aroused.

Somewhere, among the silent and intricate computations performed by the brain for perception, lies attention. Attention is a key player in this thesis, therefore it is important that the reader become well acquainted with the word attention as it is defined by cognitive scientists. The layman understands attention as a synonym of focus, an effort to channel one’s thoughts to a given task (“Pay attention to what I am saying!”); the cognitive scientist uses the word attention as a generic term to describe any imbalances in the processing of incoming information – many of these imbalances being involuntary. I cover the different forms under which attention has been described in the cognitive literature. This should be a good primer for the new student of attention, before they go browsing through the thick and ever thickening literature on the subject. Once I have laid out the taxonomy, I delve deeper into the workings of a specific form of attention: spatial or focal attention, which is the form that this thesis most explicitly addresses.

The computations that the brain performs for perception involve large-scale interactions between neuronal populations. A prominent feature of brain activity in electrical recordings are oscillations in various frequency bands, which involve large neuronal populations. In recent years, it has become clear that these oscillations are critical to brain function, hence they are of interest for us who want to uncover the hidden operations that the brain performs for perception. The next section of the introduction is thus devoted to giving the reader a general introduction to brain rhythms.

The functional importance of oscillations for brain computations naturally leads to the obvious question, which has been raised a few times without ever receiving enough empirical support to become a mainstream idea: does visual perception have an intrinsic rhythm? does it carry evidence of the underlying oscillatory neural activity on which it relies? There is in fact already some experimental data suggesting that attention may be a (quasi-) periodic process, and that perception is based on discrete samples, rather than millisecond-by-millisecond online updating. I will unveil the existing evidence

throughout this thesis and, most importantly, will present my recent work which was guided by an “astonishing hypothesis” : that our seemingly continuous perception of the world may be a mere illusion that the brain constructs from discrete samples captured by the strobe of attention.

A. From the projection on the retina to our perception of the world : a leap of faith

1. The image on the retina: a far cry from what we perceive

Light passes through the cornea and lens of our eyes. These focus an upside-down image of the world onto the retina, a thin layer of neurons and photoreceptors at the back of the eyeballs. The signal sensed and transduced by the retina is the input to the visual system. In this section, we briefly describe how distorted and incomplete this input is.

a) Photoreceptors are unevenly distributed

There are two classes of photoreceptors in the retina: about one hundred and twenty million rods and five million cones. Rods mediate vision in low light conditions: they are more sensitive and more sluggish than cones, so their photopigments are mostly bleached during normal daylight vision. Cones are thus responsible for vision most of the time for those of us who are active during the day and sleep at night. Cones are most densely packed at the fovea, the part of the retina that receives light from the direction that we are directly looking at. The density of cones falls off dramatically as one goes from the fovea to the periphery (Figure 1). In this respect, the retina is very different from the chip of a digital camera, which is tiled evenly by sensors (one per pixel). The distribution of sensors on the retina is such that the central part has very high resolution, while the periphery has very low resolution. Yet, we subjectively perceive our whole field of view as being crisp, and certainly not blurry in the periphery.

Color perception relies on three types of cones, which are tuned to different wavelengths: the short- (S), middle- (M) and long-wavelength (L) cones. Color arises from the comparison of the number of photons absorbed by each cone type. Evolutionary flaw or genius feature, the three classes of cones are not distributed evenly either. The S cones are absent from the central part of the fovea; throughout the retina, patches containing mostly M cones are interspersed with patches containing mostly L cones. Yet, if we look at an evenly colored surface, it does not appear patchy.

b) The eyes keep moving

Due to the uneven distribution of photoreceptors, we constantly need to move our eyes to bring the objects of interest onto our fovea. When both eyes move together rapidly, vision scientists talk about “saccades”. You are currently reading this text by skipping across it with a series of small saccades. This is also true when we look at a photograph, or a landscape, someone’s face, ... In fact, in between two saccades, the eyes are not completely still: they undergo low amplitude fixational eye movements (classified as tremor, drift and microsaccades). It is impossible to keep our eyes completely still. They saccade about as often as our heart beats, which means that the image projected onto the retina

changes quite abruptly every second or so. Yet, our perception of the world is not jerky – an active mechanism, saccadic suppression, is implemented by the brain to prevent motion transients.

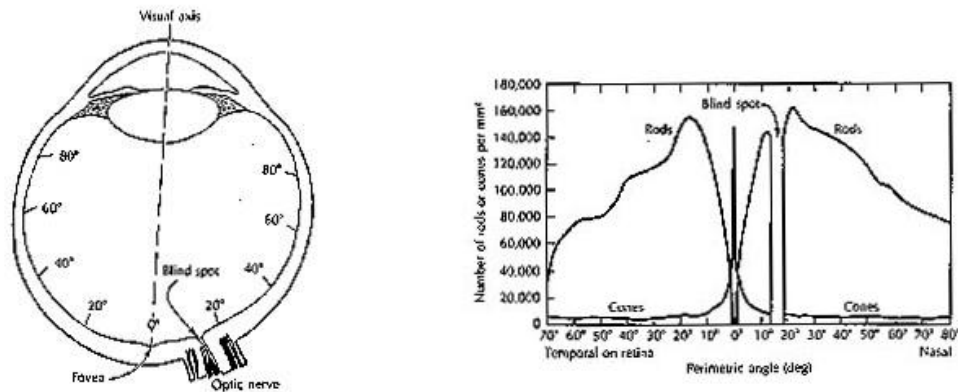


Figure 1 The uneven distribution of photoreceptors on the human retina. On the left, a schematic cross-section of the eye. The angle relative to the point of sharpest seeing (the fovea) is referred to as eccentricity. On the right, the densities of photoreceptors as a function of eccentricity are plotted. Reproduced from (Christof Koch, 2004)

c) *The blind spot*

The retina is a layered structure, containing many cell types (horizontal, bipolar, amacrine and ganglion cells are the main ones), which convert the optical signal sensed by photoreceptors into an electrical signal. The final output is relayed to the brain by the ganglion cells, whose axons are bundled together and form the optic nerve. The optic nerve actually has to go through the layer of photoreceptors (Figure 1 left), because the photoreceptors lay at the very back of the retina (yes, light has to go through the mess of retinal neurons before it reaches photoreceptors). Hence, there is a hole in the retina. You can “see” it : refer to the instructions in the caption of Figure 2 and experience it for yourself. Yet, you are not aware of your blind spots in everyday perception.

Laurent Itti, when he was a postdoc at Caltech in Christof Koch’s laboratory, created a movie that simulates the output of the retina when looking at a picture. The movie is available online (<http://www.klab.caltech.edu/~jdubois/demo/Retinal-Simulation-Itti.mpg>) and is a great way to visualize the few points that we just made. The visual input to the brain is indeed a far cry from what we perceive consciously.



Figure 2 Close your right eye. Hold the image about 20 inches away. With your left eye, look at the +. Slowly bring the image closer while looking at the +. At a certain distance, the dot will disappear from sight...this is when the dot falls on the blind spot of your retina. Reverse the process. Close your left eye and look at the dot with your right eye. Move the image slowly closer to you and the + should disappear.

2. The visual brain 101 : a hierarchy of visual processing areas organized in two main pathways

The optic nerves feed incoming visual information to the visual areas of the brain; from the retina, most (~90%) of the information flows to the lateral geniculate nucleus (LGN) of the thalamus then on through the optic radiations to the primary visual cortex (V1) and finally to higher visual areas (Figure 3).

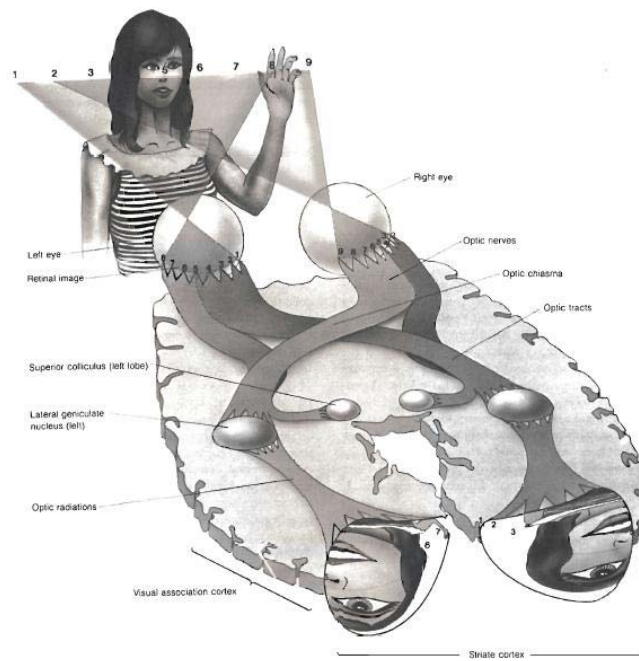


Figure 3 Neural pathways from the eye to visual cortex. Reproduced from (S. E. Palmer, 1999)

The output of the retina is analyzed and interpreted, contours are detected, motion is computed, objects are segmented from the background and recognized. In this section we give a rough overview of what is known about how the brain processes incoming visual information.

a) A hierarchy

Pretty much any textbook on visual processing features a figure from the highly cited paper Felleman and VanEssen published in 1991 (Felleman & D C Van Essen, 1991), and it is an important picture to have available when thinking about visual processing. A slightly simpler version was published by Maunsell and Newsome in 1987 (J. H. Maunsell & W T Newsome, 1987)(Figure 4). Most of our knowledge of the connectivity of visual cortex is based on macaque anatomy; however there is much evidence that human anatomy is quite similar.

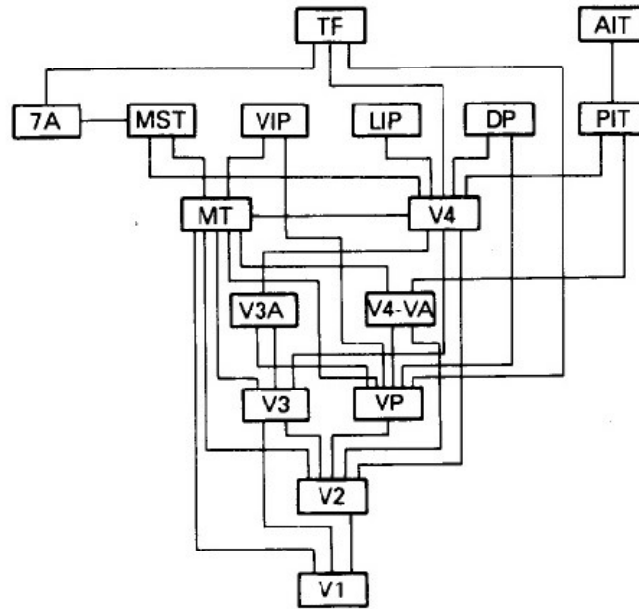


Figure 4 Organizational chart of the monkey's visual system (cortical). Reproduced from (J. H. Maunsell & W T Newsome, 1987)

The cortex is a sheet-like structure with six major anatomically defined layers (and several more sublayers defined physiologically). Two major kinds of connections are found between different areas of the cortex: on the one hand, connections that originate from neurons in the superficial layers and terminate majoritarily onto neurons in layer 4 (called forward projections, by analogy with the connections from the LGN to V1 which terminate in layer 4 of V1); on the other hand, connections that originate and terminate outside of layer 4 (called backward, or feedback projections). The hierarchy depicted in Figure 4 is based on this distinction between forward and feedback connections : if an area receives forward projections from another area, it is located one level above in the hierarchy.

Does this hierarchy reflect something about function? As one goes from the lower to the higher levels of the hierarchy, individual neurons seem to respond to input from larger retinal areas; in other words, the receptive field size increases as one moves up in the hierarchy (for macaque monkeys the size of the receptive fields increase from about 0.1-0.5 degrees of visual angle in V1 to 0.5-1 degrees in V2 to 1-4 degrees in V4 to more than 25 degrees in IT (Robert Desimone, Moran, & Spitzer, 1988)).

b) The main stages of visual processing

Retinal ganglion cells come in two major classes, the midget cells (a.k.a. P ganglion cells) and the parasol cells (a.k.a. M ganglion cells). P ganglion cells are more sensitive to color than black and white (they receive input just from cones), and the reverse is true of M ganglion cells (which receive input from both cones and rods).

The LGN consists of 6 layers; axons from M ganglion cells synapse onto neurons in the lower two layers (magnocellular layers), while axons from P ganglion cells synapse onto neurons in the upper four layers (parvocellular layers). Clearly, the

selectivity of the two classes of ganglion cells carries on to the LGN cells. The layers alternate between receiving input from the left and the right eyes; there are no neurons receiving input from both eyes at this stage yet. In each layer of the LGN, the geometrical relationships between cells are qualitatively the same as the geometrical relationships between the ganglion cells they receive their input from (retinotopic mapping). The connections between the retinae and the LGN are organized so that the LGN in the left hemisphere receives input from the right hemifield of view, while the LGN in the right hemisphere receives input from the left hemifield of view (Figure 3).

The LGN projects forward to the primary visual cortex. Layer 4 neurons receive most of the visual input from the LGN : sublamina 4C α receives most magnocellular input (which goes on to layer 4B, where cells respond to input from both eyes and are selective to direction of motion); sublamina 4C β receives input from parvocellular layers. Studies of the architecture of the primary visual cortex show that, in each hemisphere, there is a retinotopic map of the contralateral hemifield of view. The map is distorted, with an overrepresentation of the central part of the visual field. The inputs from the two eyes are still segregated but are organized in a semi-orderly fashion into ocular dominance columns; cells in V1 are selective to orientation of lines, and in a given column, all cells are selective to the same orientation. Cells in neighboring columns are selective to similar orientations. V1 is thus organized into hypercolumns, which cover a surface of about 1mm² and within which all orientations are cycled through. Not all cells are orientation selective, however; hypercolumns are further separated into blobs and interblobs, with neurons in blobs selective to wavelength and featuring center-surround, sometimes color opponent receptive fields. Color and form processing are thus seemingly segregated in V1, and they are segregated from depth and motion processing (the magnocellular pathway).

The segregation continues as information goes up to the next level of the hierarchy. V1's layer 4B neurons project to neurons in V2's thick stripes and directly to MT (the middle temporal area) which mediates motion perception. V1's blob neurons project to V2's thin stripes and V1's interblob neurons project to V2's interstripes.

From V2's thin stripes, information goes on to V4, which mediates color perception; from V2's interstripes, information eventually goes on to IT (inferotemporal cortex) which mediates object recognition. From V2's thick stripes, information about depth is routed to V3, MT and on to parietal areas.

At this point, I owe an apology to the vision scientists among us who are trying to achieve a very detailed understanding of what goes on in early areas of the visual cortex. The description I offer is oversimplified, maybe even partially wrong. In practice nothing is as clear cut as I present it, but I try to stick to the big picture without getting involved in the numerous debates that specialists fight over in this area of research.

c) Two pathways

Lesion studies in monkeys performed by Ungerleider and Mishkin (Ungerleider & Mishkin, 1982) showed that inferior temporal areas are involved in identifying objects (ventral pathway, or "what" pathway) while parietal areas are involved in locating objects (dorsal pathway, or "where" pathway). The extent to which this distinction holds and the question of how information eventually comes together are the subject of some controversies, but it is a useful coarse description to keep in mind when you think about vision.

3. Perception is a constructive act (a.k.a. a con job)

We showed that the input to the visual processing areas of the brain is heavily distorted and incomplete, and that information is further distorted and segregated as it is processed. Moreover, the problem of reconstructing the 3D scene from a 2D projection (actually, two 2D projections, one in each eye) is inherently underconstrained. Our brain constantly needs to interpret the input, and it settles on one particular interpretation – what we perceive. This interpretation may not always be the most accurate representation of reality, and I give you a brief taste for this here.

a) *Perception is fallible*

We vision scientists love to fool the brain. Beyond the ludic aspect of our job, each new way of tricking the brain teaches us something about the mechanisms at work in constructing conscious perception. Figure 5 shows four classical visual illusions, in which seemingly obvious properties are incorrectly perceived.




Figure 5 Some visual illusions demonstrating how the brain's heuristics can be fooled by clever displays. **a** Kaniska triangle: the white triangle, pointing down, is perceived to have well defined edges, but they are completely illusory. **b** Hermann grid: gray dots are perceived at the intersections of the white lines, but they do not physically exist. **c** White's illusion: the two gray bars have the exact same luminance, but the one of the left is perceived as brighter. **d** Ponzo illusion: the two horizontal lines have the exact same length, but the top one is perceived as longer.

Another demonstration of the failure of perception to faithfully represent visual input is found in the realm of afterimages. At the Vision Science Society meeting held in Naples, FL in May 2009, Daw-An Wu and Patrick Cavanagh demonstrated how the use of a very powerful flash (a.k.a. the "monster flash") can trigger a positive afterimage that can take a few tens of seconds to fade away. Less intense afterimages and other perceptual aftereffects can be demonstrated quite easily – Figure 6 is a striking example of a color aftereffect (negative afterimage). Afterimages and aftereffects are examples of nonveridical perception – you perceive something that isn't there. What you see is not always what you get.




Figure 6 Stare at the dot in the image on the left for approximately 30s, trying to keep your eyes as still as possible. Then, stare at the corresponding dot in the image on the right. The black and white photograph is perceived as though it were a color photograph!

b) Resolving ambiguity

Ambiguous figures, such as the Necker cube (Figure 7a) and other stimuli designed by vision scientists (I find the animated dancer, available at <http://www.klab.caltech.edu/~jdubois/demo/dancer.gif>, to be the most compelling ambiguous percept), have competing interpretations that are equally likely. In the face of such ambiguity, the brain makes a choice, and we are consciously perceiving only one of the alternatives at a time. Following intrinsic fluctuations, the interpretation chosen by the brain may switch between the different possibilities.




Figure 7 Some ambiguous stimuli used by vision scientists **a.** Necker Cube : the cube can be seen in two different orientations **b.** Binocular rivalry : you need to fuse the two gratings, by crossing your eyes – start by fusing the dots above the gratings, crossing your eyes in such a way that you can see three dots; then, slowly move your eyes down and look at the central image. You should see the horizontal and vertical grating alternatively taking over.

A related phenomenon is binocular rivalry. This is a phenomenon encountered when the inputs to the two eyes do not match; for instance, one eye is presented with a picture of a vertical grating while the other eye is presented with a picture of a horizontal grating (Figure 7b). The resulting percept is usually either a vertical grating or a horizontal grating, but not a mixture of the two (except in some cases when piecemeal rivalry occurs). The brain favors the input from one of the eyes at any given time.

It is important to understand that the mechanisms evidenced by such illusions are tremendously useful in ecologically valid situations – meaning, in the real world. In the case of the Ponzo illusion (Figure 5d), the real world interpretation is of a straight road vanishing in the distance; if two lines drawn on the road appeared like they do on the drawing, the farther one would indeed need to be longer than the closer one... So, the point is not that perceptual mechanisms are flawed; rather, it is that perception is not an accurate description of the visual input (which is a good thing after all, given what the visual input consists of!).

It is evident that the brain performs a wealth of transformations and interpretations, which we do not consciously experience. We do not have access to the process of perceiving, and are just presented with the end result instantaneously; introspection does not appear then to have any value whatsoever as a tool to understand how visual perception works. This is the take home message of this section.

B. The “mind’s eye” : how attention filters incoming visual information

Back in January 2011, my wife and I were hosting a friend of ours and her 2-year-old son in our apartment; as I was just about to leave and drive to work, my wife asked me to tell something to the kid (who is usually obedient when I tell him to do or not do something). Our friend was in the bathroom giving him a bath, and I walked right in, looked straight at the kid and talked to him for a minute or so. Then I turned around, on my way out of the apartment. I heard our friend and my wife laughing, and when I asked what the matter was, they said that while I was talking to the boy, our friend was standing naked right next to him. Believe it or not, I did not notice it, and she is a good looking woman. This should have been an extremely salient visual stimulus – even though I did recently marry the most beautiful woman.

The truth is, we may think that we see everything, but we don’t; this is the lesson taught by my little anecdote, or in more controlled conditions by the “inattentive blindness” paradigm used by Mack and Rock (Mack & Rock, 1998). Though I am now a fairly well trained vision scientist, I still tend to forget it. Another striking example is the “change blindness” paradigm, which is demonstrated in any Cognitive Science 101 course. Change blindness is the failure to detect changes in visual scenes when two versions of a scene are presented in a sequence separated by a brief disruption (Rensink, J. K. O. Regan, & Clark, 1997). The phenomenon is really easy to demonstrate; it can take minutes to detect the change – and observers who finally see it never fail to take their head between their hands and say to themselves, or out loud: “how could I miss THAT!”.

The importance of attention in shaping perception becomes all the clearer in patients with parietal lesions. Patients suffering from “hemineglect” have perfectly normal visual abilities, yet they often fail to notice anything in one hemifield – proof that without attention, perception is severely limited. I will not enter the debate of whether attention is strictly necessary for conscious perception or not – attention and consciousness have been shown to be different processes, however they definitely are tightly linked (Christof Koch & Tsuchiya, 2007).

1. Too Much Information?

With more than one million axons bundled in each optic nerve, the output of the retinae is of the order of a few megabytes per second. This is more information than can be processed by the visual system; at least, this is what most vision scientists agree upon, and you will find this statement in the introduction of pretty much any research paper studying attention (Figure 8). The axiom of attention research, if you will.

Broadbent (Broadbent, 1958) was the first to describe our brain as a limited-capacity information processing system; he based his claim on empirical evidence from selective hearing research, not vision, but it is not a far stretch to generalize it.

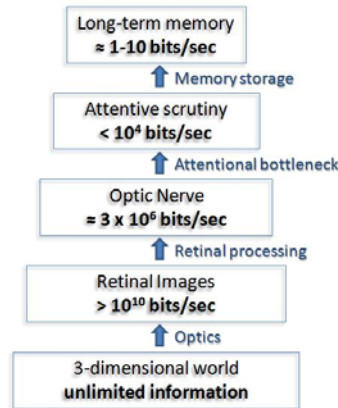


Figure 8 Information pyramid for the visual system. The estimates of information available in retinal images, encoded in the optic nerves and passed through the window of attention are based on information rates of 3 bits/s for each neuron. Reproduced from (C. H. Anderson, David C Van Essen, & Olshausen, 2005)

In science, it is not enough to make qualitative statements, they need to be supported by quantitative data. Tsotsos, “frustrated” by the pervasive vague discussions of bottlenecks and other resource limits, attempts to quantify the capacity limits of vision, within the theoretical framework of computational complexity. He defines the problem of vision as :

given a sequence of images for each pixel, determine whether it belongs to some particular object or other spatial construct, localize all objects in space, detect and localize all events in time, determine the identity of all the objects and events in the sequence, and relate all objects and events to the available world knowledge (Tsotsos, 2005)

He shows this problem to be decidable (a solution exists, hence it can be modeled computationally) but intractable (the problem can not be solved in a reasonable time), unless task and domain guidance is permitted (Tsotsos, 1990). The brain must thus perform approximations and attentional selection to reach an acceptable solution to the problem of vision.

2. What is attention? come on, everybody knows...

A computational definition of attention is thus the mechanisms that reduce the amount of information to be processed by the brain. Anderson, Van Essen and Olshausen (C. H. Anderson et al., 2005) hold the view that attention is in fact an imprecise term which (arbitrarily) applies to some of the subprocesses at work to achieve efficient dynamic routing of information in the brain.

A user-centered definition of attention would be the set of processes that enable us to direct resources towards processing some aspects of the retinal input more in depth than others. Eye movements (overt attention) determine what the input to our brain is. Covert attention (meaning, independent of eye movements) determines what subset of the information gets processed. At least, that’s how we vision scientists see the world. In most situations but laboratory experiments, observers will move their eyes towards the focus of their attention (unless it is socially inappropriate to do so). The machinery for eye

movements and the machinery for attention are tightly linked : the premotor theory of attention actually states that the current focus of attention is the locus of the next saccade (Rizzolatti, Riggio, Dascola, & Umilta, 1987).

So, we now have a general idea of what attention is. But the Devil is in the details (or is it God?). The word has been associated with such a diversity of processes that already in 1937 Spearman (Spearman, 1937) wrote:

But [attention’s] towering growth would appear to have been achieved at the price of calling down upon its builders the curse of Babel, “to confound their language that they may not understand one another’s speech”. For the word “attention” has been associated ... with a diversity of meanings that have the appearance of being more chaotic even than those of the term “intelligence”.

And let me tell you, there have been thousands of research papers on attention since then... My goal here is to point out the main concepts that emerged from the very rich field of attention research. This diversity of opinions is rather daunting for the new student of attention. I have said many times in my first couple of years as a graduate student : “this is a mess! you people don’t even know what you are studying! ‘attention’ doesn’t mean anything!”; my periodic peaks of anxiety translated into incoherent emails, in response to which Rufin always had the right words. My (superficial) wisdom now tells me that this diversity may reflect the fact that a full understanding of attention can only be achieved through consideration of a multitude of viewpoints.

a) Bottom-up and Top-down, a.k.a. exogenous and endogenous attention

The most basic distinction in the attention literature is between attentional selection automatically driven by the properties of the input (exogenous/bottom-up) and attentional selection of volitional origin, motivated by the goals of the subject (endogenous/top-down). The boundary between them may not be as clear cut as the distinction would suggest; at the single-cell level, their effects may actually be indistinguishable. However, they have different spatio-temporal properties, with exogenous attention being rapid, short-lived and fairly coarse, while endogenous attention is slower to deploy, but can be maintained for a long time and be very focused in space.

An influential school of thought, led by a seminal paper by Koch and Ullman (Christof Koch & Ullman, 1985), proposes that the conspicuity of different locations in the visual field is explicitly represented in a “saliency map”. The saliency map is a summary of multiple topographic maps, computed preattentively and in parallel for different features (orientation contrast, luminance contrast, color contrast, motion, ...) at different scales – the idea clearly derives from Treisman’s feature integration theory (A. M. Treisman & Gelade, 1980). In a purely feedforward fashion, the focus of exogenous attention chooses the most conspicuous location in the visual field, in a winner-take-all manner; as the saliency of the selected location wanes (by some inhibition of return mechanism) the focus moves to the next most salient location, and so on. Computer code has been generated since the original proposal (Itti & Christof Koch, 2000); the results, i.e. predicted locations of human eye movements when looking at a still picture or at a movie, are in fairly good agreement with human behavior : the area under the Receiver Operating Characteristic was 84% in the original implementation, and achieved 98% in (Harel, Christof Koch, & Perona, 2007). In that framework, endogenous attention may act to increase the saliency at some location or the weight of some channels in the saliency map. The lateral intraparietal area LIP (J. P. Gottlieb, Kusunoki, & M E Goldberg, 1998), frontal eye fields FEF (Tirin Moore, Armstrong, & Mazyar Fallah, 2003), pulvinar

(Petersen, D. L. Robinson, & Morris, 1987), and superior colliculus SC (Robert H Wurtz & Michael E Goldberg, 1972) have all been proposed as possible neural substrates of the saliency map.

Though the saliency map hypothesis is the most classically used in models of attention, some authors have suggested alternatives, or at least complementary ways of thinking about it. One influential framework was proposed by Desimone and Duncan (R Desimone & Duncan, 1995) : biased competition. In their model, there is no such thing as a master map of behaviorally relevant locations scanned by a mental spotlight; instead, attention arises from competition at multiple levels along the hierarchy of visual processing areas. This idea was furthered and a neural implementation using spike timing was proposed by Rufin (Rufin VanRullen, 2003).

b) Spatial, feature-based, object-based, what have you...

Perhaps the most well known demonstration of the cognitive effects of attention comes from Michael Posner and his colleagues (Posner, 1980) with their attentional cueing paradigm. The task was very simple : press a button as soon as you detect a brief flash of light. The flash was presented to the left or to the right of a central fixation point. The endogenous attentional manipulation came in the form of a central cue: an arrow pointing to the left or to the right, or a neutral cross, 1s before the flash appeared. The arrow was valid 80% of the time (meaning, the flash indeed appeared on the side where the arrow was pointing in 8 trials out of 10). The results showed faster reaction times in valid trials than in trials with a neutral cue; also, reaction times were slower in invalid trials than in trials with a neutral cue. Further experiments showed that it takes about 400ms to fully deploy attention to the cued side. This powerful attentional cuing paradigm was used with various cues to study exogenous attention as well (pull cues, as opposed to push cues (Jonides, 1981)). Pull cues fully attract attention in about 100ms.

These experiments inspired Posner to use the metaphor of a spotlight for attention : the location where attention is focused is illuminated and processed more effectively than surrounding locations that are less illuminated. The spotlight of attention is shifted from one location to the next (disengage/move/engage sequence). This metaphor and its predictions has led to a wealth of studies, with such questions as : how fast can the spotlight of attention move (Tsal, 1983)? does it illuminate objects along its trajectory from one location to the next (G L Shulman, Remington, & McLean, 1979)? how large is the beam (B. A. Eriksen & C.W. Eriksen, 1974)? is there a unique spotlight (C.W. Eriksen & Yeh, 1985)? Other metaphors have emerged, such as the zoom lens metaphor (C.W. Eriksen & St James, 1986), which allows for larger areas to be selected while keeping the amount of processing constant – a spotlight whose total power is constant, which shines brightly on a small region or dimly on a larger region.

The spotlight metaphor assumes that attention selects a region of space. Though this space-based vision of attention has dominated the field (and you will see that we keep faithfully to this tradition in our own work), it is but one view of attentional selection at the cognitive level.

There is ample evidence that attention selects objects, not just spatial locations. Duncan's (Duncan, 1984) experiments are highly cited in this regard. He used a stimulus consisting of two objects, a line and a box. Each object had two attributes : the line was either dotted or dashed, and tilted to the left or to the right; the box was either short or long, with a gap to the left or to the right. He showed subjects the stimulus briefly and asked them to report one or two attributes. In the case when two attributes had to be reported, the critical manipulation was whether they belonged to the same object or not.

Duncan found that if the two attributes belonged to the same object, subjects were just as good at reporting the two attributes as they were at reporting only one attribute. However, if they were part of different objects, performance on the report of the second attribute incurred a cost. The key feature in the design is that the two objects (the line and the box) are superimposed in space. A spotlight model can not account for these results. An object-based model of attention emerges as the best explanation. Since 1984, there have been many more demonstrations of an object-based form of attention (e.g. (O’Craven, Downing, & N Kanwisher, 1999)).

Yet another form of attention was evidenced a little later (Rossi & Paradiso, 1995; Gordon L Shulman & J. Wilson, 1987) : it is now known as feature-based attention. Rossi and Paradiso showed that performing a grating orientation discrimination task at the fovea lowered the threshold contrast to detect a grating with a similar orientation in the periphery; a grating with the same spatial frequency as the foveal grating did not however show any advantageous processing. Only when performing a spatial frequency discrimination task in the fovea did the threshold for detecting a grating of a similar spatial frequency in the periphery get lower. Other studies have shown how attention may be directed to other features such as color or direction of motion (M. Saenz, Buračas, & Boynton, 2002).

Are these views of attention compatible? Space-based, object-based, feature-based : these apparently contradictory views may be reconciliated if attention describes a collection of processes that operate at different levels of processing, as in the biased competition framework. The earliest selection may be described in terms of spatial attention: Francis Crick (Crick, 1984) hypothesized that the thalamic reticular complex, a thin shell surrounding the dorsal thalamus which contains almost exclusively inhibitory neurons and through which most of the traffic between the thalamus and the cortex is routed, may be where spatial selection operates, and recent neurophysiological evidence seems to support this view (McAlonan, J. Cavanaugh, & R.H. Wurtz, 2008). Feature-based attention is likely to operate at the levels where the various features are extracted (V1 for orientation, V4 for color, V5/hMT for motion...) or higher, while object-based attention probably operates at the levels where the scene is organized into objects (V2, V3, V4) or higher. Which brings us to the next question, which the biophysicist may have considered first : what exactly are the effects of attention on neural activity? and, what neural activity controls attentional allocation (we briefly alluded to some areas in our discussion of the saliency map framework)?

3. The neural basis of attention

Two recent reviews by Tirin Moore’s group (Noudoost, Chang, Steinmetz, & Tirin Moore, 2010) and James Bisley (Bisley, 2010) offer a good picture of the neural basis of attention, and I inspired myself largely from them (as well as the less recent but highly cited review by Corbetta and Shulman (M. Corbetta & G.L. Shulman, 2002)).

a) What areas of the brain guide attention allocation?

There is now a large body of evidence in monkeys suggesting that a network involving parietal (more specifically the lateral intraparietal area, or LIP) and frontal areas (more specifically the frontal eye fields, or FEF), as well as the superior colliculus, or SC, is involved in guiding covert attention (overt attention seems to rely on the same areas – this conforms to the previously mentioned premotor theory of attention). A schematic brain with the locations of these areas is shown in Figure 9.

There are two main lines of evidence to support this claim. On the one hand, a correlational line relying on studies of visual search (LIP (Ipata, Gee, Bisley, & Michael E Goldberg, 2009; N. W. D. Thomas & Paré, 2007); SC (McPeck & Keller, 2002); FEF (Buschman & Miller, 2007; Juan, Shorter-Jacobi, & Schall, 2004; K. G. Thompson, Biscoe, & Sato, 2005)). On the other hand, a causal line, using either inactivation (usually with muscimol; LIP (Balan & J. Gottlieb, 2009; Liu, Yttri, & L. H. Snyder, 2010; Wardak, Olivier, & Duhamel, 2004); FEF (Monosov & K. G. Thompson, 2009; Wardak, Ibos, Duhamel, & Olivier, 2006); SC (McPeck, 2008)) or microstimulation (LIP (Cutrell & Marrocco, 2002); SC (James Cavanaugh & Robert H Wurtz, 2004; J. R. Müller, Philiastides, & William T Newsome, 2005); FEF (Armstrong, Fitzgerald, & Tirin Moore, 2006; T Moore & M Fallah, 2001; Tirin Moore & Armstrong, 2003)). In humans, functional neuroimaging seems to point to a similar network (Maurizio Corbetta, Patel, & Gordon L Shulman, 2008).




Figure 9 Illustration of the primate brain. The locations of visual areas containing neurons whose responses are modulated by visual attention are shown in light grey. Cortical areas involved in the allocation of attention and the guidance of eye movements are shown in dark grey. The superior colliculus, also involved in this process, is not visible. *Reproduced from (Bisley, 2010)*

b) *The effects of attention on neural activity*

The most compelling effects of spatial attention have been evidenced (in V2, V4, IT, MT) when two stimuli, typically one preferred and one non-preferred, are simultaneously presented in the receptive field of a neuron : attention to the preferred stimulus enhances the response of the neuron while attention to the non-preferred stimulus decreases it. It is almost as if the receptive field of the neuron is shrunk around the attended stimulus (Figure 10).

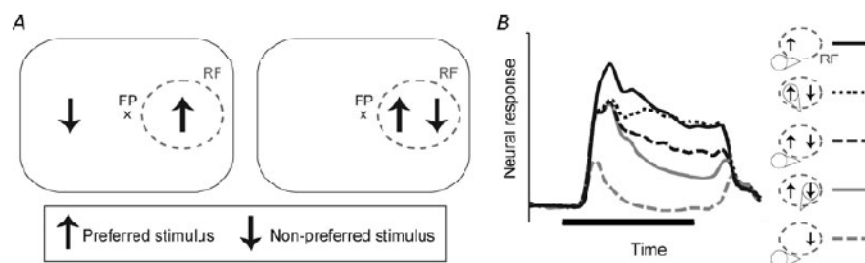


Figure 10 Studying the effects of attention on MT neurons. **A** the two main ways that the effect of attention have been studied in visual cortices. Usually at least 1 preferred stimulus (a direction of motion or oriented bar) and 1 non-preferred stimulus (the opposite direction of motion or a bar rotated 90 deg from the

preferred orientation) are presented on the screen. Either a single stimulus is placed in the receptive field (left panel), with the other stimulus in an opposite location, or both stimuli are placed in the receptive field (right panel). **B** the response of a population of 70 MT neurons under 5 attentional conditions (illustrated on right). The dashed grey oval represents the receptive field (RF), the cross represents the fixation point (FP), and the black cone illustrates the focus of attention. *reproduced from* (Bisley, 2010)

What is the effect of attention in the absence of competition, when a single stimulus is presented in the receptive field of a neuron? The response of neurons to stimuli in their receptive field has generally been found to be slightly enhanced when attention is focused on the location of the receptive field (case of spatial attention) : this is usually the case throughout the visual hierarchy, from LGN to V4 (ventral) and MT (dorsal) (Bisley, 2010).

Two mechanisms were classically proposed to model the observed modulations of neural activity: one held that attention modulates the gain of the neural response above baseline (Figure 11B), while the other suggested that there was a shift of the stimulus response curve (Figure 11A). These mechanisms yield similar predictions in most paradigms and are difficult to differentiate empirically.

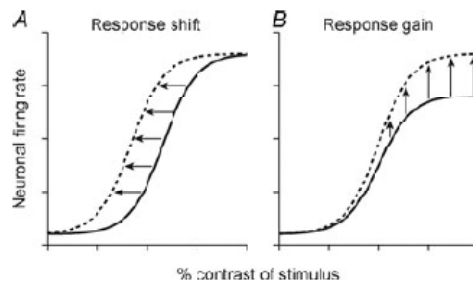


Figure 11 Theoretical effects of attention on the contrast response function. **A** the response shift model predicts that the response to an attended stimulus of a given contrast is the same as the response to an unattended stimulus of higher contrast. Thus, attention is similar to turning up the contrast of the stimulus. **B** the response gain model predicts that the response to an attended stimulus of a given contrast is a multiplicative increase in response to that contrast. Thus, attention is similar to just turning up the gain of the response. Continuous lines, unattended contrast response functions; dashed lines, attended contrast response functions. *reproduced from* (Bisley, 2010)

Recent modelling work (Reynolds & Heeger, 2009) proposed that both mechanisms may in fact be implemented at the neural level: which mechanism is at work depends critically on whether the stimulus of interest is smaller or larger than the attended area. This normalization model of attention has been very successful in predicting attentional modulations in recent experimental work. However, it still fails to explain some effects of attention which I will not cover in detail (e.g., shifting and shrinking of receptive fields); an extremely simple model was recently developed in our group which can account for all observed effects of attention, using only a few very simple organizational principles (namely, feedback connections and short range inhibition) – this “null model of attention”, by Miconi and VanRullen (Miconi & VanRullen, 2011), could well take over the normalization model.

Attention has been found to have many other effects on neuronal activity than firing rate modulations, which were reviewed in (Noudoost et al., 2010). Of particular interest to us, it affects oscillatory neuronal activity. Local power

increases have been evidenced in the gamma band in V4 (Fries, Womelsdorf, Oostenveld, & Robert Desimone, 2008; Gregoriou, Gotts, Zhou, & Robert Desimone, 2009; K. Taylor, Mandon, Freiwald, & Kreiter, 2005) and FEF (Gregoriou et al., 2009), as well as in the beta band in FEF (Buschman & Miller, 2007, 2009); these imply a local increase in synchronous synaptic activity, which has the potential role of increasing the effect of action potentials from the synchronized population onto downstream areas (this is but one possible role – see page 21 for more). In V4, it was shown that the spike-field coherence increases at frequencies >35Hz, while becoming less synchronized at frequencies <17Hz (Fries, Reynolds, Rorie, & Robert Desimone, 2001) – the spike-field coherence is a measure of synchronisation between the phase of the LFP and neuronal spiking. This decorrelation parallels the recent finding of an attentional reduction of the correlated variability of neurons in area V4, which results in an improved SNR (Cohen & J. H. R. Maunsell, 2009; Mitchell, Sundberg, & Reynolds, 2009). Gamma-band activity (induced by stimuli) in humans is also upmodulated by attention, as evidenced in the visual cortex (T. Gruber, M M Müller, Keil, & Elbert, 1999) and in the somatosensory cortex (M. Bauer, Oostenveld, Peeters, & Fries, 2006) (see (Ole Jensen, J. Kaiser, & J.-P. Lachaux, 2007) for a review). Moreover, attention may lead to increases in cross-areal, long range synchrony : FEF/LIP in 22-55Hz range (Buschman & Miller, 2007), FEF/V4 in 40-60Hz range (Gregoriou et al., 2009), LIP/MT in 25-45Hz range (Saalman, Pigarev, & Vidyasagar, 2007); this can provide a common temporal reference frame for communication – setting narrow windows in time to gate information flow, as a way to resolve competition.

Oscillations are low amplitude, coordinated modulations of the membrane potentials of thousands of neurons. I argued earlier that they may be the scale at which the study of neural activity currently has the best chance of yielding some insights about cognitive functions. The next section aims at familiarizing the reader with neuronal oscillations – which are of tantalizing importance in my work.

C. Neural oscillations and their functional significance

1. Oscillations in the EEG

Placing electrodes over the scalp is one of the cheapest ways to get some information about what the brain is doing. Electroencephalography (EEG), as this technique is called, picks up synchronized activity in neuronal networks. Neuronal communication relies on transmembrane ionic currents; the simultaneous release of charged ions by many neurons oscillating together leads to ionic currents in the extracellular space (volume conduction), which eventually reach the scalp and can be recorded – albeit at a very low signal to noise ratio, due e.g. to attenuation. The electrical potentials generated by single neurons are far too small to be identifiable in EEG recordings. Hence, the signal that we record reflects the summation of the synchronous activity of thousands of neurons, which, critically, must have similar spatial orientations (otherwise, the ionic current waves would cancel out): this is the reason why cortical pyramidal neurons are thought to be at the origin of most of the EEG signal – they are well aligned in the cortical layers. Another limitation of EEG is that deeper structures may not contribute much to the signal, because of the voltage field falling off with the square of the distance.

Perhaps the most famous EEG trace is the one published by Hans Berger in 1929 (Berger, 1929); it has a very prominent rhythmic component, at about 10 cycles per second – the alpha rhythm, the first recognized brain rhythm. Any PhD thesis which has something to do with oscillations probably should paste this trace somewhere (Figure 12).



Figure 12 First documented EEG recording (Berger, 1929). **(top)** EEG recorded from a young boy. **(bottom)** a 10 Hz frequency reference.

The alpha rhythm is the most prominent rhythm in the human EEG, in an awake subject at rest. Other rhythms have been evidenced and described following Berger's initial discovery, in clinical as well as fundamental research settings. A look at the power spectrum of the EEG may be a first step towards appreciating the diversity of rhythmic activity that the brain can generate. If spectral analysis is conducted on a long enough EEG segment, a characteristic $1/f$ power spectrum is obtained – meaning that the higher the frequency, the less it contributes to the overall waveform. However, this summarized view of brain rhythms is not very informative (though some people argue that it reflects scale-free dynamics and that it has interesting implications for brain function (Sadaghiani, Hesselmann, Karl J Friston, & Kleinschmidt, 2010)); the underlying reality is that there are discrete epochs of rhythmic activity, in different frequency bands and involving different cell assemblies.

At first sight, rhythmic activity in the EEG appears most prominent either at rest (eyes closed) or in pathology (epileptic seizures). The traditional classification into rather arbitrary bands – as established by the International Federation of

Societies for Electroencephalography and Clinical Neurophysiology in 1974 – was dictated by clinical observations of the brain in a default state. This classification is still widely used today, despite its limitations. The classical international clinical guidelines are summarized in Table 1.

Rhythm	Frequency band	localization	amplitude	behavioral correlates
α	8-13Hz	posterior	25-75 μ v	relaxed wakefulness
β	14-30Hz	anterior and middle	<20 μ v	active wakefulness/REM sleep
γ	>30Hz	anterior and middle	<20 μ v	active wakefulness/REM sleep
θ	4-7Hz	centrotemporal	20 μ v	wakefulness/REM sleep
δ	0.1-3.5Hz	diffuse	>30 μ v	slow wave sleep

Table 1 International clinical guidelines for brain rhythms. adapted from (Vion-Dury & Blanquet, 2008)

Buzsáki (György Buzsáki, 2006) argues that the international classification (α , β , ...) is of little use to the cognitive scientist. A typical example is that of the hippocampal “theta” (2-6Hz) oscillation discovered in the anesthetized rabbit. The corresponding oscillation in the awake animal, arising from the same physiological machinery, has a frequency of 5-10Hz. Should it thus be coined hippocampal “theta-alpha” in the awake rabbit? Moreover, the lower (0.1Hz) and upper (30Hz) limits on EEG frequency bands were imposed by technical recording limitations rather than real physiological limits. Penttonen and Buzsáki (Penttonen & György Buzsáki, 2003) proceeded to carefully observe oscillatory activity in the hippocampus and supragranular neocortex of the rat, and found that three distinct rhythms usually coexist – “delta” (1-3Hz), “theta” (6-9Hz) and “gamma” (30-50Hz). Both slower (M Steriade, 1999) and faster (R D Traub et al., 2003) rhythms have been found to be associated with these. Buzsáki and Draguhn (György Buzsáki & Draguhn, 2004) tentatively extrapolated existing observations to conclude that all rhythms from 0.02Hz to 600Hz may be present in the brain, and that at least ten distinct rhythmic generators are necessary to span the entire range of possible frequencies – the centers of the distinct frequency bands following a geometric progression according to the natural log (‘e’) (see Figure 13). This is likely to be an underestimate of the actual number of functionally distinct rhythmic generators. Classifying oscillations on the basis of the physiological mechanism that gives rise to them would be the most useful taxonomy – a lot of work remains to be done to this end.




Figure 13 Oscillation bands form an arithmetic progression on the logarithmic scale. For each band the frequency (Hz) or period ranges are shown together with their commonly used names. Reproduced from (Penttonen & György Buzsáki, 2003)

Roopun and colleagues (Roopun et al., 2008) used a secondary somatosensory cortical slice preparation and bathed it in various media to elicit “natural” rhythms. They found that they could generate more distinct rhythms than predicted by Buzsáki and Draguhn’s natural log progression (see Figure 14) – rather than ‘e’, the reason of the sequence may be ‘phi’,

the golden ratio... oh well. Interestingly, the co-occurrence of rhythms in a given set of conditions was the rule rather than the exception. This speaks to the likely functional importance of interactions between neuronal oscillations, which has been recognized under the names of frequency coupling, multiplexing, spectral processing...

Figure 14 Multiple modal peak frequencies of persistent rhythms generated in isolated neocortex *in vitro*. Reproduced from (Roopun et al., 2008)

2. How do oscillations come about?

I do not intend to give you a biophysics lecture here, because it would not offer too much insight at this point. It may suffice to know that there are three main factors that contribute to the emergence of oscillations in neural activity. Firstly, single neurons have oscillatory properties – a neuron which is fed a constant input will fire periodically, due to intrinsic properties such as its refractory period. Secondly, some circuit properties are known to give rise to oscillations : e.g., two reciprocally connected excitatory neurons will tend to synchronise their firing; also, shared inhibitory inputs make a cortical population alternate between periods of excitation and inhibition (this is the main mechanism for gamma oscillations). Thirdly, information transfer delays in the case of long-range connections (feedback loops) determine the activation periodicity (this is the main mechanism for low-frequency cortico-thalamic oscillations).

3. How can oscillations serve brain function?

a) *A reference signal for a spike timing neural code*

Modelling studies have shown that there is enough information in the first volley of spikes leaving the retinae for a first-pass analysis of visual scenes, using a spike timing code, in a situation where stimuli are briefly flashed on a screen (Rufin VanRullen & Thorpe, 2002). Of course, it is only possible to use a precise spike timing code if there is a reference time point. In lab conditions, the reference is provided by sudden stimulus onset. However, in real life, photons impinge continuously on the retinae, and an internal reference is therefore needed : this is a potential role for brain oscillations (Rufin VanRullen, Guyonneau, & Thorpe, 2005). There have been numerous demonstrations that precise spike timing can be used by neurons to encode information. The most famous example is provided by the phase precession phenomenon in the rat hippocampus: a so-called place cell fires when the rat is within a region of space called its place field. It has been shown that the position of the rat within the place field is encoded by the phase of the hippocampal theta cycle (4-10Hz) at which the place cell fires an action potential (Figure 15a): the theta cycle thus serves as a periodic temporal reference signal, with respect to which spike times can be used to encode information. Another case of spike time coding dependent on an oscillation has been established for whisking in rats (E Ahsissar, Sosnik, & Haidarliu, 2000): the horizontal position of an object can be retrieved from the onset of firing (induced by the contact of the whisker with the object) with respect to the onset of protraction of the whisker (Figure 15b). In the olfactory modality, a higher odor concentration generates earlier (and longer) bursts of action potentials coupled to the respiration cycle – odor concentration can thus be encoded in the firing latency with respect to the oscillatory respiratory rhythm (Cang & Isaacson, 2003). However, this finding comes

from anesthetized rats, with respiration rates much lower than observed during active sniffing (~2Hz vs ~8Hz) – to what extent it applies in behaving animals is an open question. Note that the proposal by VanRullen (Rufin VanRullen, 2003; Rufin VanRullen & Thorpe, 2002) that visual saliency is readily encoded in the timing of spikes in the first volley of action potentials following stimulus presentation is very clearly related to this finding – in the case of vision, rhythmic eye movements (or, as we will see later on, the rhythm of covert attention, which is the main object of this thesis), instead of the sniffing rhythm, could provide a temporal reference.

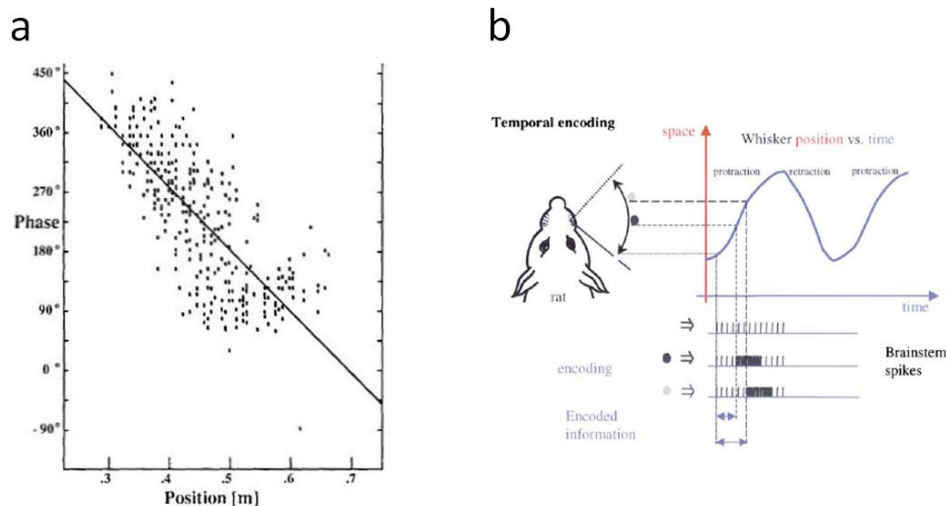


Figure 15 a Plot of the theta phase of a place cell against the position of the rat on the track. Reproduced from (O'Keefe & Recce, 1993) b Temporal encoding of horizontal position of objects with respect to the phase of the whisking cycle. Horizontal position of an object is encoded by the interval between the onset of the burst generated by protraction onset and the onset of the burst generated by touch. Reproduced from (Ehud Ahissar & Arieli, 2001)

b) Nested oscillation coding

In the rat hippocampus, there are prominent oscillations in the theta frequency band (4–10 Hz), as well as oscillations in the gamma frequency band (30–100 Hz). Both oscillations can occur together, and the phase of the theta rhythm was found to modulate gamma amplitude (Soltesz & Deschênes, 1993). Moreover, while the two rhythms can vary in frequency over quite a wide range, their frequency shifts are highly correlated (Bragin et al., 1995). Hence, the theta and gamma rhythms in the hippocampus are tightly interrelated. These observations led Lisman to propose that a sequence of items held in memory may be represented by a sequence of neuronal ensembles firing in different gamma cycles nested within a theta cycle (J E Lisman & Idiart, 1995). The theta cycle thus contains the whole memory store, while each item in memory is represented in a single gamma cycle; this model accounts well for the known capacity limit of 7 ± 2 items in human working memory. This model can also account for how a rat learns about its environment, and how information is replayed by phase precession during later navigation (O Jensen & J E Lisman, 1996). In essence, this scheme is but a special case of phase coding (as described in the previous paragraph), in which the time of spiking is constrained by the successive gamma cycles (Figure 16); it however has added explanatory power for many behavioral results.




Figure 16 **a** Hypothesized activity of four neurons (A–D) using a theta phase code. Any kind of information could be encoded by the phase of spiking (black tick marks) with respect to the theta cycle (gray line). **b** Same as **a** for a nested theta/gamma code. The ripples on the theta cycle represent gamma subcycles, and neural activity is constrained to occur within these. Reproduced from (Kepecs, Uchida, & Mainen, 2006)

Recently, this multiplexing scheme has been suggested to apply beyond the hippocampus. Jensen and Lisman, driven by the observation of theta/gamma nested oscillations in the entorhinal cortex, proposed that a cortical buffer using temporal multiplexing could explain sequence-encoding, by bringing items that are presented seconds apart within a temporal window short enough for Long Term Potentiation (Ole Jensen & John E Lisman, 2005). Lisman also proposed a more widespread use of the theta/gamma coding scheme in sensory cortex (J. Lisman, 2005). This proposal had already been formulated, for example in (F. Varela, J.P. Lachaux, Rodriguez, & Martinerie, 2001):

slower rhythms could provide the slower temporal framing for successive cognitive moments of synchronous assemblies, a slower beat within which beta and gamma rhythms operate

It was explored further in the case of visual representations by VanRullen and Koch (Rufin VanRullen & Christof Koch, 2003). VanRullen argues that slow and fast waves are common in the visual pathway (thalamus and visual cortex) - though slow waves are usually recorded in the alpha range rather than in the theta range. Alpha waves are global, traveling throughout cortex, with well-defined phase lags between distinct areas; in comparison, gamma waves generally correspond to more localized neuronal processes. VanRullen proposes that slow waves could constitute the 'context', and fast waves the 'content' of neuronal representations. He proceeds to describe some predictions of this representation which fit well with behavioral data : the number of visual objects that can be seen 'at a glance' is about 4, consistent with the frequency ratio of fast vs slow waves in the visual cortex; motor behavior may be triggered by each fast wave, predicting a ~25ms reaction time periodicity, while conscious perception would only be updated at the frequency of slow waves...

c) *Enhanced communication through phase synchrony*

Recently there has been much interest in high frequency neuronal oscillations (in the 30-100Hz range), as they have been linked to several cognitive functions (e.g., feature binding, object representations (Catherine Tallon-Baudry & O. Bertrand, 1999), attention (T. Gruber et al., 1999), memory (C Tallon-Baudry, O Bertrand, Peronnet, & Pernier, 1998)). Neurons whose membrane potentials oscillate in synchrony in the gamma band will spike together. If the synchronized neurons send their outputs to a downstream neuron, their effect will be enhanced (by temporal summation of their respective Excitatory Post Synaptic Potentials) (Figure 17a). This is a likely mechanism for attention. It has also been proposed that binding of the features of an object is done through synchronous firing of the neuronal populations coding for these features, which establishes a temporary neuronal coalition (Eckhorn et al., 1988; von der Malsburg, 1999; Singer, 1999). Finally, if the gamma rhythm synchronizes an input and an output area, the right conditions are set for Spike Timing Dependent Plasticity, a learning mechanism, to operate (Figure 17b).

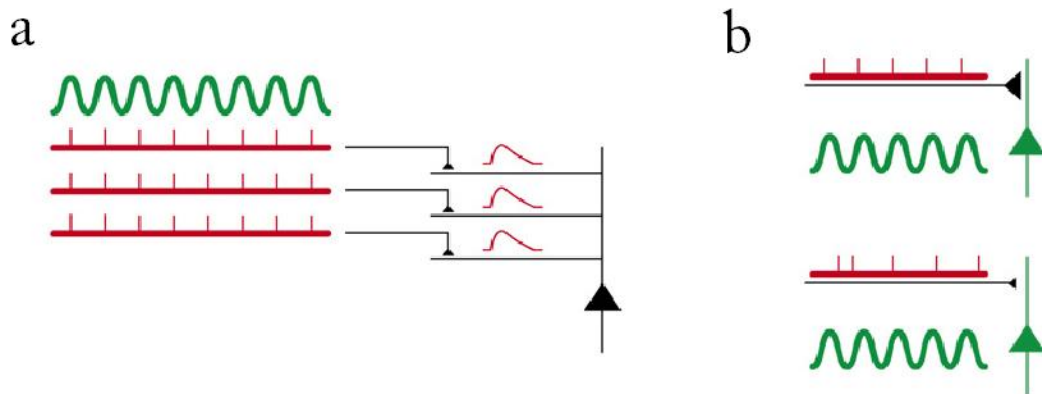


Figure 17 Examples of the functional role of neuronal gamma-frequency synchronization. **a** When the firing of the neurons in the network is synchronized, their effect on the receiving neuron increases because their inputs will add up. **b** Phase-specific synaptic input, with respect to the gamma rhythm, can facilitate synaptic plasticity and thus encoding of long-term memory. Consider a cell driven by subthreshold gamma-frequency oscillations. If incoming spikes are timed to the phase of the oscillations, this is likely to result in a stronger synaptic strengthening (top traces) compared with when the spikes are arriving independent of the phase (bottom trace). Adapted from (Ole Jensen et al., 2007)

More generally, Varela and colleagues proposed that phase synchrony in different frequency bands (not only gamma) could be the basis for large-scale integration and dynamic long-range interactions (F. Varela et al., 2001).

d) *A communication protocol*

Jensen showed that, if specific kinds of information are presented at distinct phases of the theta cycle, these can be read out selectively by a downstream area oscillating with the same theta rhythm, with the appropriate phase shift (O Jensen, 2001): this is a very simple mechanism by which oscillations can be used to gate information (Figure 18).

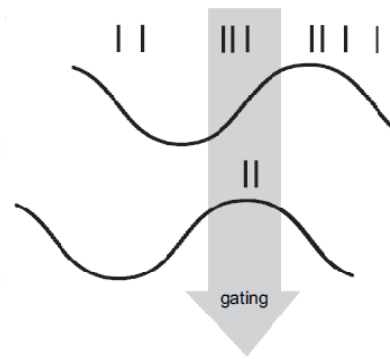


Figure 18 Two areas with coupled but phase-shifted theta rhythms could engage in selective gating of information. Adapted from (Kepecs et al., 2006)

These are just some of the functions that may be subserved by brain oscillations. As you can see, we have a good reason to study them as they may be a fundamental aspect of brain function.

D. Summary and aim of my thesis work

The brain does a wonderful job at interpreting incoming visual information, despite the many distortions that are imposed by the organization of the retinae. We are only aware of the end result, and have no conscious access to any of the intermediate stages of perceptual computations. As an example of the intermediate stages, the flow of incoming information is way above the rate that perception can handle, and there are attentional mechanisms that select the incoming information that should be processed. The computations that the brain performs for perception involve the concerted activity of widely distributed populations of neurons. In recent years, it has become increasingly clear that oscillations, routinely observed in electrical recordings of brain activity, may be shaping neuronal communication and long-range interactions, hence cognitive processes and more especially the computations underlying visual perception. If brain computations depend heavily on rhythmic activity, a question ensues that has seldom been asked in those terms: is there any evidence of rhythms in perception? We made an effort to argue in the first place that introspection wasn't going to be the best avenue of study. We shall thus seek evidence in more quantitative methods, namely psychophysics/behavior and EEG. Let the show begin.

II. Behavioral evidence of rhythms in attention and perception

Perception may rely on brain oscillations. This rather vague statement is compelling enough to ask the question of whether perception behaves periodically : this is a natural prediction. We therefore focus, in this part of the thesis, on behavioral and psychophysical observations which point to periodicities in perceptual processes.

Two major debates which are relevant to the quest for periodicities in behavior have divided the experimental psychology community for many years. First, and most obviously, the debate of whether perception relies on a continuous flow of information from the senses or on discrete samples from the stream of incoming information. Second, the debate of whether “simultaneous” focal attention to multiple locations (/objects/features) is truly achieved in parallel or relies on a serial mechanism. In a recent review (PAPER 3, page 85), we argue that these two debates are practically two facets of the same debate; the discrete VS continuous perception debate has traditionally been concerned with a single relevant stimulus at any point in time, whereas the parallel VS sequential attention debate deals with multiple simultaneously presented stimuli. Importantly, the finding that attention serially enhances concurrently attended stimuli would logically imply that attention is a periodic process, and thus that perception relies on a periodic refresh; however, if attention did not show such serial behavior, attention could still be inherently periodic, taking successive samples that include all attended locations when confronted with multiple stimuli. Hence, they remain a priori distinct questions; in this thesis we will go back and forth between them.

A. Two classical debates in the literature

In the following, I explore some of the psychophysical evidence that experiments fueled by the two debates have gathered. Rather than attempting an exhaustive enumeration of the various paradigms that have been used over the years, I concentrate on a few of the main lines of research to provide a good summary picture.

1. The discrete perception debate

An experimental psychology thesis without a quote from William James would be missing something, so here is one that I particularly like:

Suppose we were able, within the length of a second, to note 10,000 events distinctly, instead of barely 10, as now; if our life were then destined to hold the same number of impressions, it might be 1000 times as short. We should live less than a month, and personally know nothing of the change of the seasons. If born in winter, we should believe in summer as we now believe in the heats of the Carboniferous era. The motions of organic beings would be so slow to our senses as to be inferred, not seen. The sun would stand still in the sky, the moon be almost free from change, and so on. Now reverse the hypothesis and suppose a being to get only 1000th part of the sensations that we get in a given time, and consequently to live 1000 times as long. Winters and summers will be to him like quarters of an hour. Mushrooms and the swifter growing plants will shoot into being so rapidly as to appear instantaneous creations, annual shrubs will rise and fall from the earth like restlessly boiling water springs; the motions of animals will be as invisible as are to us the movements of bullets and cannon balls; the sun will scour through the sky like a meteor leaving a fiery trail behind him, etc.

This quote is taken from the chapter on the perception of time (James, 1890). The idea is that there are natural units of duration, likely to vary from one organism to the other. This proposal of quantized perceptual time originates in the belief that some kind of time sampling is necessary in order for one to detect change, or even to have an idea of time itself. Though it is originally a statement about the perception of time, the concept of a psychological unit of duration has an implicit consequence: that perception is built as a succession of discrete epochs. In the following, I review some experimental evidence that has been put forward to support this theory of discrete perceptual moments, inspiring myself largely from three existing reviews (Sanford, 1971; Rufin VanRullen & Christof Koch, 2003; C T White, 1963).

a) Studies of simultaneity and temporal rate : discrete epochs, or integration period?

If there is a perceptual unit of duration, an obvious prediction is that temporally discrete events that fall within a single unit of duration should be perceived as simultaneous, and that there should be an upper limit to the perceived rate of stimulation.

Exner (Exner, 1875) (according to (Hirsh & Sherrick, 1961)) reports a minimum interval of 44ms for two brief flashes of light occurring at the same position. Using a Savart wheel, he measured the minimum interval for perception of successiveness between two auditory clicks to be as low as 2ms. These experiments require the subject to say whether he saw a single stimulus or two distinct stimuli. Another question that may be asked is that of perceived temporal order – this question obviously requires stimuli to be distinguishable in another dimension than the temporal one. In the visual modality, Exner found that apparent movement may be seen for spatially separated flashes with intervals as short as 15 to 20ms. Hirsh and Sherrick (Hirsh & Sherrick, 1961) found that reporting the temporal order of two stimuli is quite independent of the sensory modality (they even tested cross-modal stimulation) and reaches 75% correct with a temporal separation of about 20ms. It is interesting to note that perfect performance is reached with intervals exceeding 80ms to 120ms, depending on subjects (table 1 in (Kristofferson, 1967)). Another interesting paradigm used in studies of simultaneity is to cyclically present an ordered sequence of segments of a composite form (or letters of a word), and to find the rate of presentation at which the segments appear to be presented simultaneously (can be done with the method of adjustment). The temporal spacing of the different segments may be uniform (Murphree, 1954) or not (Lichtenstein, 1961) – the finding is that the period of the full cycle is the only variable that matters, and that it needs be under 125ms.

The other line of evidence comes from studies of perceived rate of stimulation. Early studies were based on subjective estimation of the rate of a continuously flickering light (Le Grand, 1937) or perceptual matching of the rate of a continuously flickering source with a separate, adjustable flickering source (Bartley, 1951). Some methodological problems were pointed out by Pieron (Pieron, 1952) and another method was introduced by Cheatham and White (P G Cheatham & C T White, 1952). It consists in reporting the number of flashes perceived upon presentation of short trains of flashes at different rates (10, 15, 22.5 and 30hz in (P G Cheatham & C T White, 1952)). The main finding is that the number of flashes reported by subjects is primarily dependent on the time it takes to present the sequence, rather than on the actual number of flashes. Summarizing the data and the various follow-ups (Carroll T White & Paul G Cheatham, 1959), the maximum perceptual rate was found to be 12hz for all studied sensory modalities (vision, audition and touch), with little dependence on the conditions of the experiment. In the case of vision, it was shown that this was not a retinal limitation: for one thing, the retina can easily respond to flickering rates up to 125hz; moreover, the monocular presentation of two flashes at 7.5hz is easy to discriminate, whereas the interocular presentation of four flashes at 15hz (each eye receiving two flashes at 7.5hz) leads to perception of 2.4 flashes on average ((Andrews & Purves, 1997; Carroll T White & Paul G Cheatham, 1953)). The fact that all modalities share approximately the same limit seemingly points to a central process.

The paradigm with cyclic presentation of stimuli and the temporal numerosity paradigms seem to yield similar estimates of the duration of the perceptual unit of duration : between 80 and 120ms, corresponding to frequencies ranging from about 8 to 12 hz. It seems possible to make simultaneity judgements above chance at intervals as low as 20ms when only two brief stimuli are presented (2ms for auditory clicks! to be replicated experimentally...) – these judgements possibly rely on a different mechanism.

Arguably, and as pointed out in (Rufin VanRullen & Christof Koch, 2003) (and, before them, (Allport, 1968; Efron, 1973; Shallice, 1964)), these results do not constitute strong evidence for discrete perception. It is not enough to show that the brain has a limited temporal resolution – this can be accounted for by the concept of an integration period, a low pass filtering of incoming information, which has also been referred to as “iconic persistence” (Di Lollo & a E. Wilson, 1978). What the theory of discrete perception claims is that there is an ongoing brain process which organizes incoming sensory information in discrete epochs – hence, for a given SOA, the perception of simultaneity will be dependent on the exact state of the ongoing brain process, namely whether the two stimuli fall within an epoch or in two separate epochs. I should

note here that there may be a reset of perceptual cycles triggered by stimulus onset, as described in some variants of the theory, e.g. the event-initiated moment (Haber & Hershenson, 1973).

So why go through all these experiments, if they are not decisive? In fact, these experiments do give us some information on what the periodicity of sampling would be if there was such sampling. In a way they tell us where to look... Having just established that we need to know the exact state of the brain at the time of stimulus presentation to test the discrete perception theory, psychophysical methods without concurrent brain recordings may appear doomed. However, some clever experiments can be designed to test predictions unique to the discrete perception theory.

b) Some studies undoubtedly point to periodicities in perceptual processes

(1) Double detection function

If perception relies on discrete samples, implemented by brain oscillations, there is likely to be an off or down time between two perceptual epochs. The threshold for detecting a faint stimulus may thus oscillate, and depend on the state of the brain at the time of stimulus presentation. A very clever paradigm was described by (Latour, 1967) which circumvents the need to know the exact state of the brain at the time of stimulus presentation. It is impossible to control the time of presentation of a single flash with respect to ongoing perceptual cycles without monitoring brain activity (NB: it is not trivial either when you monitor brain activity, because the neural basis of perceptual cycles is not clear!); consequently, the potential effects of an underlying sampling rhythm will average out over trials, as the phase of the ongoing perceptual cycles will be random. However, the probability of detecting both flashes can be shown to oscillate with respect to the delay between them. For instance, very caricaturally, if the delay corresponds exactly to one period of the underlying excitability cycle, both flashes will fall at the same phase : if it is in the off phase, both stimuli will be missed, whereas they will both be detected if it is in the on phase. If the delay corresponds instead to a half period of the underlying rhythm, one stimulus will invariably be missed when the other one is detected – the probability of double detection falls to zero. Latour introduced the proper mathematical formalism, which allows to derive the frequency of the underlying rhythm from double detection psychometric curves. I will not get into it though, as I do not use it in the remainder of this thesis. In fact, while this method is interesting, it has very little power – the results of Latour (preliminary evidence that visual threshold oscillates in the gamma range) are difficult to replicate, as we discussed in PAPER 3 (page 85). While its applicability is limited, this method shows that it is theoretically possible to observe effects of rhythmic processing with pure psychophysical methods. This being established, it is only a matter of designing methods that are applicable in practice...

(2) Rhythmic updating of spatial conflict

A recent illusion discovered by Arnold and Johnston (D. H. Arnold & A. Johnston, 2003) may be a case of perceiving the rhythmic updating of perception. The display they used was a small green disk surrounded by an isoluminant red ring, itself embedded in a black background. As the red/green stimulus moved, the green disk was seen to jitter, as if it was lagging behind and its position was corrected periodically. In a recent investigation of the neural correlates of the illusion, Amano and colleagues (Amano, D. Arnold, Takeda, & A. Johnston, 2008) measured the flickering rate, which they found to be 10-11hz (consistent with the literature reviewed thus far). This “motion induced spatial conflict” illusion and its interpretation

have been largely ignored by the community (only 9 citations since 2003 for a Nature paper!), but recently Holcombe (Alex Holcombe, 2011) argued that it may be the least controversial evidence to date of a perceptual rhythm.

(3) Periodicities in reaction time distributions

Another finding of periodicities in behavior has been discussed in the context of the discrete perception debate : upon inspection of reaction times histograms, some authors have evidenced periodicities at 100ms (Venables, 1960) and at 25ms (Dehaene, 1993). Discrete epochs of processing may be the cause for such periodicities : if a stimulus arrives too late to be processed in a given epoch, it is deferred to the next one. Seeing periodicities in reaction time histograms however also implicates that the discrete epochs of processing are synchronized with stimulus presentation – they represent a rhythm that is reset by the stimulus. While this finding is of interest in the context of a discussion on the influence of oscillatory brain activity on behavior, it is slightly less relevant to our quest of perceptual moments.

2. The sequential attention debate

When more than one stimulus are concurrently presented, are they all represented in each perceptual sample? Some empirical evidence has suggested that, when several stimuli are relevant to a task, attention may enhance each of them sequentially, in a cyclic fashion.

Just as with the discrete perception debate, the original theory was dictated by philosophy rather than empirical experiments. Attention was initially seen as a unitary and indivisible resource. In this context, serial allocation of attention was the most sensible mechanism, implicitly accepted (Charles W. Eriksen & Spencer, 1969; Kahneman, 1973; A. M. Treisman, 1969). Different paradigms have been used which provide arguments for or against this theory, which I briefly review here.

a) *Many classical paradigms fail to provide a definite answer*

(1) Visual search

As put loosely by Jeremy Wolfe (Wolfe, 1998) :

Visual search tasks are those tasks where one looks for something.

Slightly more formally, in the visual search paradigm, subjects are presented with an array of items, and instructed to look for a target among distractors; their response is a two-alternative forced choice : is there a target in the display or is there not? For instance, subjects may be looking for a circle among squares (feature search) or for a red circle among red squares, green circles and green squares (conjunction search). Here I consider studies where reaction time (RT) is the measure of interest; the display remains visible until the subject responds. RT is generally analyzed as a function of set size (the number of distractor items), and the experimenter derives two functions - one for target present (positive) and one for target absent (negative) trials. The slopes (and intercepts) of these generally linear functions are used to infer the search mechanism. The classical finding (A. M. Treisman & Gelade, 1980) is as follows : in feature search, reaction time is short and independent of the number of distractor stimuli (“pop out”, interpreted as a parallel search mechanism); in

conjunction search reaction time can be longer and appears to be proportional to the number of distractors (interpreted as a serial search mechanism; the slope – for target absent trials – is roughly 40-60ms per item in a typical conjunction search). While the reaction time slopes in conjunction search have been interpreted as the time it takes for attention to examine each item in turn until the target stimulus is found, this interpretation is not warranted by the data – an alternate model has been proposed, according to which attention is always distributed in parallel among items, and the increase in reaction time with increasing number of distractors simply reflects the increasing task difficulty (M Carrasco & Yeshurun, 1998; M. P. Eckstein, 1998; Miguel P. Eckstein, J. P. Thomas, J. Palmer, & Shimozaki, 2000; McElree & Marisa Carrasco, 1999; J. Palmer, 1995). The visual search paradigm may never be able to distinguish between these two models (Townsend, 1990).

Wolfe in fact argues that “the serial/parallel dichotomy in visual search is deceased and ought to be allowed to rest in peace”(Wolfe, 1998): one of his stronger arguments is that, rather than a dichotomy, a continuum of RT slopes exists. This shows that, although originally attractive for our purpose of uncovering a rhythmic behavior of attention, the visual search paradigm suffers from many confounds and is not easy to work with. So, again, why mention this paradigm? This time, it is simply because it is one of the most classical paradigms for students of attention... not mentioning it would be almost as unfathomable as not mentioning Newton in a classical mechanics textbook.

(2) Simultaneous/sequential paradigm

The simultaneous/sequential paradigm was designed especially to test the parallel versus serial mechanisms of attentional deployment. It compares accuracy performance between simultaneous and sequential presentations of otherwise equivalent stimuli. Because the critical information is presented for the same time, an unlimited capacity parallel allocation of attention would predict no difference in performance for simultaneous and sequential presentations, whereas a serial model would be impaired in simultaneous presentations, missing information while it is sampling one item and ignoring the others (Charles W. Eriksen & Spencer, 1969; Shiffrin & Gardner, 1972). This is a pretty seductive idea, but in practice this paradigm suffers from many confounds, just like the visual search paradigm : some phenomena such as grouping, crowding, apparent motion, masking are difficult to control for between the two conditions, and the real effects of attentional mechanisms become very difficult to retrieve.

(3) Multiple object tracking

Another popular paradigm is multiple-object tracking : a display with multiple items is shown to the subject, a few of which are cued. The task consists in keeping track of these items while they undergo random motion for about 10s, then reporting their final positions. The fact that it is possible to track four items (and sometimes more) has been taken as evidence for multiple foci of attention, hence a parallel allocation (Patrick Cavanagh & Alvarez, 2005). However it has also been argued that a sequential model of attention could account for most of the experimental data as well (Howard & A. O. Holcombe, 2008; Oksama & Hyönä, 2008). Once again, confounding factors such as grouping may limit the interpretability of multiple object tracking experiments in terms of attentional mechanisms...

(4) Split spotlight

There is a host of studies that are interested in whether attention can be allocated in a sustained manner to two disjoint locations. It is often assumed in those studies of divided attention that attention can not move beyond a certain speed. This is usually justified by some reference to a study of the speed of voluntary attention shifts, e.g. (Duncan, Ward, & K.

Shapiro, 1994; H. Hogendoorn, Carlson, & F. A. J. Verstraten, 2007; C. M. Moore, Egeth, Berglan, & Luck, 1996). Hence, having short stimulus presentations (<200ms) is deemed enough to ensure that attention is truly divided rather than sampling relevant stimuli in turn (Awh & Pashler, 2000; Bichot, K R Cave, & Pashler, 1999; McMains & Somers, 2004; M M Müller, Malinowski, T. Gruber, & S. A. Hillyard, 2003).

All these classical paradigms in the study of attention fail to provide definite evidence in favor or against a natural rhythmic behavior of attention. The inherent difficulty is that any effects of “set size” can usually be explained equally well by a serial strategy or a limited-capacity parallel strategy (Townsend, 1990). Is our quest for periodic behavior of attention a practical dead-end? There is a need for a paradigm in which different attentional mechanisms would make different predictions.

b) The blinking spotlight of attention

A recent psychophysical study (Rufin VanRullen, Carlson, & Patrick Cavanagh, 2007) systematically varied set size and presentation time independently on a large number of trials in a detection task. From the performance on a given set size and duration of the target, exact predictions could be made as to what the performance should be with other set sizes and durations, according to a few models of attentional deployment (serial, unlimited/limited capacity parallel, ...). Critically, these predictions differed, and thus the correct model could be derived – as the one whose predictions offered the best fit to the data. While a classical serial allocation of attention was shown to be outperformed by the parallel model (see also (J. Palmer, 1995)), a “sample always” model of attention was the overall winner. When attention is divided, successive samples naturally focus on different stimuli, but when it is concentrated on a single target, the samples continue to occur repeatedly every ~150ms, accumulating evidence for this one stimulus. This is reminiscent of the previous debate about discrete epochs of perception. The authors coined the term “blinking spotlight” to describe attentional mechanisms. It is thus possible (not easy, but possible) to evidence a periodic behavior of attentional processes, using psychophysical methods.

In section B (page 34), I will present a method that has had some success recently in uncovering an inherent rhythm in perceptual processes : the use of temporal aliasing artefacts to evidence an underlying sampling process. I will then describe my attempts to extend the initial findings derived from this paradigm, in a variety of settings. In section C (page 52), I will present evidence against a sustained division of attention (and thus indirectly in favor of a serial behavior of attention), through the critical revisiting of an existing paradigm. In section D (page 70), I propose a promising avenue of research, complementary to the approaches that I described in normal subjects. A large part of our knowledge of the brain was derived from observations of cases when it stopped functioning properly, due to disease or lesions. It is not until your car breaks down and you have to open the hood that you really start understanding how it works... I will argue that some disorders may reflect periodic mechanisms at work in perceptual processes.

B. Temporal aliasing artefacts as evidence that perception relies on periodic processes

1. Aliasing in the continuous Wagon Wheel Illusion

a) A definition of temporal aliasing

In information theory, aliasing is tightly linked to the sampling theorem. The Nyquist-Shannon sampling theorem states that to be able to fully reconstruct any given signal, you need to sample it at at least twice the highest frequency present in the signal : this critical frequency is called the Nyquist frequency (this is a very simplified, layman version of the theorem). If you fail to do so, interpolating the samples will lead to the construction of an aliased signal – a signal different from the original signal.

Figure 19 illustrates the very simple case of a sine wave oscillating at a given frequency. By sampling it inappropriately (every $\frac{3}{4}$ of its period), reconstruction yields a much slower sine wave (with a period equal to three times the period of the original signal) which can also be seen to rotate in the “wrong” direction in the complex domain. Take a simple device that samples an input signal at sampling frequency F_s and outputs a signal reconstructed from the samples it took. If this device is fed a signal that contains frequencies higher than $F_s/2$, it will necessarily output an aliased signal.




Figure 19 Sampling a temporal signal using too low a sampling rate leads to systematic errors about the signal, known as "aliasing errors". Reproduced from PAPER 3 (page 85)

b) The continuous wagon wheel illusion (cWWI): a case of temporal aliasing?

(1) A bistable percept with maximal illusory strength at 10Hz

The “classical” wagon wheel illusion consists in the apparent reversal of the direction of rotation of wheels (or propellers) as a consequence of the framed nature of movies. If a one-blade propeller is rotating clockwise with a period T , and the frames of the movie are taken with a period $3T/4$, the blade will seem to have rotated 90 degrees counterclockwise from one frame to the next, hence the spurious apparent motion. This is a case of temporal aliasing.

Surprisingly, some variant of this phenomenon occurs under continuous conditions of illumination, for example in plain sunlight (Purves, Paydarfar, & Andrews, 1996). This *continuous* version was further investigated by a few groups (K. Kline, AO Holcombe, & Eagleman, 2004, 2006; Piantoni, K. A. Kline, & Eagleman, 2010; Rojas, Carmona-Fontaine, López-Calderón, & Aboitiz, 2006; Rufin VanRullen, 2006, 2007; Rufin VanRullen, Reddy, & Christof Koch, 2005, 2006). It differs from the cinematographic wagon wheel effect on one major point: the percept of reverse motion is not stable, it occurs transiently and only for about 20% of the viewing time. It is thus a bistable percept. A temporal frequency of 10Hz was found to maximize the illusion strength (percentage of the viewing time spent reporting the illusory direction of motion), and that was independent of the actual velocity of the wheel’s spokes. The illusion was observed for both first-order and second-order motion stimuli (Figure 20)

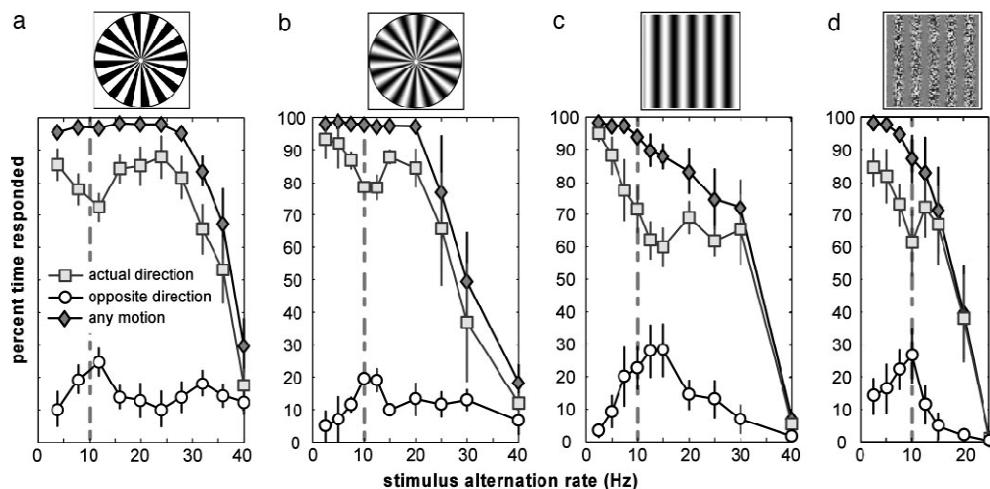


Figure 20 Moving stimuli were displayed continuously for 60 s while subjects reported the perceived direction of motion using computer keyboard arrows. Although the dominant direction of motion was always the actual direction, observers also reported perceiving the opposite direction for a considerable fraction of the time. **a** rotating sunburst patterns under natural, continuous illumination. **b** rotating radial sinusoidal patterns. **c** horizontally drifting luminance-modulated (first-order motion). **d** horizontally drifting contrast-modulated gratings (second-order motion). **b, c** and **d** were displayed on a computer screen with 160hz refresh rate. Figure and caption reproduced from (Rufin VanRullen, Reddy, et al., 2005)

The illusion requires attention: if attention is directed to another task (e.g., a Rapid Serial Visual Presentation, or RSVP, of randomly rotated letters at fixation), subjects paradoxically perform better at reporting the motion of an ambiguous

stimulus. This was shown using unbalanced counter phase gratings, which consist of two gratings of the same spatial frequency with slightly different contrasts drifting in opposite directions (Rufin VanRullen, Reddy, et al., 2005).

The only component of the electroencephalogram power spectrum that was affected (negatively) during the illusion was found at ~13Hz, mostly in the right parietal regions (Figure 21) (Rufin VanRullen et al., 2006); the changes in this frequency band two seconds before a reversal were enough to predict reversal occurrences above chance.

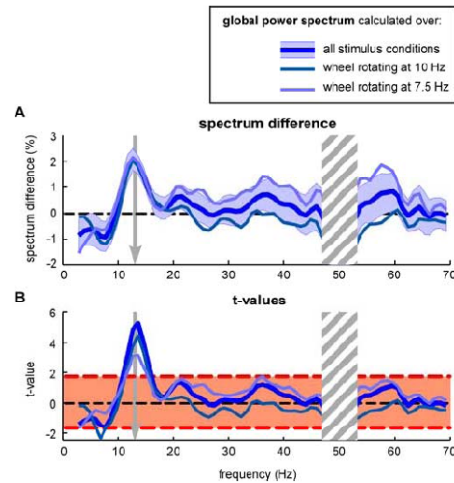


Figure 21 The power difference at 13Hz is independent of the temporal frequency of the stimulus. **A** The difference in global power spectrum (averaged over all subjects and electrodes) between the periods of real and illusory motion percepts is computed for all stimulus conditions together and also separately for the wheels rotating at 7.5 and 10 Hz. The peak difference at 13Hz is common to both curves. **B** The t values confirm that the difference at 13Hz is highly significant for both stimulus conditions. The dashed red lines mark the 0.05 significance threshold. Figure and caption reproduced from (Rufin VanRullen et al., 2006)

It was further shown that the illusion is object-based, rather than occurring at a specific location (Rufin VanRullen, 2006). It depends on neuronal adaptation, as other bistable percepts, but is not totally explained by it – the illusion and adaptation, as indexed by the motion aftereffect, do not always covary (Rufin VanRullen, 2007).

(2) The mechanisms of visual motion perception : a quick primer

(a) *Passive versus Active motion perception*

It was recognized in 1974 by Braddick (Braddick, 1974) that motion perception involved at least two distinct systems : the short range system and the long range system. The main distinction between the two proposed subsystems was their spatio-temporal range, the short range system operating over short distances and brief durations and the long range system operating over long distances and long durations. The observations that led Braddick to posit the existence of these two systems relied largely on random dot kinetograms as a typical example for the short range system and apparent motion of isolated stimuli for the long range system. Additional distinctions were made over the following years between the two systems in terms of aftereffect production (only for short range), handling of chromatic stimuli (short range can't do it) and dichoptic stimulation (short range can't do it). The short range system was seen as an automatic, passive system thought to occur at early levels of processing (tentatively, directionally sensitive cells in striate cortex) and working in

parallel over the entire visual field, while the long range system was considered a more cognitive, interpretive mechanism, an active tracking system taking place at higher levels of processing. The original distinction became problematic as short range like motion was found to occur over long distances and long range like motion over short distances; also, short range like motion was eventually found to occur for chromatic stimuli and even dichoptic stimulation, while long-range motion was found to yield aftereffects. Cavanagh and Mather (P Cavanagh & Mather, 1989) pointed out these inconsistencies and proposed to rid the dichotomy of its original implications. After proposing a new dichotomy, again based on stimulus properties (first-order: luminance and color, low level attributes; as opposed to second order: texture, relative motion, binocular disparity, higher level attributes), they realized that sensing of motion for these two systems is likely performed passively by similar arrays of sensors, in parallel over the whole visual field (Reichardt-like comparators). In a later contribution, Cavanagh (P Cavanagh, 1991) emphasizes the most important distinction that can be made between motion systems: that between a passive motion perception system (responding to both first order and second order defined motion) and an active motion perception system. The latter amounts to attentive tracking of the position of moving stimuli. The active motion perception system can respond to the same stimuli that activate the passive system, although their sensitivity may not be the same. Critically, Cavanagh demonstrates a stimulus in which the two systems can disagree on the perceived direction of motion – proving that they are indeed distinct systems : the stimulus consists of counterrotating radial gratings, one defined by luminance (small modulations) and one defined by color. While the passive system signals the direction of rotation of the luminance grating, the active system clearly perceives the opposite direction of rotation when tracking the colored spokes. There are other examples of competition between the passive and the active motion perception systems . Interestingly, Lu and Sperling (Z.-L. Lu & George Sperling, 2001; Z.-lin Lu & George Sperling, 1996) refer to the active motion perception system as the “third order” system, in analogy to the first and second order systems originally proposed and discarded by Cavanagh. They maintain that there is a fundamental distinction between the first and the second order systems – however, this is not a view shared by all, and the distinction between an active and a passive motion perception system is in any case a far deeper one.

(b) *Motion sensor models*

How does passive motion perception work? Passive motion sensors detect motion, after necessary feature extraction / texture grabbing operations have been performed on the input. Three models of motion sensors have been proposed. The first computationally based model of motion detection was proposed by Reichardt, as a theory for beetle vision (Reichardt, 1961). It was then adapted to human vision (van Santen & G Sperling, 1985). The elaborated Reichardt detector performs a simple spatio-temporal correlation of signals from neighboring points in an image. The idea is simple: if there was a movement of a (luminance) pattern along the axis connecting the two points under scrutiny, the signal present at one point is a time shifted version of the signal at the other point. By comparing the evidence for motion in either direction along the axis, an elaborated Reichardt detector can decide in which direction the stimulus moved (Figure 22a). Alternatively, motion can be seen as orientation in a space time diagram. This has led other authors to propose the existence of neurons with receptive fields selective for orientation in space-time (Adelson & Bergen, 1985). These receptive field properties arise readily from simple summation of units with different spatio-temporal profiles, as illustrated in (Figure 22b). Two linear motion sensors, whose responses are 90deg out of phase, can be combined by squaring and summing to produce motion energy filters. The final stage consists in comparing the outputs of the motion energy filters for opposite directions. Formally, this model is equivalent to the elaborated Reichardt detector, in the sense that it will output the same result given the same input. Yet another approach has been proposed based on spatiotemporal gradients, which consists in taking the ratio of the temporal and spatial gradients of luminance in the moving image (Fennema & W. Thompson, 1979). Though the motivation is different from the previous two approaches, and appears more suited to

computer vision problems than to biological systems, there are great similarities between the motion energy model and this spatiotemporal gradient approach (Bruce, P. R. Green, & Georgeson, 1996).

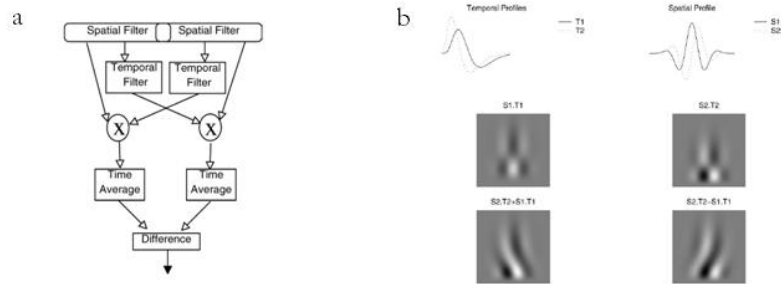


Figure 22 Models of passive motion sensors **a** Two-stage elaborated Reichardt detector **b** Motion energy filters can be constructed by combining units with different spatio-temporal profiles.

In essence, all proposed computational models of motion sensing have the same input-output relationship, and it is impossible to tease them apart on the sole basis of psychophysical experiments. While physiology could potentially help decide between these models, currently available data is not enough yet to clearly elect one of the models.

(c) *Where in the brain is motion computed?*

Studies in the rabbit retina have revealed a neural organization which provides all the wiring required for calculating the direction of image motion. Many ganglion cells were found to be directionally selective. Interestingly, the structure of motion detectors in the rabbit retina does not correspond exactly to the theoretical models I just described: without entering in too much detail, motion in the preferred and anti-preferred directions is actually driven by different systems rather than identical systems with mirror symmetric tuning. Presumably, the structure of the retina in other mammals is fairly similar to that of the rabbit. In cats and primates, direction sensitive cells are found in the pretectal nucleus of the optic tract and in the accessory optic system, nuclei which project to motor areas in charge of stabilizing eye movements; it is likely that direction sensitivity comes mostly from their inputs (retinal ganglion cells, and cortical neurons). In all species studied (cats, primates, rabbit, opossum, wallaby), directional responses are common in the primary visual cortex. Other areas in the visual cortex are known to specialize in coding motion information, notably the middle temporal area (MT/V5) in primates. The majority of cells in MT are direction-selective, but they do not seem to compute motion themselves; rather, they receive directional signals from earlier stages in the visual system. These earlier stages include V1, but also subcortical structures such as the colliculus and the pulvinar. The evidence points to V1 as a major location for motion computation, and suggests that directional properties of some V1 cells are similar to those predicted by stages of the Energy model. This is a very schematic view of the field, and I refer you to the excellent review by Clifford and Ibbotson (Clifford & Ibbotson, 2002) and the papers they cite for further details.

(3) Competing interpretations of the cWWI: quasi-periodic attentional sampling or aliasing at the level of motion detectors?

(a) *Quasi-periodic attentional sampling*

This hypothesis is driven by the observation that attention is necessary for the illusion to occur. It posits that the illusion arises from a competition between a passive motion perception system, which mostly signals the veridical direction of motion (and maybe to some extent the opposite direction, as argued by Eagleman and colleagues (K. Kline et al., 2004)) and an active motion perception system which tracks the spokes of the wheel. The passive motion perception system is prone to adaptation and, when adapted, may lose the competition in favor of the attentive tracking system. Critically, attention does not operate continuously, but takes quasi-periodic samples. For attentional samples to lead to the observed tuning of the illusion, with a maximum around 10Hz, the samples should be taken at the rate of about 13 per second (in a motion energy model, the opposite direction of motion has the most energy when the wheel turns $\frac{3}{4}$ of a turn between two samples – it then looks as if it had turned $\frac{1}{4}$ of a turn in the opposite direction). Such attentional sampling leads to temporal aliasing, and to motion reversals when adaptation allows active motion perception to take over temporarily. Recently, evidence from neuroimaging and transcranial magnetic stimulation both point to the involvement of the right parietal area in generating the illusion, which is in line with the causal importance of attentional processes (Reddy, Rémy, Vayssière, & Rufin VanRullen, 2011; Rufin VanRullen, Alvaro Pascual-Leone, & Battelli, 2008).

(b) *Spatial aliasing at the level of Reichardt-like motion detectors, and motion “during-effect”*

Eagleman and colleagues (2004, 2006, 2008) argue that illusory motion reversals (note that they refuse to use the term “continuous wagon wheel illusion”, because they do not want to emphasize the analogy with the stroboscopic wagon wheel illusion experienced in movies) are the consequence of spatial aliasing at the level of Reichardt-like motion detectors. A periodic stimulus is likely to activate Reichardt-like detectors that signal the direction of motion opposite to the real direction, in addition to activating detectors signaling the real direction. Follows a competition for interpreting the motion of the object under consideration, initially dominated by the detectors tuned to the real direction of motion. But as the observer keeps looking at the moving stimulus, adaptation to the perceived direction of motion occurs and there may be transient episodes when the spurious direction of motion takes over. This interpretation can explain most observations. Eagleman and colleagues argue that even the tuning of the illusion (at about 10Hz) is not surprising, as this frequency is the one to which the visual system is most sensitive (S. J. Anderson & D C Burr, 1985; Snowden & R F Hess, 1992). The involvement of attention is not surprising either. Attention is known to increase adaptation to motion stimuli (Alais & Blake, 1999; Chaudhuri & Chaudhuri, 1990; Culham, F. A. J. Verstraten, Hiroshi Ashida, & Patrick Cavanagh, 2000; Lankheet & F. a Verstraten, 1995; Nishida & H Ashida, 2000; Rezec, Krekelberg, & Dobkins, 2004; Gordon L. Shulman, 1991), so more attention means faster adaptation which will increase the occurrence of the illusion...

Holcombe (AO Holcombe & Seizova-cajic, 2008) argues that spatial aliasing of motion detectors may not even be necessary. When the population of cells responding to the true direction of motion has been adapted enough, spontaneous activity in the population of cells tuned to the opposite direction of motion may be enough to lead to illusory motion reversals : he thus proposes a motion ‘during-effect’.

(c) *No agreement yet*

Many experiments have been conducted to feed the debate over the mechanisms accounting for the continuous Wagon Wheel Illusion. Every argument made by detractors of the sampling theory was tentatively rejected, but the follow-up experiments that were used as counter-arguments can be criticized in turn. For example, there is no doubt now that

attention plays a role in the illusion, and that activity in the right parietal cortex is important for the illusion to occur (fMRI then TMS). But the role of attention may simply be to increase adaptation, while evidence in favor of the opposite direction of motion comes from another source (aliased motion detectors or spontaneous activity). This possibility, though it was rejected on the basis of a study which found an increase in the switching rate of rivalrous stimuli upon rTMS stimulation to the right Inferior Parietal Lobe (Carmel, Walsh, Lavie, & Rees, 2010) (meaning that rTMS on the rIPL may increase adaptation), can not be completely ruled out yet. Importantly, VanRullen demonstrated that the continuous wagon wheel illusion and adaptation do not always covary (Rufin VanRullen, 2007) – arguably the strongest argument to date against alternative accounts. VanRullen’s argument was based on opposite effects of eccentricity, and of contrast, on the illusion and on motion adaptation. Detractors may however argue that other factors may change with eccentricity, and with contrast, that influence the perception of the illusion (e.g., distribution of motion sensors for eccentricity). Yet another major challenge to the attentional sampling theory was posed recently by the perception of reversals in a non-periodic stimulus (drifting random dots (K. Kline & Eagleman, 2008)).

One clear merit of the c-WWI paradigm is that it puts forward temporal aliasing as a possible sign of periodic processing for perception in the brain. It is only natural to try and extend the findings of the c-WWI paradigm to other motion stimuli, to visual features other than motion and even to sensory modalities other than vision. If evidence for temporal aliasing could be found in these cases, the corresponding sampling frequencies may then be compared to one another and further inform our understanding of perception. Is it the case that there is a single rhythm, a central (attentional) clock that samples all sensory inputs? Or is information from any single channel of sensory information read out periodically at its own rate, independently from other channels? While the first proposition reflects the understanding that most have of the theory of discrete perception (K. Kline & Eagleman, 2008), the latter may be a much more faithful description of reality; additionally, the sampling rate for a given channel may depend strongly on task demands and attentional state, further blurring intrinsic periodicities.

c) How to probe the brain for temporal aliasing artefacts

The simple generic paradigm which we advocate to probe the brain for temporal aliasing is as follows. Human observers are presented with a periodic time varying input which physically evolves in an unambiguously defined direction; they are asked to make a two-alternative forced choice judgment on the direction of evolution of this input, whose frequency is systematically varied by the experimenter across trials. A consistent report of the wrong direction at a given input frequency may be taken as a behavioral correlate of temporal aliasing, and the frequencies at which this occurs inform the experimenter about the underlying sampling frequency of the brain for this input.

2. Seeking aliasing artifacts in other sensory modalities – the case of auditory motion perception

Humans have five senses : vision, audition, touch, smell and taste. Audition and touch are arguably the senses other than vision that have the best temporal resolution. They are thus a priori the best candidate senses to investigate processing of a time-varying signal. Incidentally, about 6 months after I started this thesis, Holcombe and Seizova-Cajic published a study (AO Holcombe & Seizova-cajic, 2008) where they report illusory reversals in touch and proprioception (see also (Seizova-

cajic, Smith, J. L. Taylor, & Gandevia, 2007)). In the case of touch, there's not much more than a report in the abstract, really; 6 subjects out of 7 think that they may have perceived reversals when placing three fingers on a textured rotated drum... the authors did not investigate it further, which may point to the unreliability and lack of controllability of this finding. The illusory reversals that they reported upon proprioceptive stimulation are of more interest. Subjects perceive arm extension when their biceps is stimulated with a vibrator. After prolonged stimulation, 12 out of 15 subjects report perceiving reversals of the illusory arm motion (arm contraction). The authors suggest that this is a proof of concept for a motion "during-effect". The point is taken but high-level influences on the cWWI, as discussed previously, are hardly compatible with such low level adaptation of sensors. As others had already investigated the touch modality, without too much success, I preferred to focus on the auditory modality.

a) Auditory motion in the spatial domain

(1) The mechanisms of auditory motion perception : a quick primer

(a) Perceiving the location of sound sources

There are two main cues that the auditory system can use to compute the spatial location of auditory sources. First, sound from the right side reaches the right ear earlier than the left ear – this leads to an interaural time difference, or ITD, which the auditory system can evaluate (using phase delays or group delays depending of the carrier frequency). Second, sound from the right side has a higher level at the right ear than at the left ear, because the head shadows the left ear – this leads to an interaural level difference, or ILD (ILD increases with increasing frequency, because higher frequency signals carry less energy and consequently dampen faster). For frequencies below 800Hz, mainly ITD are evaluated; for frequencies above 1600Hz, mainly ILD are evaluated. Between 800Hz and 1600Hz there is a transition zone, where both mechanisms play a role. These mechanisms cannot be used to differentiate between a sound source ahead of the hearer or behind the hearer; therefore additional cues have to be evaluated. The human outer ear, i.e. the structures of the pinna and the external ear canal, form direction-selective filters. These patterns in the ear's frequency responses depend highly on the shape and size of the individual's outer ear. If sound is presented through headphones, and has been recorded via another head with different-shaped outer ear surfaces, the directional patterns differ from the listener's own, and problems will appear when trying to evaluate directions in the median plane. As a consequence, front-back permutations or inside-the-head-localization can appear when listening to dummy head recordings. Finally, the human auditory system has only limited possibilities to determine the distance of a sound source. In the close-up-range there are some indications for distance determination, such as extreme level differences (e.g. when whispering into one ear) or specific pinna resonances in the close-up range.

(b) How spatial motion is sensed remains unclear

We know relatively little about pathways for auditory motion processing. Given how sound source location is computed, it seems plausible that some cells could respond to changes of ILD and ITD over time, thereby directly sensing sound source movement, in a manner similar to the elaborated Reichardt detectors. In fact, such cells have been found in the primary auditory cortex of cats and monkeys (and in some subcortical structures) as reviewed in (J. W. Lewis, Beauchamp, & E. a DeYoe, 2000), but they still have to be found in primates. Functional MRI (Warren, Zielinski, G. G. R. Green, Rauschecker, & Griffiths, 2002) in normal subjects shows a bilateral posterior network of activation to sound source motion, including

planum temporale (PT) and parieto-temporal operculum (PTO), just posterior to Heschl's gyrus. Functional MRI in sight recovery patients who were blind from an early age recently uncovered the involvement of hMT (V5) in the perception of auditory sound (Melissa Saenz, L. B. Lewis, Huth, Fine, & Christof Koch, 2008); however, this is not the case in healthy sighted subjects. It is unclear whether there is a passive auditory motion sensing system in humans. The functional imaging experiments that have been conducted are not conclusive on the existence of such a system. There is however evidence for the ability to attentionally track the spatial position of sound sources in time : auditory "apparent motion" shares many characteristics with its visual counterpart (Huddleston, J. W. J. W. Lewis, Phinney, & E. A. E. DeYoe, 2008). If attentional tracking is involved in spatial motion perception and if extraction of position information is fast enough, we may have a chance to find evidence of aliasing when the sound sources move fast enough, by analogy with the visual cWWI. There is of course no guarantee that attentional sampling for the auditory modality will occur at the same rate as attentional sampling for the visual modality.

(2) A dead end

After creating many stimuli, using ILD and ITD cues, translational and rotational motion, and presenting them through headphones and even a semi-circular array of eight speakers, we finally came to the realization that the auditory perception of motion in the spatial domain was limited to very slow speeds. In fact, a recent paper was published that reports temporal limits similar to the ones we observed for the judgement of motion of an auditory source, namely 3-4Hz (Féron, Frissen, Boissinot, & Guastavino, 2010).

The temporal resolution for discriminating the direction of the time-varying input under consideration should be at least as good as the hypothesized sampling frequency. If the psychometric function is already at chance at the frequency where aliasing is expected to take place, this aliasing will simply not be observed – whether the perceptual process relies on periodic sampling or not. It was immediately apparent that perception of spatial auditory motion did not satisfy this condition – although we did not realize straight away that this was the problem.

This failure may have been expected, considering that spatial information is not an explicit feature of audition, but a constructed one (as we described earlier). In the primary visual cortex, each neuron has a receptive field well defined in retinotopic coordinates. Hence, position is an explicit feature for vision; each neuron corresponds to a spatial location. In the primary auditory cortex, it is different: each neuron has a receptive field which is well defined in tonotopic space; meaning, it responds to sounds in a limited frequency band. Frequency is the explicit feature for audition. Thus, time-varying signals in frequency space for audition should be the closest equivalent to time-varying signals in retinal space for vision. This reasoning provides the motivation for the next attempt.

b) Auditory motion in the frequency domain

(1) The Shepard scale and the Risset glissando

How do we create a periodic stimulus in motion in the frequency domain, with a constant direction of evolution? In 1964, Shepard designed an auditory sequence based on 12 tones, which played continuously gives the impression to continuously ascend (or descend) – it has been likened to the visual illusion of the barber pole, or to the Penrose stairs (Figure 23a). Each tone in the Shepard sequence (Figure 23b) consists of a superposition of 4 to 10 octaves – the

components with the lowest and highest frequencies have low amplitudes while the middle frequencies have higher amplitudes. The fundamental frequency is increased (decreased) from one tone to the next, and 12 tones constitute the whole scale. Low (high) frequency harmonics progressively appear in the ascending (descending) sequence, while high (low) frequency harmonics progressively disappear – the 13th tone would correspond exactly to the first tone, the sequence is periodic and can be played repetitively. The illusion of an endlessly ascending (descending) pitch with the Shepard scale works best if there is a gap between successive tones. The Risset scale is a continuous version of the Shepard scale. Instead of having discrete steps, it is a continuous frequency sweep with several superimposed harmonics. Finding the right amplitude for each harmonic so as to create a perfect illusion is difficult – and in fact, when paying close attention to most examples out there, the illusion is never really perfect...




Figure 23 a The Penrose stairs, a visual analog to the Shepard scale. b A Shepard scale spanning 4 octaves.

(2) A dead end again

The trick used for the Shepard and Risset scales breaks down as temporal frequency is increased beyond 3-4Hz, and artefacts become prevalent which prevents further experimentation. It is possible that judgements of direction could be performed in frequency space at very high temporal frequencies, but we could not design a periodic stimulus continuously evolving in the same direction.

An interesting parallel to the limitation we encountered, both with the spatial and frequential stimuli, is the midstream order deficit (AO Holcombe, Nancy Kanwisher, & A. Treisman, 2001): if you present a sequence of 4 sounds repeatedly, or a sequence of 4 letters repeatedly, humans are impaired at perceiving the relative order of the sequence, compared to a situation when the sequence is presented only once (even though they can report the individual elements). Various experiments (variations on pre-masks, ...) show that the deficit is due to the lack of a salient “start” to the sequence, the lack of an exogenous anchor point. Their interpretation is that in the absence of an exogenous anchor one has to deploy endogenous attention to define a first element for the sequence but as endogenous attention takes too long then by the time it is deployed the element has changed...




Figure 24 The midstream order deficit. Mean probability of reporting the correct relative order of a sequence of 4 letters, as a function of letter duration, for the cycling (3 cycles, preceded by rapid presentation which progressively slows down) and single-presentation (1 cycle, embedded in a stream of symbols) conditions. If the participants knew which four letters were presented, but not their order, the relative order accuracy would be 16.6%, since there are six possible relative orders for four items (represented by the dotted line). Reproduced from (AO Holcombe et al., 2001)

Once this was realized and I had mourned the idea of designing an auditory analog of the wagon wheel illusion, it was time to come back to good old vision and start looking for more evidence of aliasing.

3. Does attentive (third order) visual motion perception rely on a sampling process?

According to the interpretation of the cWWI that we have favored, the attentional motion perception system is the one that suffers from aliasing – the automatic (first-order) system provides the correct direction of motion, but a competing, wrong interpretation coming from the attentional (third-order) system takes over transiently. A simple prediction that follows is that a moving stimulus designed to be invisible to the first-order system should be strongly affected by aliasing. The performance of the subject should look something like the psychometric curve depicted in Figure 25, assuming a 13hz sampling rate and an envelope such that at alternation rates of 50hz and higher, the subject cannot discriminate the direction of motion any more.




Figure 25 Theoretical output of a motion perception system sampling position information at 13hz. This is based on a simple motion energy model, similar to the one depicted in (Rufin VanRullen, Reddy, et al., 2005), and an overall envelope which imposes an upper limit to motion perception.

a) *Interocular contrast-defined motion*

(1) Stimuli and paradigm

We used a contrast-defined grating (the mean luminance is the same everywhere, only the contrast of a random dot texture changes), which moved horizontally. A whole period of the motion stimulus consisted of four frames, with a 90 degree phase shift from one frame to the next. If the subject is presented with such a stimulus binocularly (both eyes see all the frames), this is a typical example of what Lu and Sperling have called second order motion. By alternating frame presentation between the two eyes (the left eye may for instance see frames 1 and 3 while the right eye sees frames 2 and 4), we get a situation in which the direction of motion is ambiguous for each eye taken separately. Only by integrating the information registered in both eyes can the brain perceive the true direction of motion. According to Lu and Sperling (Z.-L. Lu & George Sperling, 2001; Z.-lin Lu & George Sperling, 1996), the perception of motion with such interocular motion presentation is handled by the third order system, i.e. the attention-based system. Figure 26 offers a depiction of this paradigm.




Figure 26 Display for the interocular second-order motion experiment

Thirteen subjects were run with the following paradigm : 2-second trials, 90 trials per frequency, with 11 different frequencies chosen according to the limitations of a screen with 160hz refresh rate : 1, 2, 3, 4, 5, 6.7, 8, 10, 13.3, 20 and 40hz). A binocular control, in which all frames were presented to both eyes, was also run for all subjects.

(2) Results

As seen in Figure 27, the performance averaged over all 13 subjects falls earlier for interocular motion than for binocular motion, an expected result if interocular motion relies mostly on the attentional motion perception system which is known to have lower limits, while binocular motion presentation allows automatic monocular motion detection mechanisms to contribute to the percept. However, the presence of aliasing is not evident from the curve corresponding interocular motion presentation. What is happening? Does this dismiss the theory of discrete attentional sampling as the basis of the cWWI?

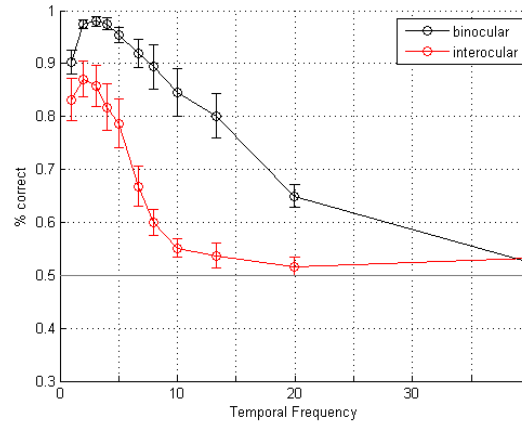


Figure 27 Grand mean performance of thirteen subjects on the interocular and binocular motion paradigms. Both curves show an overall sigmoidal shape, with the interocular curve falling earlier than the binocular curve.

It could be that interocular motion perception does not truly rely on attentional motion perception alone, unlike what Lu and Sperling have claimed. Our use of very short trials does not allow any adaptation to occur, and the contribution of the attentional system to performance may be masked by the contribution of another mechanism. To address this possibility, we decided to try other stimuli that were reported to be processed by the attentional motion perception system alone.

b) Other third order motion stimuli

Further attempts with stimuli such as stereo-defined motion, isoluminant motion, salience-defined motion were all aborted early when we realized that most observers could not discriminate the direction of motion above chance at temporal frequencies higher than 3-4hz. In essence, we encountered the same limitation that we had with auditory stimuli earlier. It is impossible to see the effects of aliasing, expected to occur around 10hz for motion stimuli, if there can't be a reliable judgement of direction beyond 3-4hz... It should be kept in mind that the limitation comes from the time it takes to extract features – this is already the argument that we put forward in the case of auditory stimuli. We can't extract features fast enough to compute motion with attention, in the case of stereo or salience-defined motion. In the case of luminance or contrast-defined motion, features may be extracted more rapidly and the temporal frequency cutoff for attentional motion perception may be higher, allowing it to contribute to the percept at higher frequencies (at least 10hz).

4. A model fitting approach to detect aliasing in psychometric curves

In the previous paragraphs, we have described obstacles that we encountered in our quest for temporal aliasing. There are three main hurdles which may have prevented us from seeing temporal aliasing artefacts in psychometric functions so far:

- It could be that the sampling frequency is slightly variable over the time of the experiment, depending on the subject's motivation and arousal for example. A small percentage of variation could make the aliasing disappear in the average performance.
- It could be that performance falls to chance too quickly to detect any aliasing (see the argument that we developed in the case of auditory motion perception and third-order motion). Aliasing is expected to occur around 10hz, however performance is already pretty much at chance at that frequency in all paradigms that we have tried. We have called this the 'curse of seeing slow', in reference to Holcombe's classification of visual paradigms according to their temporal limits, 'seeing slow' and 'seeing fast' (AO Holcombe, 2009)
- It could finally be that the effects of aliasing are hidden by the contribution of another system, working continuously, to perception in the situations that we investigated.

To deal with these nuisance factors, we introduced the following method (Figure 28): psychophysical performance should be modelled as the sum of two processes, one sampled (aliasing function with two parameters, the mean sampling frequency μ_{al} and the standard deviation of the sampling frequency, σ_{al}) and one not sampled (cumulative gaussian envelope with two parameters, the mean of the gaussian μ_{env} and the standard deviation of the gaussian, σ_{env}). The contribution of the sampled process is C.




Figure 28 Modelling performance as the sum of two processes : on the left, a “continuous” process whose output is an inverted cumulative normal distribution characterized by its cutoff temporal frequency (μ_{env}) and how sharply performance falls off around it (σ_{env}); on the right, a “discrete” process that can be characterized by its sampling frequency (μ_{al}) and the variability of the sampling frequency (σ_{al}). The evidence from the discrete and continuous processes is summed linearly, the discrete process contributing a fraction C of the overall performance.

The aliasing function we used was derived from a Fourier energy model, similar to the one used in (Rufin VanRullen, Reddy, et al., 2005). Its output was fitted with a damped sine function (4 parameters). Simulations allowed us to generate a lookup table, reducing the aliasing function to only 2 parameters. The parameter search bounds that we enforced were as follow: $2 < \mu_{al} < 20$; $0.1 < \sigma_{al} < 1$; $3 < \mu_{env} < 30$; $0.05 < \sigma_{env} < 0.5$; $0 < C < 1$. The fitted performance for a representative subject is shown in Figure 29.




Figure 29 Performance of a representative subject on the interocular second-order motion task (blue stars), fitted by our model (red curve). The values of the parameters are reported.

Fitting 5 parameters to 11 datapoints may seem quite questionable without a proper statistical assessment. Under the null hypothesis, the psychometric curve is expected to have a classical sigmoid shape (Figure 30). Such a theoretical curve, under the null hypothesis, can be fitted to the experimental data (2 parameters for a simple inverted cumulative gaussian). Taking the proportion correct given by the cumulative gaussian fit at each tested frequency as the true probability of correct direction discrimination, we can simulate many surrogate experiments by performing N binomial trials at each tested frequency. These surrogate datasets are submitted to the 5-parameter fit and the contribution of aliasing is estimated for each of them. The p-value for the contribution of aliasing in the real data can be estimated in this fashion (as well as the p-values of the other parameters). For each subject, 1000 surrogate experiments were generated; then, they were randomly combined to yield 100M surrogate grand means for the contribution of the sampling process to motion discrimination performance.

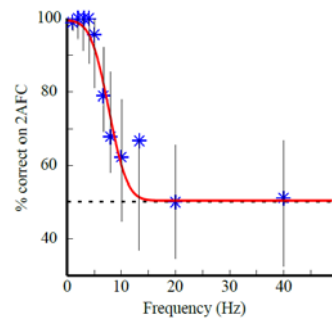


Figure 30 Under the null hypothesis, the true theoretical curve which describes performance is an inverted cumulative gaussian (red curve), and any departures from this curve that were observed are manifestations of noise. 1000 surrogates datasets were generated using binomial trials at each temporal frequency, based on the true underlying probability (red curve). The black lines at each temporal frequency tested in the experiment are the 95% confidence intervals.

We found that there was a significant contribution of a discrete process, with a grand mean frequency of 12.5Hz, to the judgement of motion direction in both the interocular and the binocular second-order motion discrimination paradigms (Figure 31); moreover, the interocular paradigm yielded a much larger contribution for the discrete process, which is in keeping with its purported increased reliance on attentional motion perception mechanisms. Finally, the 16% contribution in the binocular paradigm corresponds quite well with the reported 20% occurrence of direction reversals in the cWWI.




Figure 31 The contribution of the discrete process (red line) compared to its distribution under the null hypothesis (blue histogram), for the interocular paradigm on the left and the binocular paradigm on the right.

The results that this method uncovered were not readily available upon inspection of the mean psychometric curve. Of course, this method relies on many assumptions (models for the shapes of the performance predicted from a continuous process and from a discrete process, linear summation of information from the two sources) and their applicability should be a matter of criticism. The interplay of the various parameters is also a matter of concern in this approach (Figure 32). Furthermore, though the grand mean results concerning the contribution of the discrete process are encouraging, the individual results are extremely variable – fitting 5 parameters to 11 datapoints does leave a lot of room for variability and misleading local minima. All things considered, we still consider that this approach is an encouraging first step towards detecting the influence of discrete processes in psychometric functions.




Figure 32 Single subject results. Comparison of the value of each parameter of the fit (red line) to their estimated null distribution (blue bars). Some interactions between the various parameters should be further investigated.

5. Summary and Discussion

As I mentioned previously, it may be that some residual processing by automatic motion detection mechanisms is available for the purported ‘third order’ stimuli that we considered. In fact, this is what our modelling analysis seems to point to. This rescues the attentional sampling account of the cWWI...

At some point along our quest, we tried to tackle attention directly, outside of the realm of motion perception. In tasks that require attention to multiple locations, the attentional sampling account logically would predict that the spotlight should switch rapidly, at the intrinsic frequency of attentional processes. By manipulating the presentation of information in time, we should thus be able to use aliasing as a signature of the natural rhythm of attention. One task that we explored along this avenue was binding of spatially separated color and orientation features, which was shown to be performed by attention (AO Holcombe & Patrick Cavanagh, 2001). I ran many pilot versions until I finally realized that the same problem was again plaguing my attempts; as you probably expect by now, the ‘curse of seeing slow’ hit us again... indeed, beyond an alternation rate of about 3-4hz, the task could hardly be performed above chance level, hence potential artefacts due to temporal aliasing were not detectable.

What of the intriguing motion standstill phenomenon reported by Lu and colleagues (Z. L. Lu, L. A. Lesmes, & G Sperling, 1999; Z. L. Lu, L. a Lesmes, & G Sperling, 1999), in the case of isoluminant chromatic stimuli? Could it be a manifestation of temporal aliasing? I have not looked at the details to check whether discrete sampling could explain the illusion. I contacted Zhong-Lin Lu (back in 2007) and told him that I would like to experience the illusion, that I would come to his laboratory to see it. He told me that it was a very difficult experiment, and that it took extremely precise individual tuning – which I pretty much took as an informal no. I have not tried very hard to replicate the chromatic illusion, so I have never experienced it so far... I did some work on the easier to set up stereo-standstill stimuli (Tseng, Gobell, Z.-lin Lu, & George Sperling, 2006), but when replicating their paradigm I ran into some ‘real’ temporal aliasing artefacts – as the stimuli are based on discrete frames (4 per period as here), increasing the temporal frequency of the motion happened to yield stimuli sequences that were physically moving in the direction opposite to the intended one... Hence I am not fully convinced of their demonstration of motion standstill in the case of stereo stimuli. In fact, the authors themselves claim that the quality of the standstill was much better in the case of isoluminant motion. But like St Thomas, I need to see it with my own eyes to believe that it exists.

How can we further exploit the study of temporal aliasing artifacts, if we are limited to the cWWI paradigm? One avenue of research that we started exploring is to test the cWWI paradigm in patient populations to gain a better understanding of their deficits – this is of course highly dependent of the interpretation of the cWWI, but let’s consider that it is established.

Migraine patients reportedly have impairments for motion perception (Granziera, DaSilva, J. Snyder, Tuch, & Hadjikhani, 2006; McKendrick, D. R. Badcock, J. C. Badcock, & Gurgone, 2006); we predict that they will perceive the illusion more readily as they are presented with a rotating wagon wheel for the first time (in healthy controls, the time-to-first-inversion –with no guidance from the experimenter other than fixating and watching the stimulus for an extended time – is quite variable, on the order of minutes), and that they will experience reversals more often than healthy controls (automatic motion perception mechanisms are malfunctioning and patients’ motion perception relies more heavily on attentional processes). We do not have a prediction in terms of the optimal frequency for the illusion.

Impaired neuronal synchronization (K J Friston & Frith, 1995; Tononi & Edelman, 2000) and a disturbance of the sense of continuity (Fuchs, 2007) have been reported in schizophrenic patients. Recently, it has been shown that their ‘temporal window of integration’, as assessed using simultaneity thresholds, is longer than in healthy subjects (J. R. Foucher, Lacambre, Pham, a Giersch, & M a Elliott, 2007; A. Giersch et al., 2009). From these observations, one may hypothesize that attentional sampling in the schizophrenic population may be slower than in the healthy population : the optimal frequency for reversals in the cWWI paradigms is expected to be lower than it is for healthy subjects.

We started collaborations with laboratories in Strasbourg, France (A. Giersch) and in Taipei, Taiwan (C.-T. Wu). These projects are still in their infancy.

C. Psychophysical evidence against a sustained split of the attentional spotlight

To account for the observations of the cWWI paradigm, we hypothesized that attention takes periodic samples of the incoming sensory information to judge motion direction. More generally, attention has been shown to behave as a “blinking spotlight”: according to (Rufin VanRullen et al., 2007), attention operates in a periodic fashion at all times – and particularly, it oscillates between concurrently attended locations in a situation where more than one location in the visual field contains relevant information. This periodic, serial view of attention is not shared by all, as I mentioned previously. I reviewed earlier a few paradigms that have been used over the years to feed the debate over parallel VS sequential attention. In this section, I propose to focus again on the evidence that comes from the “split spotlight” studies, and provide new experimental results that support the serial model.

1. Can the spotlight of attention be split?

a) *Existing experimental evidence is inconclusive*

Can the spotlight of attention be divided in a sustained fashion, enhancing processing at two separate locations simultaneously while ignoring a location in between? In 1998, McCormick and colleagues (P. A. McCormick, Klein, & S. Johnston, 1998) reviewed the literature and concluded that

most of the evidence favors the unified model.

In 2008, Adamo and colleagues (Adamo, Pun, Pratt, & Ferber, 2008) stated, in the introduction of their paper, that

there is abundant evidence showing that observers can deploy attention to separate locations.

A thorough recent review by Jans and colleagues (Jans, Peters, & De Weerd, 2010) states that none of the evidence put forward in the last 15 years conclusively demonstrates that divided spatial attention can be achieved in a sustained manner. This review is also thoroughly biased, judging from the initial statement :

the notion that spatial attention can easily be split is counterintuitive.

I propose to review here some of the key findings and key arguments in the debate – I too will necessarily present a biased sketch of the literature... It is important to stress that in the case of feature-based attention or object-based attention, attentional enhancement is readily seen at disjoint spatial locations – however, I have argued in the general introduction that these forms of attention operate at higher levels in the visual processing hierarchy. The split spotlight debate should thus be limited in scope to purely spatial attentional selection. The boundary between spatial attention and feature-based or object-based attention remains a thin and ill-defined one. For example, Gobell and colleagues (Gobell, Tseng, & George Sperling, 2004) cued attention with colored patches, using feature-based attention as a means to distribute spatial attention; in such an experiment, when the cues disappear, does the finding of a processing enhancement at

noncontiguous locations constitute receivable evidence that spatial attention can sustainably select more than one location at a time? We are treading muddy waters... The debate of whether attention can be divided or not could in fact be extended to feature-based or object-based attention, if one does not lose from sight what is being attended to. In the case of spatial attention, the question is whether disjoint retinal coordinates can receive simultaneous processing enhancement. In the case of feature-based attention, the question should be whether two disjoint coordinates in feature space can be simultaneously attended to; for example, can one attend to red and blue at the same time? In the case of object-based attention, the question will be about two or more objects whose processing is simultaneously enhanced. For the time being, let us focus on the debate around spatial attention – the split spotlight studies.

It is difficult to demonstrate a split of attention experimentally while definitely ruling out all alternative explanations. The finding of two disjoint regions of enhanced processing may indeed be the consequence of a truly parallel allocation of attention, but can in most cases be explained as easily by serial switching of a unitary spotlight of attention between the two locations. If the landscape of attentional enhancement is not sampled properly, it is also theoretically possible that a single (possibly oddly shaped) focus of attention encompasses the seemingly noncontiguous locations. These are the main concerns one should keep in mind when designing an experiment which aims at resolving the debate. According to Jans and colleagues (Jans et al., 2010), most studies that have claimed to provide definite evidence for or against a sustained division of spatial attention failed to properly control for alternative models, in some way or another... They proposed a list of criteria which any experiment should fulfill before its results can be considered as evidence for or against a sustained division of spatial attention (see also (Kyle R Cave, Bush, & T. G. G. Taylor, 2010)), which includes the concerns that I just raised (the target stimuli must be presented rapidly to prevent shifts of attention; the attention probes must eliminate all possible configurations of a unitary region of attention). Additionally, they argued that the task should be difficult enough to require selective attention – a criterion that Cave and colleagues (Kyle R Cave et al., 2010) dismissed easily: if there is evidence of disjoint foci of enhanced processing with no enhancement at an intervening location, a criterion of task difficulty is irrelevant. Finally, they propose to limit the preparation time between the attentional cue and the target – the rationale is not clear, though it may be a good way to prevent eye movements in an experiment in which they are not controlled. As argued in (Kyle R Cave et al., 2010), some studies of divided attention may “survive the cut” : a few psychophysical studies – Experiment 6 in (Bichot et al., 1999), (Awh & Pashler, 2000), (Kawahara, Enns, & Di Lollo, 2006); and an electrophysiological (M M Müller et al., 2003) and a couple of brain imaging studies (McMains & Somers, 2004, 2005). I will come back to the EEG and fMRI studies in part II of this thesis. In this chapter, I chose to focus on Experiment 6 of (Bichot et al., 1999). I will proceed to describe it in some detail, to finally show that it fails to provide the evidence that it claims to provide.

b) The Bichot et al 1999 study

(1) Methods

Quoting directly from the methods section in (Bichot et al., 1999) – the display is shown on Figure 33:

[...] Each primary stimulus contained a circular array of eight shapes, two of which were of the target color and the remaining six of the distractor color. [...] The primary stimulus shapes were either unfilled circles (2.1° visual angle diameter) or unfilled squares (2.1° visual angle across). In half the trials, the two target shapes were the same (and of these trials, half had two circles and half had two squares); in

the other half of the trials, the two target shapes were different (one was a square and the other a circle). Each distractor shape was chosen randomly. [...] In each trial, the two targets could be either adjacent to one another or separated by one, two, or three distractors. [...] Across trials the two targets were presented at all possible locations on the imaginary circle. The primary task was to determine whether the two target shapes were the same or not. After a fixed stimulus onset asynchrony (SOA) of 105 msec, eight black letters were presented, each centered within one of the primary stimuli. [...] The letters were chosen randomly from all possible consonants, and no two letters were the same. The secondary (probe) task was to report as many of the letters as possible. [...] Both primary stimuli and letters were removed after 60 msec, simultaneous with the presentation of masking stimuli consisting of a pound (#) sign (same size as the letters) at each of the locations previously occupied by the letters. After a 250-msec delay, a display containing all possible letters appeared and, using a mouse, the subjects selected the letters they had seen. They were encouraged to report letters even if they were not certain of their presence [...]. After this response, a question display appeared, asking whether the two target shapes were the same. Furthermore, the subjects were instructed at the beginning of the experiment that their most important task was to accurately determine the match (or the lack of it) between the two target shapes.




Figure 33 Stimulus display and paradigm of experiment 6 in (Bichot et al., 1999)

(2) Results

Subjects performed the primary shape comparison task with an accuracy of 80% (slightly better when target shapes were adjacent). The secondary task is the one that the authors were most interested in. Indeed, the probability of report of the probe letters is a proxy for the amount of attention that was allocated to each location in this paradigm. A simple analysis of the probabilities of report as a function of location (Figure 34a) showed that attention was preferentially allocated to target locations over distractor locations. For further analyses, Bichot used conditional probabilities. Given that a letter at a target location was reported, the other target location still enjoyed a greater letter report probability than distractor locations, as shown in Figure 34b. Finally, and critically, letter report at distractor locations intervening between the two targets was no better than at other distractor locations.

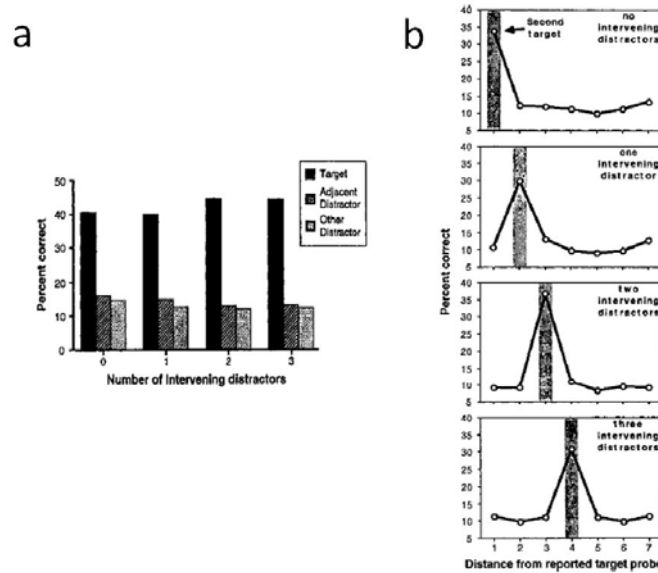


Figure 34 a Probability of letter report as a function of attentional deployment. **b** Conditional probability of letter report as a function of the distance to the correctly reported target letter. Reproduced from (Bichot et al., 1999)

As the letter presentation time (60 ms) was not deemed long enough for the spotlight to switch from one location to another, the logical conclusion was that the spotlight must have been divided between the two targets without encompassing distractor locations between them. The evidence appears flawless, indeed. However, it relies on an inadequate use of conditional probabilities. To understand what went wrong, let us examine how Bichot computed conditional probability estimates. For each trial, when a letter at a target location was correctly reported, they considered each other location in turn (the location adjacent to the correctly reported target and closest to the second target was assigned position 1, and position number was incremented moving away from the correctly reported target around the circle) and incremented a counter if the letter at that location was correctly reported. In trials with two correctly reported targets, they applied this procedure twice, i.e., for each target in turn. After going through all the trials, the final counts for each position were divided by the number of correctly reported target letters to obtain the conditional probability estimates. In this procedure, the conditional probability of reporting the other target letter given that one target letter is reported is thus given by $2 \times N(\text{BOTH}) / (N(\text{T1}) + N(\text{T2}))$, which can also be written as $2 \times P(\text{BOTH}) / (P(\text{T1}) + P(\text{T2})) = 2 \times P(\text{T1}) \times P(\text{T2}) / (P(\text{T1}) + P(\text{T2}))$ under the assumption of independence. Now consider a situation in which no attentional split occurs, so that only one of the two targets receives attention (letter report probability of 0.6) while the other is processed at the same level as a distractor (letter report probability of 0.2). This yields an overall probability for reporting a letter at a target location of 0.4. The calculation performed by Bichot et al. based on conditional probabilities would conclude that, given that a letter was correctly reported at one target location, there is a 0.3 probability ($= 2 \times 0.6 \times 0.2 / (0.6 + 0.2)$) of correctly reporting a letter at the other target location, against 0.2 at a distractor location. Hence, they would reach the erroneous conclusion that both targets simultaneously received attentional resources above the distractor level. The argument put forward in this study is thus inconclusive, and the possibility remains that attention was not divided between the two simultaneous target locations in this paradigm. But we thought of a proper way to analyze the data collected in this paradigm and draw proper conclusions...

2. PAPER 1 : Dubois, Hamker and VanRullen (2009) - Journal of Vision

Attentional selection of noncontiguous locations: The spotlight is only transiently “split”

Julien Dubois

Universite de Toulouse, Universite Paul Sabatier, &
CNRS, UMR5549, Faculte de Medecine de Rangueil,
Toulouse, France, &
California Institute of Technology, Pasadena, CA, USA



Fred H. Hamker

Institute of Psychology and Otto-Creutzfeldt Center for
Cognitive and Behavioral Neuroscience, Westfälische
Wilhelms-Universität, Münster, Germany, &
Artificial Intelligence, TU Chemnitz, 09107 Chemnitz,
Germany



Rufin VanRullen

Universite de Toulouse, Universite Paul Sabatier, &
CNRS, UMR5549, Faculte de Medecine de Rangueil,
Toulouse, France, &
California Institute of Technology, Pasadena, CA, USA



It is still a matter of debate whether observers can attend simultaneously to more than one location. Using essentially the same paradigm as was used previously by N. P. Bichot, K. R. Cave, and H. Pashler (1999), we demonstrate that their finding of an attentional “split” between separate target locations only reflects the early phase of attentional selection. Our subjects were asked to compare the shapes (circle or square) of 2 oddly colored targets within an array of 8 stimuli. After a varying stimulus onset asynchrony (SOA), 8 letters were flashed at the previous stimulus locations, followed by a mask. For a given SOA, the performance of subjects at reporting letters in each location was taken to reflect the distribution of spatial attention. In particular, by considering the proportion of trials in which none or both of the target letters were reported, we were able to infer the respective amount of attention allocated to each target without knowing, on a trial-by-trial basis which location (if any) was receiving the most attentional resources. Our results show that for SOAs under 100–150 ms, attention can be equally split between the two targets, a conclusion compatible with previous reports. However, with longer SOAs, this attentional division can no longer be sustained and attention ultimately settles at the location of one single stimulus.

Keywords: attention, divided spotlight, time course, premotor theory

Citation: Dubois, J., Hamker, F. H., & VanRullen, R. (2009). Attentional selection of noncontiguous locations: The spotlight is only transiently “split”. *Journal of Vision*, 9(5):3, 1–11, <http://journalofvision.org/9/5/3/>, doi:10.1167/9.5.3.

Introduction

The processing of a full scene in parallel up to a high level description of each of the objects it contains is an overwhelming task for our brains. The mechanism in charge of reducing the perceptual load according to behavioral constraints is termed “selective attention.” It is not clear at which stage of sensory processing it operates (Driver, 2001). A classical metaphor for selective spatial attention is that of a spotlight shining on some part of the visual field (Crick, 1984; Eriksen & Eriksen, 1974; Eriksen & St James, 1986; Eriksen & Yeh, 1985; LaBerge, 1983; Posner, Snyder, & Davidson, 1980; Tsal, 1983). This widely used image naturally leads one to wonder whether the attentional spotlight is necessarily undivided or whether it can be split,

thereby allowing the brain to process in parallel two or more noncontiguous foci. In recent years, evidence that multiple locations can be simultaneously attended has accumulated (Awh & Pashler, 2000; Baldauf & Deubel, 2008; Bichot, Cave, & Pashler, 1999; Carlson, VanRullen, Hogendoorn, Verstraten, & Cavanagh, 2007; Gobell, Tseng, & Sperling, 2004; Godijn & Theeuwes, 2003; Kraft et al., 2005; Kramer & Hahn, 1995; McMains & Somers, 2004, 2005; Müller, Malinowski, Gruber, & Hillyard, 2003; Pylyshyn & Storm, 1988), though most studies lack the temporal resolution necessary to distinguish between a truly split spotlight enhancing multiple locations in parallel and an undivided one switching quickly between concurrent locations (Townsend, 1990; VanRullen, Carlson, & Cavanagh, 2007).

A seemingly convincing argument in favor of a divided spotlight was made by Bichot et al. (1999). They

presented human subjects with eight shapes arranged on a circle, two of which—the target shapes—were of a different color than the remaining distractor shapes. After a delay (105 ms), they briefly (60 ms) presented letters inside the shapes, followed by a mask. The primary task of the subjects involved a shape judgment on the target shapes, ensuring that attention was effectively directed to them. The secondary task was to report as many letters as possible from distractor and target locations alike. The probability of correctly reporting the letter inside a shape was used as a proxy for the amount of spatial attention that was allocated to that location. The authors reported that:

1. attention was preferentially allocated to target locations over distractor locations;
2. given that a letter at a target location was reported, the other target location still enjoyed a greater letter report probability than distractor locations;
3. letter report at distractor locations intervening between the two targets was not better than at other distractor locations.

As the letter presentation time (60 ms) was not deemed long enough for the spotlight to switch from one location to another, they concluded that the spotlight must have been divided between the two targets without encompassing distractor locations between them.

While their argument appears strong, it relies in part on an inadequate use of conditional probabilities. The reasoning flaw is best understood through a numerical example. Consider a situation in which no attentional split occurs, so that only one of the two targets receives attention (letter report probability of 0.6) while the other is processed at the same level as a distractor (letter report probability of 0.2). This yields an overall probability for reporting a letter at a target location of 0.4. The calculation performed by Bichot et al. (1999) based on conditional probabilities would conclude that, given that a letter was correctly reported at one target location, there is a 0.3 probability ($= 2 * 0.6 * 0.2 / (0.6 + 0.2)$, see [Methods](#) section) of correctly reporting a letter at the other target location, against 0.2 at a distractor location. Hence, they would reach the erroneous conclusion that both targets simultaneously received attentional resources above the distractor level. To summarize, the argument put forward by Bichot et al. (1999) is inconclusive, and the possibility remains that attention was not divided between the two simultaneous target locations in this paradigm. There exists however a proper way to analyze data collected with this paradigm, which determines the probability of letter reports at the most and least attended target locations, not from conditional probabilities but through the resolution of a simple second-degree equation (see [Methods](#) section). With access to these two probabilities, conclusions can be drawn about the integrity or split of the spotlight. In this

paper, we present a replication of Bichot et al.'s (1999) experiment, which we analyze in this novel way; in addition, by varying the delay between the presentation of the shapes and the appearance of the letters, we are able to track the deployment of the attentional spotlight(s) over time. We find that for delays under 100–150 ms, attention can be equally split between the two targets, a conclusion compatible with previous reports. However, with longer delays, attention ultimately settles at the location of one single stimulus.

Methods

Subjects

Nineteen subjects participated in this study: 8 females, age range 18–32, mean 22.3; 11 males, age range 22–34, mean 26.2. Four out of our 19 subjects performed an insufficient number of trials (less than 25 trials per SOA were left after all discards mentioned later on). The following analysis thus includes only 15 subjects.

All subjects had normal or corrected-to-normal vision, including normal color perception. Subjects provided informed consent before starting the experiment and received financial compensation afterward. The experimental procedure received IRB approval from the California Institute of Technology.

Apparatus

Subjects were seated in a dark room specially designed for psychophysics experiments. The eyes were approximately 120 cm from a computer screen ($1,024 \times 1,286$ pixels, 3×8 bit RGB), connected to a Silicon Graphics (Mountain View, CA) O2 computer running a custom-made OpenGL software. The refresh rate of the monitor was 75 Hz, and the display was synchronized with the vertical retrace of the monitor. Black and white luminances were set to the minimum and maximum values achievable by the monitor.

Experimental procedure

We closely followed the paradigm used by Bichot et al. (1999). Before starting the experiment, subjects were given written instructions and a number of practice trials. Each subject was given a particular target color (red or green): 8 subjects searched for red targets among green distractors, 11 searched for green among red. Subjects had 1 hour to complete the task, including the practice trials (except for subjects who were tested on 7 SOAs, who completed two

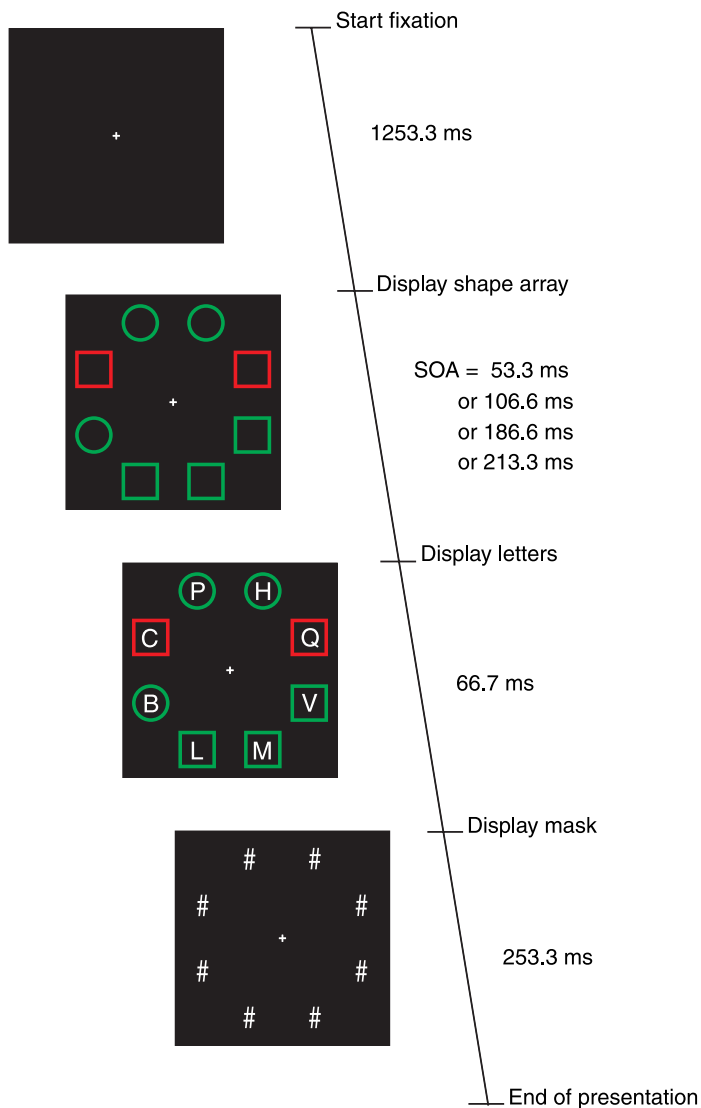


Figure 1. Stimulus sequence for one trial. The primary task consisted in comparing the shapes of the two odd-colored items (targets). The secondary task involved reporting as many letters as possible from the subsequent array. The letter report probability was used as an indication of the amount of attention allocated to the corresponding location. The use of variable stimulus onset asynchronies (SOAs) between the onsets of the shape and letter arrays allowed us to investigate the time course of attentional deployment.

1-hour sessions). Therefore, different numbers of trials were collected for the different subjects, and the different trial types were fully randomized. The sequence of displays in this experiment is shown in Figure 1.

The background was black. Each trial began with the presentation of a central fixation cross that the subjects were instructed to fixate during the entire duration of the trial. The instructions explicitly warned against eye movements. After a delay of 1200 ms, the primary

stimulus display was presented. It contained a circular array of eight shapes, two of which were of the target color and the remaining six of the distractor color. The shapes were equally spaced on an imaginary circle with a radius of 4.8 degrees of visual angle centered on the fixation cross. To avoid potential confounds associated with having stimuli on horizontal and vertical midlines, all the shapes were placed so that each shape was 22.5 degrees off a horizontal or vertical midline. The primary stimulus shapes were either unfilled circles (2.1 degrees visual angle diameter) or unfilled squares (2.1 degrees visual angle across). In half the trials, the two target shapes were the same (and of those trials, half had two circles and half had two squares); in the other half of the trials, the two target shapes were different (one was a square and the other was a circle). Among the six distractor shapes, three were circles and three were squares, randomly assigned. In each trial, the two targets could be either adjacent to one another or separated by one, two, or three distractors. All four target separations were equally likely to occur, and across trials the two targets were presented at all possible locations on the imaginary circle. The target separations, as well as the target positions, were selected randomly across trials so that the subjects could not anticipate a particular display configuration. The primary task was to determine whether the two target shapes were the same or not. It was emphasized that the primary task had to be performed correctly for the trial not to be discarded. Since a failure to perform the shape comparison task might potentially reveal an attentional wandering, only trials for which the subjects correctly performed the primary task were considered for further analysis.

After a variable stimulus onset asynchrony (SOA), eight white letters were presented, each centered within one of the shapes. We started out with seven SOAs: 0, 40, 53.3, 80, 106.6, 160, 186.6, and 213.3 ms. After running 7 subjects and looking at the data, we reduced this number to four SOAs: 53.3, 106.6, 186.6, and 213.3 ms in order to limit the required number of trials per subject. Only the 4 SOAs that were common to all subjects are considered in the following analysis. The responses to a probe at a particular SOA may vary depending on which other SOAs the subject is exposed to on other trials; likewise, extra practice or fatigue may entail differences in the responses across these two subject groups. We initially performed all subsequent analyses with a factor for the subject's group. As no main effect of the group nor any interaction was ever found, we decided to consider that all subjects came from the same population, and dropped the factor group, so as to provide the reader with a less complex picture. The letters were chosen from all possible consonants but V, W, X, and Z for the 7-SOA group, and from all possible consonants for the 4-SOA group; no two letters were the same in a given trial. Each letter subtended a visual angle of 0.6 degrees vertically and 0.4 degrees horizontally. The secondary task was to report as many letters as possible. Both shapes and

letters were removed 66.7 ms after letter onset, simultaneous with the presentation of masking stimuli consisting of a pound (#) sign (same size as letters) at each of the locations previously occupied by the letters. Masking stimuli were used to prevent further processing of the letters from iconic memory after their removal. After a 240-ms delay, a display containing all possible letters appeared on the left of the screen, and using a mouse, the subjects selected any number of letters they had seen, from zero to eight. They were encouraged to report letters even if they were not certain of their presence but were warned against any wild guesses. After clicking on the letters (secondary task), they had to click on one of two buttons presented on the right-hand side of the display, “same” or “different,” referring to the target shapes (primary task). The importance of accuracy over speed was emphasized.

Probability estimates

In the following, let T1 be the event “the letter at the most attended target location is reported” and let T2 be the event “the letter at the least attended target location is reported” (of course, this definition is only for mathematical purposes and does not preclude the possibility that in practice both target locations might receive comparable amounts of attention). We assume that the probability of reporting a behaviorally relevant target only depends on the amount of attention that it receives. It follows that the events T1 and T2 can be considered independent. Let BOTH be the event “both target letters are reported,” and let NONE be the event “none of the target letters is reported.”

Conditional probabilities

Let us first examine how Bichot et al. (1999) computed conditional probability estimates. For each trial, when a letter at a target location was correctly reported, they considered each other location in turn (the location adjacent to the correctly reported target and closest to the second target was assigned position 1, and position number was incremented moving away from the correctly reported target around the circle) and incremented a counter if the letter at that location was correctly reported. In trials with two correctly reported targets, they applied this procedure twice, i.e., for each target in turn. After going through all the trials, the final counts for each position were divided by the number of correctly reported target letters to obtain the conditional probability estimates. In this procedure, the conditional probability of reporting the other target letter given that one target letter is reported is thus given by $2 \cdot N(\text{BOTH}) / (N(\text{T1}) + N(\text{T2}))$, which can also be written as $2 \cdot P(\text{BOTH}) / (P(\text{T1}) + P(\text{T2})) = 2 \cdot P(\text{T1}) \cdot P(\text{T2}) / (P(\text{T1}) + P(\text{T2}))$. As described in the [Introduction](#) section, this procedure can artificially foster the conclusion that both targets are simultaneously

attended, even when only one of them actually receives prioritized processing.

Target locations

Our improved analysis method, described below, does not rely on the use of conditional probabilities. For simplicity, $P(\text{T1})$ and $P(\text{T2})$ are assumed to be constant across trials (and potentially equal, if the hypothesis of an attentional split is verified); but the identity of the most and least attended locations is unknown on any given trial. Can we still get the probability estimates for the two events $P(\text{T1})$ and $P(\text{T2})$?

The following equations hold:

$$P(\text{BOTH}) = P(\text{T1} \cap \text{T2}) = P(\text{T1}) \cdot P(\text{T2}), \quad (1)$$

$$\begin{aligned} P(\text{NONE}) &= P(\sim \text{T1} \cap \sim \text{T2}) \\ &= (1 - P(\text{T1})) \cdot (1 - P(\text{T2})) \\ &= 1 + P(\text{T1}) \cdot P(\text{T2}) - (P(\text{T1}) + P(\text{T2})). \end{aligned} \quad (2)$$

Probabilities $P(\text{BOTH})$ and $P(\text{NONE})$ can be estimated directly from the experimental data—unfortunately, the corresponding values were not reported in the study by Bichot et al. (1999), preventing a reevaluation of their conclusions.

Let $\Sigma = P(\text{T1}) + P(\text{T2})$ and $\Pi = P(\text{T1}) \cdot P(\text{T2})$. [Equations 1](#) and [2](#) can be rewritten as

$$\Pi = P(\text{BOTH}), \quad (3)$$

$$\Sigma = 1 + P(\text{BOTH}) - P(\text{NONE}). \quad (4)$$

Here we appeal to a general theorem of polynomial equations: for the equation $a_n x^n + a_{n-1} x^{n-1} + a_{n-2} x^{n-2} + \dots + a_0 = 0$, with $a_n \neq 0$, the sum of the roots is $-a_{n-1}/a_n$ and the product of the roots is a_0/a_n if n is even, and $-a_0/a_n$ if n is odd. From this theorem and [Equations 3](#) and [4](#), $P(\text{T1})$ and $P(\text{T2})$ are the solutions of the following second-degree equation:

$$X^2 - \Sigma \cdot X + \Pi = 0. \quad (5)$$

Solving [Equation 5](#) requires calculating its discriminant Δ , defined as

$$\Delta = \Sigma^2 - 4 \cdot \Pi. \quad (6)$$

If the discriminant Δ is zero, Equation 5 only has one solution and $P(T1) = P(T2)$: attention is equally allocated to the two targets. If Δ is positive, there are two distinct real solutions; the greatest solution is assigned to $P(T1)$ and the other solution to $P(T2)$: the attentional spotlight favors one of the two locations. Last, if Δ is negative, there are no real solutions to Equation 5—the solutions are complex numbers. If this were to happen with our data, it would likely be due to noise, as probabilities are real numbers: to obtain the solutions of the equation in this case, we artificially set Δ to zero; however, the proper (negative) value of Δ was retained in order to compare the distribution of discriminant values over subjects with the null hypothesis of $\Delta = 0$, as described below.

To determine whether there is an attentional split, we perform the estimation of $P(T1)$ and $P(T2)$ as explained above for each subject. The null hypothesis is that there is an equal split of attentional resources, i.e., that Δ is zero. Since we are dealing with experimental data, even if the null hypothesis is true we do not expect every subject's Δ to be exactly zero, but rather we expect individual Δ s to be distributed around zero. Our criterion for an attentional split thus consists in testing the mean of the distribution of Δ s against zero, using a one-sided t -test. If Δ is significantly positive ($p < 0.05$), the equal attentional split interpretation cannot hold any more.

It is important to note that subject-by-subject estimates will artificially separate the mean values of $P(T1)$ and $P(T2)$. Indeed, consider a situation in which $P(T1) = T + \varepsilon_1$ and $P(T2) = T + \varepsilon_2$, where T is the true probability of detecting a target and ε_1 and ε_2 are normally distributed error terms across subjects. The true means of $P(T1)$ and $P(T2)$ are the same, implying that attention is equally allocated to the two targets (in other words, $\Delta = 0$). For each subject, however, our method estimates two probability values and always assigns the greatest to $P(T1)$. The final mean estimates of $P(T1)$ and $P(T2)$ across subjects will thus artificially differ, even though they are drawn from distributions with the same mean. Therefore, we always refer to the distribution of Δ s first, before making any conclusions regarding $P(T1)$ and $P(T2)$: even when the mean values across subjects of $P(T1)$ and $P(T2)$ significantly differ, the hypothesis of an equal split of attention can only be rejected when Δ is significantly positive.

Finally, another way to look at the data is to pool all trials as if they came from a single subject and perform the estimation of $P(T1)$ and $P(T2)$ for this pooled data. Assuming that each subject's data is drawn from the same population distribution, we then obtain a better estimate of the underlying probabilities (i.e., smaller error terms ε_1 and ε_2); the estimates will be more reliable and less susceptible to the above-mentioned bias, providing a less misleading picture than the mean of individual estimates (with the obvious caveat that no estimate of inter-subject variability can be computed in this case). In the following, we analyze and report both the pooled data and the means of individual values across subjects.

Distractor locations

As described by Bichot et al. (1999), it is interesting to compare performance at distractor locations that are between the targets to performance at distractor locations that are outside the targets (“inside” and “outside” can only be defined when the two targets are separated by one or two distractors; when the two targets are separated by zero or three distractors, we put all distractors in the “outside” category); indeed if a single spotlight encompasses the two target locations and what is in the midst of them, then performance at distractor locations *between* targets should be greater than at distractor locations *outside* of targets.

Chance level

The average probability for a subject to report a given letter by chance on one trial depends on the number of letters they report. Say a subject reports three letters per trial on average, among the 20 possible letters. The number of three-letter reports comprising the letter of interest is $\binom{19}{2}$. The total number of three-letter reports is $\binom{20}{3}$. Hence, the probability for that subject to report a letter by chance is on average $\binom{19}{2} / \binom{20}{3} = 0.15$ (the same




Figure 2. Letters in squares were more likely to be reported than letters in circles (data collapsed over SOAs).

calculation applied to the 7-SOA group, for which only 16 letter choices were given, would yield a chance level of $\binom{15}{2} / \binom{16}{3} = 0.1875$.

Results

Primary task (shape comparison)

All subjects performed well on the primary task (average correct: 86.7%, standard error of the mean 1.8%). As the SOA increased, performance on the task got slightly but significantly better (82.4% correct at 53.3 ms, 86.0% at 106.6 ms, 88.0% at 160 ms, and 90.6% at 213.3 ms; 1-way ANOVA for the effect of SOA, $F(3, 42) = 13.88$, $p < 10^{-5}$). Trials in which subjects responded incorrectly were discarded.

Secondary task (letter report)

We observed a stimulus-driven bias: letters presented inside squares were, on average, more likely to be reported than letters presented inside circles, as shown in Figure 2 (2-way ANOVA, $F(1, 14) = 74.06$, $p < 10^{-6}$ for [Circle/Square], $F(1, 14) = 59.65$, $p < 10^{-5}$ for [Target/Distractor], with a significant interaction $F(1, 14) = 34.04$, $p < 10^{-4}$ showing that the effect is more pronounced for targets than for distractors). This bottom-up advantage for squares could have lead to spurious conclusions regarding

the distribution of attention on trials in which the target shapes were different. We therefore only considered trials in which the target shapes were identical for subsequent analysis: that way, both target locations were expected to receive comparable amounts of exogenous attention.

First, we replicated the conditional probability results presented as the main evidence for a split attentional spotlight by Bichot et al. (1999), as shown in Figures 3 and 4. In Figure 3, we observed that, on average, the probability of reporting a letter at a target location was higher than at distractor locations, and this difference increased significantly with increasing SOA (2-way ANOVA, $F(1, 14) = 74.36$, $p < 10^{-6}$ for [target/distractor] and $F(3, 42) = 7.16$, $p < 10^{-3}$ for [SOA], with a significant interaction $F(3, 42) = 8.31$, $p < 10^{-3}$). Letters at distractor locations were reported significantly less often than they would have been by chance (2-way ANOVA, $F(1, 14) = 7.57$, $p < 0.05$ for [distractor/chance] and $F(3, 42) = 0.44$, $p = 0.73$ for [SOA], with a significant interaction $F(3, 42) = 3.37$, $p < 0.05$). This is due to the manner chance is calculated, depending on the average number of letters reported by each subject, and it shows that the primary task was effective in making subjects attend selectively and almost exclusively to the target locations. Distractor locations that were between target locations did not receive more attention than other distractor locations at any SOA (2-way ANOVA, $F(1,14) = 2.85$, $p = 0.11$ for [between/outside] and $F(3, 42) = 0.46$, $p = 0.71$ for [SOA], with no significant interaction $F(3, 42) = 0.69$, $p = 0.57$), thus arguing against the possibility of a single extended spotlight encompassing both target locations. In Figure 4, probabilities of letter report conditional on the




Figure 3. Probability of correct letter report as a function of SOA. Letters at target locations (red) are reported more often than letters at distractor locations (blue and cyan). Furthermore, performance is not better at distractor locations “between” (blue) than at distractor locations “outside” (cyan) of the two targets.

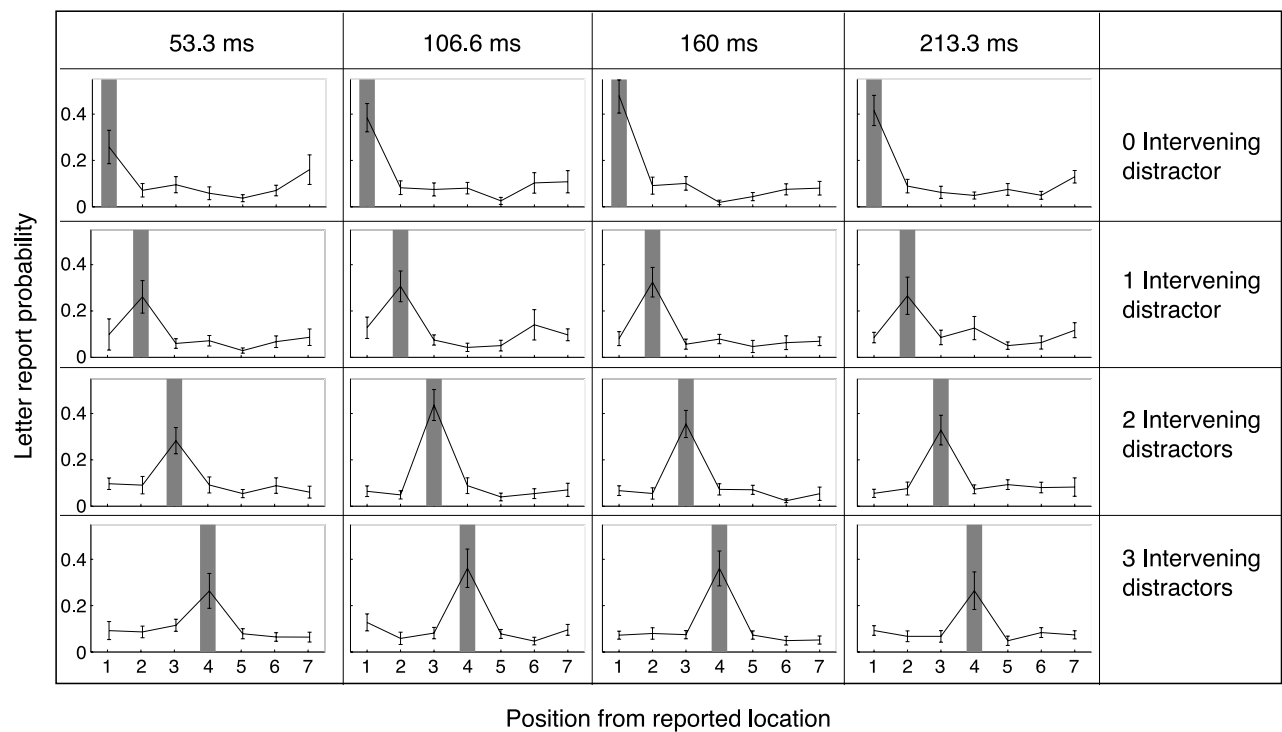


Figure 4. Conditional probabilities (computed as in Bichot et al., 1999). Each graph represents the probability of letter report as a function of the positional distance from a target letter that was correctly reported. Position 1 is the shape right next to the correctly reported target location. The shaded gray bar indicates the location of the second target. Positions are counted around the circle so that the second target always falls within positions 1 to 4.

report of one target are plotted. The conditional probability of reporting the other plotted was significantly higher than the conditional probability of reporting a distractor, and this difference significantly increased with SOA, which reflects the previous observation that target letter report performance increases with SOA (see Figure 3; 3-way ANOVA, $F(1, 14) = 22.99$, $p < 10^{-3}$ for [target/distractor], $F(3, 42) = 4.12$, $p < 0.05$ for [SOA], and $F(3, 42) = 3.46$, $p < 0.05$ for [target separation], with a significant interaction between [target/distractor] and [SOA] $F(3, 42) = 4.63$, $p < 0.01$). Likewise, in the cases of one or two intervening distractors, the conditional probability of report for distractors between targets is not significantly different from the conditional probability of report for distractors outside targets at any SOA or any target separation (3-way ANOVA, $F(1, 14) = 0.752$, $p = 0.40$ for [distractor between/distractor outside], $F(3, 42) = 0.57$, $p = 0.64$ for [SOA], and $F(1, 14) = 1.98$, $p = 0.18$ for [target separation], with no significant interactions). At first sight, our data are thus compatible with the data that Bichot et al. (1999) based their conclusions on.

We then analyzed the results with our improved calculation procedure, as described in the Methods section (Figure 5). We estimated the probabilities P(BOTH) and P(NONE) for each subject in turn, which allowed us to solve a simple second-degree equation and estimate P(T1) and P(T2), the probabilities of correct letter report at the

most and least attended target locations, respectively. Because potential measurement errors across individual subjects do not cancel out but instead add up in this analysis (due to our systematically assigning the largest of the 2 computed probabilities to P(T1), as explained in detail in the Methods section), we also performed the same analysis on pooled data, using all trials from all subjects as though they came from a single subject. This pooled calculation, illustrated by thick lines in Figure 5, will be less prone to measurement error. Only trials in which there was at least one intervening distractor were considered; indeed, when the target locations are adjacent, the spotlight need not be divided to illuminate both target locations simultaneously.

Critically, the spatial distribution of attention between the two target locations was found to depend on SOA (2-way ANOVA, $F(3, 42) = 9.50$, $p < 10^{-4}$ for [SOA], $F(1, 14) = 50.181$, $p < 10^{-5}$ for [T1/T2] with a significant interaction $F(3, 42) = 3.95$, $p < 0.05$). This result goes together with a significant effect of SOA on the discriminant (1-way ANOVA, $F(3, 42) = 3.25$, $p < 0.05$). A post-hoc analysis using Tukey's Honestly Significant Differences test shows that the discriminant at the longest SOA (213.3 ms) is significantly higher than at the shortest SOAs (53.3 ms and 106.6 ms). We can already conclude that the bias of attention for one target location over the other increases with time. Further tests show that for

SOAs at or below 106.6 ms, the two target locations apparently received comparable amounts of attention: the discriminants across subjects did not statistically differ from zero (one sided t -test, $t(14) = 1.27$, $p = 0.11$ at 53.3 ms and $t(14) = 0.93$, $p = 0.18$ at 106.6 ms); the discriminant estimated from the pooled data was close to zero (-0.01 at 53.3 ms and -0.01 at 106.6 ms), providing graphical confirmation. Note that this result is consistent with the conclusions of Bichot et al. (1999), who investigated a single SOA of 105 ms. For SOAs of 160 ms or greater, one of the target locations was found to win the attentional competition, receiving more attention as the SOA

increased: the discriminants across subjects were statistically larger than zero (one sided t -test, $t(14) = 3.54$, $p < 0.01$ at 160 ms and $t(14) = 3.36$, $p < 0.01$ at 213.3 ms) and the discriminant estimated from the pooled data was larger than zero (0.01 at 160 ms and 0.04 at 213.3 ms). These results are summarized in Figure 5.

Discussion

In this experiment, two target locations were cued with odd-color shapes and were behaviorally relevant for the subject's primary task (shape comparison). Letters were briefly presented inside the shapes, and we estimated the probabilities of correct letter report at the most and least attended target locations as a proxy for the amount of spatial attention that each of these locations received. It turned out that for short stimulus onset asynchronies (the time between the presentation of the cues and the appearance of the letters inside them) the probabilities of correct letter report at the two target locations were comparable; that is, resources appeared to be equally allocated to the two concurrent target locations. However, this early state evolved into the selection of one target location over the other when the stimulus onset asynchrony increased: one of the spotlight's initial two beams died out as the other beam gradually used more resources. Our results thus point to two successive modes for the deployment of attention in space: an early, transient mode in which multiple locations can be attended and a later, maybe more sustained mode in which attention needs to focus on a single location at a time.

Given our procedure and instructions, it is difficult to determine whether the odd-color feature cues acted in an exogenous (bottom-up) or endogenous (top-down) manner. Some evidence to this effect can be derived from a pilot study (as yet unpublished) in which 8 participants performed the experiment without a primary shape

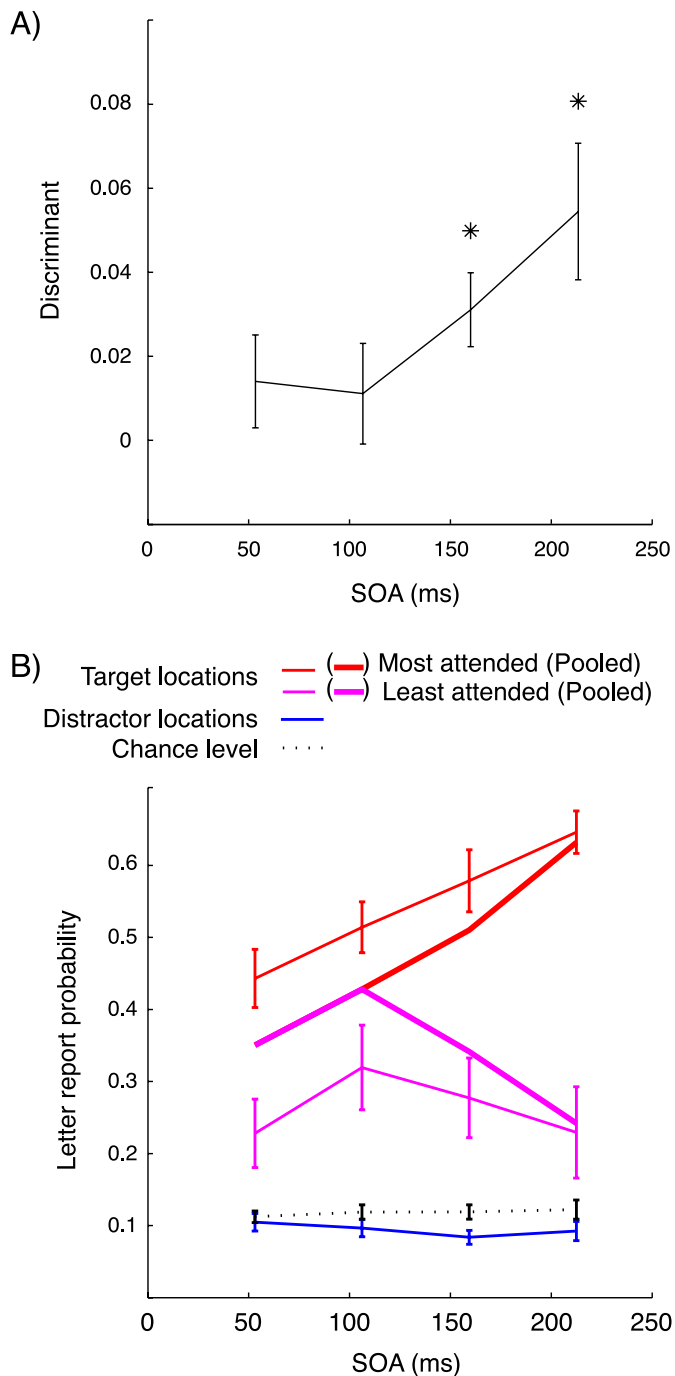


Figure 5. Deployment of the spotlight of attention as a function of SOA. (A) The average discriminant across subjects is not significantly different from zero for SOAs shorter or equal to 106.6 ms, whereas it is significantly above zero for SOAs longer or equal to 160 ms. This implies that, at least for the two largest SOAs, the conclusion that the 2 targets receive equal amounts of attention is not warranted. (B) The probability of report at the most attended location keeps increasing with increasing SOAs, whereas the probability of report at the least attended location starts decreasing: attention focuses most of its resources on one target at longer SOAs. The pooled (thick lines) and individual data (thin lines), respectively, provide a more conservative and a more liberal estimate of the "split attention" hypothesis. In both cases, however, this hypothesis is found not to hold for SOAs larger than 106 ms.

comparison task; preliminary analysis revealed that the two odd-color shapes in this case were treated no differently from the distractors, indicating that our present findings may be limited to behaviorally relevant targets.

One may argue that the observed selection of one single location at SOAs between 107 ms and 160 ms is an artifact of saccades that subjects made to one of the targets, despite the clear instructions to not move their eyes. We did not record eye movements so we do not have definite evidence to reject this argument. Average saccadic latencies between 180 and 250 ms have been reported in humans in experimental conditions similar to ours, i.e., with no offset of the fixation cross prior to the saccade (McPeck, Maljkovic, & Nakayama, 1999; Song & Nakayama, 2007; Yang, Bucci, & Kapoula, 2002). The duration of an ~ 5 degree amplitude saccade is in the range of 30–45 ms (Carpenter, 1988). It would thus take on average 210 to 300 ms for subjects to fixate one of the targets if they were instructed to do so. We argue that the observed selection of a single location between 107 ms and 160 ms occurs too early to be accounted for by the execution of an eye movement. It is reasonable to suppose, however, that the oculomotor system may have been *planning* a saccade. Indeed, our results are a natural prediction of a certain class of computational models of attention (Hamker, 2004) in which the planning of saccadic eye movements guides attentional selection, in line with the premotor theory of attention (Rizzolatti, Riggio, Dascola, & Umiltà, 1987). In this model of saccadic decision making and attention, the frontoparietal network receives the current output of a “refined” distributed saliency map (Hamker, 2006) and selects the unique location of an eye movement by a competition over time. Activity from this frontoparietal network is fed back continuously to extrastriate visual areas. Thus, the SOA determines the state of this competition at the time the letters are flashed, and ultimately the distribution of attention at different locations. While the amount of motor in the premotor theory has been a subject of intense debate (Chambers & Mattingley, 2005; Hamker, 2005; Juan et al., 2008; Juan, Shorter-Jacobi, & Schall, 2004; Thompson, Biscoe, & Sato, 2005) our results argue for feedback signals from cells in oculomotor areas that have knowledge about the motor plan.

In light of our previous distinction between transient and sustained attention modes, it is likely that some of the evidence reported in favor of a split of the spotlight only addressed the transient attention mode (Baldauf & Deubel, 2008; Bichot et al., 1999; Carlson et al., 2007; Godijn & Theeuwes, 2003; Kramer & Hahn, 1995). Other reports that clearly pertained to the sustained attention mode (Awh & Pashler, 2000; Gobell et al., 2004; Kraft et al., 2005; McMains & Somers, 2004, 2005; Müller et al., 2003; Pylyshyn & Storm, 1988) may not have possessed the temporal resolution necessary to distinguish between multiple spotlights and a unique, rapidly switching one. In particular, the main argument to discard the switching

spotlight model in the SSVEP study by Müller et al. (2003) and in the fMRI studies by McMains and Somers (2004, 2005) is that of limited stimulus presentation times (around 170–180 ms). But in a recent study involving sustained monitoring of multiple locations, VanRullen et al. (2007) reported evidence for an undivided spotlight in constant periodic motion, sampling its multiple targets at a rate of seven items per second. Of course, there might be circumstances outside of those tested in the present study in which a split attentional spotlight may be sustained. In any case however, our results suggest that the seemingly disparate literature on attentional allocation to multiple targets could be reconciled by considering time (transient or sustained attention mode) as a critical factor.

Acknowledgments

Part of this work has been presented earlier in abstract form (Hamker, F. H., & VanRullen, R. (2002). The time course of attentional selection among competing locations [Abstract]. *Journal of Vision*, 2(7):7, 7a, <http://journalofvision.org/2/7/7/>). This work was supported by the German Science Foundation (DFG HA2630/2-1) and the Federal Ministry of Education and Research Grant (BMBF 01GW0653), the CNRS (Grant CNRS-USA), the ANR (06JCJC-0154), the Fyssen Foundation, and the European Science Foundation (EURYI). We wish to acknowledge Kyle Cave and an anonymous referee for helpful advice and suggestions and Christof Koch for his continued support.

Commercial relationships: none.

Corresponding author: Julien Dubois.

Email: jdubois@klab.caltech.edu.

Address: 1200 East California Boulevard, MC 216-76, Pasadena, CA 91125, USA.

References

- Awh, E., & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 834–846. [[PubMed](#)]
- Baldauf, D., & Deubel, H. (2008). Properties of attentional selection during the preparation of sequential saccades. *Experimental Brain Research*, 184, 411–425. [[PubMed](#)]
- Bichot, N. P., Cave, K. R., & Pashler, H. (1999). Visual selection mediated by location: Feature-based selection of noncontiguous locations. *Perception & Psychophysics*, 61, 403–423. [[PubMed](#)]

- Carlson, T., VanRullen, R., Hogendoorn, H., Verstraten, F., & Cavanagh, P. (2007). Distinguishing models of multifocal attention: It's a matter of time [[Abstract](#)]. *Journal of Vision*, 7(9):641, 641a, <http://journalofvision.org/7/9/641/>, doi:10.1167/7.9.641.
- Carpenter, R. H. S. (1988). *Movement of the eyes*. London: Pion.
- Chambers, C. D., & Mattingley, J. B. (2005). Neuro-disruption of selective attention: Insights and implications. *Trends in Cognitive Sciences*, 9, 542–550. [[PubMed](#)]
- Crick, F. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 81, 4586–4590. [[PubMed](#)] [[Article](#)]
- Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, 92, 53–78. [[PubMed](#)]
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16, 143–149.
- Eriksen, C. W., & St James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception and Psychophysics*, 40, 225–240. [[PubMed](#)]
- Eriksen, C. W., & Yeh, Y. Y. (1985). Allocation of attention in the visual field. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 583–597. [[PubMed](#)]
- Gobell, J. L., Tseng, C. H., & Sperling, G. (2004). The spatial distribution of visual attention. *Vision Research*, 44, 1273–1296. [[PubMed](#)]
- Godijn, R., & Theeuwes, J. (2003). Parallel allocation of attention prior to the execution of saccade sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 882–896. [[PubMed](#)]
- Hamker, F. H. (2004). A dynamic model of how feature cues guide spatial attention. *Vision Research*, 44, 501–521. [[PubMed](#)]
- Hamker, F. H. (2005). The reentry hypothesis: The putative interaction of the frontal eye field, ventrolateral prefrontal cortex, and areas V4, IT for attention and eye movement. *Cerebral Cortex*, 15, 431–447. [[PubMed](#)] [[Article](#)]
- Hamker, F. H. (2006). Modeling feature-based attention as an active top-down inference process. *Biosystems*, 86, 91–99. [[PubMed](#)]
- Juan, C. H., Muggleton, N. G., Tzeng, O. J., Hung, D. L., Cowey, A., & Walsh, V. (2008). Segregation of visual selection and saccades in human frontal eye fields. *Cerebral Cortex*, 18, 2410–2415. [[PubMed](#)]
- Juan, C. H., Shorter-Jacobi, S. M., & Schall, J. D. (2004). Dissociation of spatial attention and saccade preparation. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 15541–15544. [[PubMed](#)] [[Article](#)]
- Kraft, A., Müller, N. G., Hagenhof, H., Schira, M. M., Dick, S., Fendrich, R. M., et al. (2005). Interactions between task difficulty and hemispheric distribution of attended locations: Implications for the splitting attention debate. *Brain Research: Cognitive Brain Research*, 24, 19–32. [[PubMed](#)]
- Kramer, A. F., & Hahn, S. (1995). Splitting the beam: Distribution of attention over noncontiguous regions of the visual field. *Psychological Science*, 6, 381–386.
- LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 371–379. [[PubMed](#)]
- McMains, S. A., & Somers, D. C. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron*, 42, 677–686. [[PubMed](#)] [[Article](#)]
- McMains, S. A., & Somers, D. C. (2005). Processing efficiency of divided spatial attention mechanisms in human visual cortex. *Journal of Neuroscience*, 25, 9444–9448. [[PubMed](#)] [[Article](#)]
- McPeck, R. M., Maljkovic, V., & Nakayama, K. (1999). Saccades require focal attention and are facilitated by a short-term memory system. *Vision Research*, 39, 1555–1566. [[PubMed](#)]
- Müller, M. M., Malinowski, P., Gruber, T., & Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature*, 424, 309–312. [[PubMed](#)]
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, 109, 160–174. [[PubMed](#)]
- Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3, 179–197. [[PubMed](#)]
- Rizzolatti, G., Riggio, L., Dascola, I., & Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, 25, 31–40. [[PubMed](#)]
- Song, J. H., & Nakayama, K. (2007). Fixation offset facilitates saccades and manual reaching for single but not multiple target displays. *Experimental Brain Research*, 177, 223–232. [[PubMed](#)]
- Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *Journal of Neuroscience*, 25, 9479–9487. [[PubMed](#)] [[Article](#)]

- Townsend, J. T. (1990). Serial vs. parallel processing—Sometimes they look like tweedledum and tweedledee but they can (and should) be distinguished. *Psychological Science*, *1*, 46–54.
- Tsal, Y. (1983). Movements of attention across the visual field. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 523–530. [[PubMed](#)]
- VanRullen, R., Carlson, T., & Cavanagh, P. (2007). The blinking spotlight of attention. *Proceedings of the National Academy of Sciences of the United States of America*, *104*, 19204–19209. [[PubMed](#)] [[Article](#)]
- Yang, Q., Bucci, M. P., & Kapoula, Z. (2002). The latency of saccades, vergence, and combined eye movements in children and in adults. *Investigative Ophthalmology & Visual Science*, *43*, 2939–2949. [[PubMed](#)] [[Article](#)]

3. Summary and Discussion

The paradigm pioneered by Bichot and colleagues (Bichot et al., 1999), analyzed in a fashion which allows to truly detect differences in attentional allocation to the two target shapes, warranted the conclusion that attentional resources can not remain spatially divided in a sustained manner. I wish to point out that their conclusions were not erroneous, given that they had tested a SOA of 105ms, which we showed to belong to a transient split mode. However, as SOA increases, the passage to a sustained mode seems to be accompanied by a reunification of the spotlight of attention. It is not the case either that Bichot and colleagues simply neglected the influence of SOA – they rely on an earlier study by Kim and Cave (Kim & K. R. Cave, 1995) using a similar task which showed no significant differences in attentional allocation across SOAs of 60, 105, and 135 ms – the highest SOA in this range, 135ms, arguably still belongs to the transient mode.

Does this mean that spatial attention can truly be transiently divided? We mentioned in the discussion of the article that the initial transient effects may be a signature of exogenous attention – the odd-colored shapes initially pop out, cueing the locations of the targets. There is another interpretation, similar to the point I raised earlier in the case of (Gobell et al., 2004). The cueing of attention in this experiment is done through a feature-based attentional cue – for a given subject, the target shapes are always the same color (red or green, counterbalanced across subjects). Hence, the subject is initially attending to a color, cueing the locations of the to-be-compared shapes. Comparing the shapes of the two targets, as well as reporting letters, requires spatial attention to take over – whether it takes over exogenous or feature-based endogenous attention. What we may be seeing then is a switch from exogenous/endogenous feature-based attention to endogenous spatial attention (which takes roughly 150ms). This is a psychologist's view of the data, which is difficult to back up with empirical evidence.

Fred Hamker, a co-author on this study, has another take on our results. In a recent publication from his group, they write (Zirnsak, Beuth, & Hamker, 2011):

Research on attention has been much influenced by the idea of a spotlight and a multitude of attentional phenomena have often been described in qualitative terms. However, without a close consideration of the neural correlates, attention cannot really be constrained and arbitrary shaped foci can be taken into account such as elliptic foci to argue against a split of attention (e.g. (Jans et al., 2010)). Computational approaches, particularly when linked to neuroanatomical and physiological observations, have the advantage of making more precise assumptions about the underlying mechanisms.

Hamker has developed a detailed model of attention, represented in Figure 35a (a basic description of the model is in the caption of the figure). Feeding a slightly modified version of our experiment to their model, they found that it predicted qualitatively our results when two locations were cued (feedback from IT) – the initial split and the later selection of one of the attended locations (Figure 35b). They therefore argue that our results are a simple consequence of feedback from FEF.

Arguably though, the psychological interpretation of these results may remain unchanged – global feature-based selection turning into spatial selection, which can not sustainably be allocated to two locations.

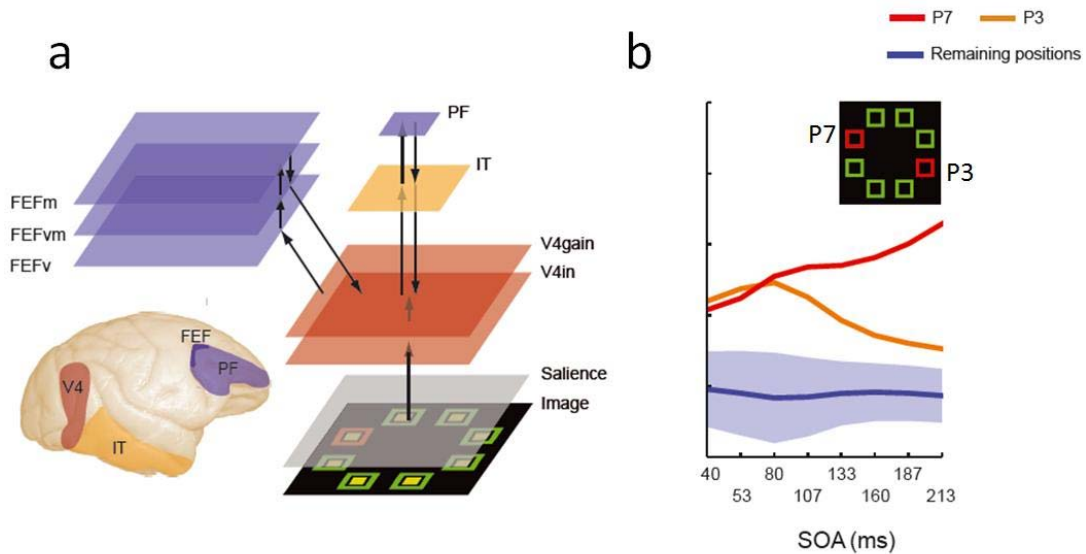


Figure 35 **a** Model layout. The input image is processed by a set of filters resulting in feature conspicuity values across different spatial scales for a given channel, which provides the V4in activity. The activity in V4gain is a function of the bottom-up input provided by V4in and top-down modulations provided by IT (which in turn can be biased by the PF) and the top-down modulations of the FEF. Thereby, the feedback signal from IT is feature specific whereas the feedback signal from the FEF is space specific. The FEF consists of three layers whose cells differ in their response characteristics. Cells in the FEFv layer are responsive to a visual stimulus but show no motor-related activity. Cells in the FEFm layer show no early transient response to visual stimulation but built up their activity during later stages of the selection process if not suppressed by fixation cells. Finally, cells in the FEFvm layer vary along a continuum of stimulus- and motor-related responses. **b** The output of the model when presented with a slightly modified version of our task matches the psychophysical data which we obtained. Adapted from (Zirnsak et al., 2011).

In conclusion, we provided psychophysical evidence against a sustained division of the spotlight of attention, and in favor of the unified model. One natural consequence is that in the case of sustained attention to multiple disjoint locations, the spotlight needs to rapidly oscillate between concurrently attended targets... Hence, this finding can be taken as an argument in favor of a periodic behavior of attention!

D. Pathological behavioral evidence for rhythms in perceptual processes

1. Akinetopsia : a rare and striking disturbance

There are some case reports of neurological patients that anyone who has followed an introductory neuroscience class has heard of : H.M. (anterograde amnesia), G.Y. (blindsight), S.B (visual agnosia), S.M. (no fear), etc... These famous patients present fascinating disorders and have been repetitively studied by various research teams in order to precisely understand the underpinnings of their deficits. One such famous patient, L.M., does not perceive motion. Zihl and colleagues (Zihl, Von Cramon, & Mai, 1983), who studied her, wrote:

The visual disorder complained of by the patient was a loss of movement vision in all three dimensions. She had difficulty, for example, in pouring tea or coffee into a cup because the fluid appeared to be frozen, like a glacier. In addition, she could not stop pouring at the right time since she was unable to perceive the movement in the cup (or a pot) when the fluid rose. Furthermore the patient complained of difficulties in following a dialogue because she could not see the movements of the face and, especially, the mouth of the speaker. In a room where more than two other people were walking she felt very insecure and unwell, and usually left the room immediately, because 'people were suddenly here or there but I have not seen them moving'. The patient experienced the same problem but to an even more marked extent in crowded streets or places, which she therefore avoided as much as possible. She could not cross the street because of her inability to judge the speed of a car, but she could identify the car itself without difficulty. 'When I'm looking at the car first, it seems far away. But then, when I want to cross the road, suddenly the car is very near. She gradually learned to 'estimate' the distance of moving vehicles by means of the sound becoming louder.

The term akinetopsia was coined to describe L.M.'s disorder (Zeki, 1991). Many psychophysical experiments were performed to exactly quantify her deficit (Baker, Robert F Hess, & Zihl, 1991; R. Hess, Baker, & Zihl, 1989; Rizzo, Nawrot, & Zihl, 1995) : contrast sensitivity for motion discrimination with drifting luminance gratings, coherence thresholds to discriminate motion in random dot kinetograms, etc. These are all very interesting, but what really matters to me remains her subjective experience of freeze frames. If you told me that someone does not perceive motion any more, I would not necessarily predict that they should see freeze frames. A naive prediction would be that objects in movement would appear as they do in a long exposure photograph, with ill-defined edges due to visual persistence. Perception of sharp freeze frames is not the most obvious prediction, yet it is what her subjective experience appears to be. The one piece of critical information that I can however not find is the duration of the freeze frames that she experienced (though according to (Sacks, 2004), they were several seconds long – I have no idea where he got that from).

The reason why the perception of freeze frames is so interesting to me is that it fits in the framework of a construction of perception based on discrete samples. If the annihilation of motion processing mechanisms leads to the perception of freeze frames, it points to motion perception as the glue between discrete snapshots.

Motion perception deficits have been experienced by other patients, though never to the degree of completely losing the sense of motion (Blanke, Landis, Mermoud, Spinelli, & Safran, 2003; Vaina, Cowey, Eskew, LeMay, & Kemper, 2001). It could be that unilateral lesions are not disruptive enough to lead to the dramatic disturbance experience by L.M. Unilateral (W. Newsome & Pare, 1988) and bilateral (Marcar & Cowey, 1992) lesions were performed in monkeys, which were subsequently tested on coherence thresholds and found to be extremely impaired, but no data was collected to get a quantitative estimate of the subjective freeze frames the monkeys may have been experiencing. In order to get some data on the duration of freeze frames, I initially fancied the idea of using repetitive Transcranial Magnetic Stimulation (TMS) to bilaterally inactivate V5 (the region that was found to be most affected in L.M. (Shipp, Jong, Zihl, Frackowiak, & Zeki, 1994)). I spent some time reading up on TMS, and found a study that claimed to have induced akinetopsia in normal subjects (Beckers & Homberg, 1992). However, it presented results of the delivery of a single pulse of TMS to V5 during presentation of a motion stimulus, which does not really qualify to be called akinetopsia... There are in fact many studies using single pulses on V5 to disrupt the perception of motion (Hotson & Anand, 1999; Matthews, Luber, Qian, & Lisanby, 2001; Alvaro Pascual-Leone & Walsh, 2001; Sack, Kohler, Linden, Goebel, & Muckli, 2006; Silvanto, Cowey, Lavie, & Walsh, 2005), yielding interesting results but not recreating akinetopsia. Some studies used repetitive TMS (Cowey, Campana, Walsh, & Vaina, 2006; B. Thompson, Aaen-Stockdale, Koski, & Robert F Hess, 2009) but again the effects were pretty disappointing. Though I played with stimulation of V5 for a few weeks, I was not satisfied by the effects that I was seeing, and gave up on my idea. I believe that the lesion approach in monkeys is the best way to get a quantitative assessment of the duration of the purported freeze frames – one could for example imagine a task akin to pouring a virtual cup of tea and stopping at a defined level, with the level of the ‘water’ rising at a speed that fluctuates slightly so as to prevent exact predictions.

While I was exploring the literature, I came across a deficit that differed from L.M.’s, but seemed more common. Already in 1959, Teuber (Teuber, 1959) writes about patients who suffered gun wounds :

in impaired portions of defective visual fields, the perception of a continuous motion is frequently dissected into a series of multiple stationary images, quite analogous to the phenomena obtained for normal observers in Brown’s experiments when the target speed exceeded certain values. (Thus, one patient with a gunshot wound of the right occipitotemporal region complained that when a motorcycle passed him to his left, he saw instead ‘a string of motorcycles standing still’).

While the freeze frames seem to succeed to each other in the case of L.M., Teuber’s patients experience a series of stationary images instead of moving objects. This phenomenon, which I refer to as ‘the trailing phenomenon’, is in fact experienced by many other categories of patients, whether under the influence of drugs or linked to a neurological condition. I was so intrigued by it that I ended up writing a whole review about it ...

2. A more common, poorly studied disorder : Visual trails

a) *PAPER 2 : Dubois, J. and VanRullen, R. (2011) – PLoS Biology*

Unsolved Mystery

Visual Trails: Do the Doors of Perception Open Periodically?

Julien Dubois^{1,2*}, Rufin VanRullen^{1,2}

1 Université de Toulouse, Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, Toulouse, France, **2** Centre National de la Recherche Scientifique, UMR5549, Faculté de Médecine de Purpan, Toulouse, France

“Visual trailing” is a transient but dramatic disturbance of visual motion perception of unknown origin: the subject perceives a series of discrete stationary images trailing in the wake of otherwise normally moving objects. Although this phenomenon is most frequently encountered after ingestion of prescription and/or illicit drugs (most commonly with lysergic acid diethylamid, or LSD), it has also occasionally been reported following brain damage or neurological disorders. A quantitative account of visual trails is lacking; we argue that careful experimental investigation could potentially reveal how our brains update conscious visual perception in time.

What Do Visual Trails Look Like?

Ask any LSD user: they know the drug is taking effect when the “good trails” kick in. Trailing is a visual perceptual effect commonly experienced during LSD consumption and as a long-lasting side effect of the drug (hallucinogen persisting perception disorder) [1–4]. LSD users perceive a series of discrete positive afterimages in the wake of moving objects, a percept that has been likened to a multiple-exposure stroboscopic photograph, somewhat akin to Etienne-Jules Marey’s chronophotographs [5] from 1880, or to more recent digital art produced in a few clicks (Figure 1).

Trailing is a visual disturbance that has been observed under various conditions. While some authors use the terminology of LSD users [1–3,6,7], others describe the phenomenon in more clinically suited terms: examples are “akinetopsia” [8,9], “polyopia” [9], “palinopsia” [4,10–13], or “visual perseveration” [14]. In Table 1, we clarify the meaning of each of these terms (according to [5,8,15–21]). However, none of these terms precisely captures the specific elements of visual trails: they occur only for moving objects (unlike polyopia); moving objects are still perceived as moving (unlike akinetopsia); duplicate images are perceived only in the presence of real objects (unlike palinopsia); and, most importantly, visual trails are not mere streaks ([22]), as would appear on a long-exposure photograph of moving objects (visual perseveration). Rather, trailing consists of a discrete series of snapshots of the moving object along its past trajectory, as if successive frames of a movie had been superimposed. These phenomenological features of trailing are clearly portrayed in the case reports that we collected (Box 1). The discreteness of the trailing percept, in particular, is one of its most characteristic, and most mysterious, aspects.

Who Experiences Visual Trails?

While LSD induces trailing [1–3,23], consumption of other hallucinogens like psilocin and mescaline has not been linked with such visual disturbances in the clinical literature (to our knowledge). Either their interaction with the receptors that

mediate trailing is not as potent, or LSD acts at sites that other hallucinogens do not bind with (for a review of the pharmacology of hallucinogens, see [24] and for LSD more specifically, see [25,26]). Reports pertaining to the past use of illicit drugs are, however, poorly controlled (a good example is a report of illegal marijuana leading to trailing by Levi and Miller [27], in which they explicitly acknowledge that their cases may be confounded by the presence of impurities in illegal marijuana, as well as by past use of LSD).

A better understanding of the pharmacology of trailing may arise from the report of similar visual disturbances following the use of prescription drugs. The chemistry and dosage of prescription drugs is well controlled, which is not the case for illicit drugs. Nefazodone, for instance, is a recent antidepressant drug with multiple reports of episodes of visual trails as a side effect [6–8,28]. A closely related antidepressant drug, trazodone, can induce similar side effects [10]; so can mirtazapine [29], and, possibly, risperidone (the study is inconclusive due to the concurrent administration of trazodone [30]). These four antidepressant drugs all antagonize 5-HT_{2a} and 5-HT_{2c} receptors, and it has been proposed that this is at the origin of the trailing phenomenon [28]. However, the first recognized inducer of visual trails, LSD, has the opposite effect (agonist activity) at both these receptors [24–26]. Common to these antidepressant drugs and LSD is an increase in extracellular serotonin levels (indeed, all these drugs have been linked to serotonin syndrome [31], a form of poisoning due to excess serotonergic activity that causes a range of symptoms, including hallucinations, elevated body temperature, and fast heart beat). Although the evidence is suggestive of a link between serotonin and visual trailing, the widespread action of serotonin throughout the brain prevents us from speculating further for the time being.

Visual trails have also been experienced with other drugs with very different pharmacology. The antiepileptic drug topiramate, which, among other actions, blocks voltage-dependent sodium channels and increases GABA activity at some GABA-A receptor subtypes, has been involved with visual trails [11,12]. Increased inhibition, mediated by higher levels of serotonin or GABA

Citation: Dubois J, VanRullen R (2011) Visual Trails: Do the Doors of Perception Open Periodically? *PLoS Biol* 9(5): e1001056. doi:10.1371/journal.pbio.1001056

Published: May 10, 2011

Copyright: © 2011 Dubois, VanRullen. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: EURYI (<http://www.esf.org/activities/euryi/awards/2007.html>) ANR (<http://www.agence-nationale-recherche.fr/>) 06JCJC-0154. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

Abbreviations: LSD, lysergic acid diethylamid.

* E-mail: jcrdubois@gmail.com

Unsolved Mysteries discuss a topic of biological importance that is poorly understood and in need of research attention.



Figure 1. Multiple-exposure stroboscopic photograph. This illustrates (inasmuch as possible with a static image) the perception experienced during trailing.
doi:10.1371/journal.pbio.1001056.g001

potentiation, stands out as a likely common action of this drug and the ones described previously. Topiramate also inhibits cytochrome isoform CYP2D6 [32], which has been associated with serotonin syndrome [31].

Besides pharmacologically induced trailing, some clinical reports from neuropsychological populations deserve to be mentioned. A

Box 1. Selected Case Reports

Lysergic acid diethylamide (LSD) “[When asked] if they ever ‘saw any trails’, [the subject] would wave a hand across the visual field and say ‘Oh, you mean these?’ The subject could then describe a trail of images of the moving hand, much like the frames of a piece of motion picture film frozen in space long enough for the subject to see them” [3].

Nefazodone “A 47-year-old man reported seeing streams of multiple, frozen images trailing in the wake of moving objects. As soon as motion ceased, the images collapsed into each other. He compared his vision to a scene lit by a flashing strobe, except that stationary elements were perceived normally. In fact, if nothing was in motion and he held perfectly still, his vision was entirely normal. The moment anything moved, however, it left a stream of static copies in its path. For example, while out for an evening stroll, he saw a pack of identical dogs lined up behind his West Highland terrier. Driving was impossible because he was confused by multiple snapshots of cars, streets, and signs. Moving lights were followed by a long comet tail” [8].

Trazodone “He began to have numerous morning episodes of ‘strobe images’. They were most apt to occur in dim illumination. He described looking at the door of his bedroom and seeing multiple images of the door ‘march’ across his field of vision when he shifted his gaze to the bureau on the other side of the room. The false images were intense, and he had many such episodes every morning, each lasting 15 minutes or less” [10].

Mirtazapine “As she watched her husband walk past her, she saw multiple afterimages of him as if he were leaving a visual trail. These afterimages were less color intensive than the normal visual image, slightly blurred, and faded away after 30 seconds to 1 minute. The phenomenon repeated itself with most moving objects and was generally more pronounced with objects in Ms. A’s lateral visual fields” [29].

Topiramate “Ms. A reported seeing ‘picture in picture’ images, like she was in a ‘discotheque’, or in a place with stroboscopic lights. Those persistent ‘frozen pictures’ faded away after a few seconds. The phenomenon repeated itself with most moving objects” [11].

crisp account of migraine patients experiencing trailing is given by Sacks [5], where he estimates the frequency of snapshots at “six to twelve frames per second” (see also [33]). We found one other report of visual trails (combined with akinetopsia and palinopsia) in migraineurs [34] (see also medical doctor Klaus Podoll’s Web site [35]). On the whole, trailing is a very rare disturbance for migraine sufferers, the migraine being a very heterogeneous condition that affects roughly 20% of the population at one time or another; the possibility exists that the disturbance occurs only for a very specific and rare form of migraine, or is a side effect of drugs that patients take to relieve their symptoms [12]. As a side note of interest, controlled studies have found migraine patients to have impairments for global motion perception [36] (see also [37]).

A recent case study [9] described visual trails associated with another neurological condition: the posterior cortical atrophy variant of Alzheimer’s disease [38]. The symptoms are referred to as akinetopsia/polyopia by the authors, but their description corresponds to a direction-specific version of trailing (visual trails are perceived when objects move from right to left, but not when they move from left to right). Direction-specificity restricts the mechanistic models that can be put forward to explain visual trails; however, the etiology is very different from previously discussed cases, and in the absence of other reports of directional trailing, it is premature to draw conclusions from this report.

Finally, a very rare disturbance of motion perception, for which Zeki coined the word akinetopsia [15], can follow bilateral lesions in the occipito-temporo-parietal cortices [39]. Patient L. M. permanently lost motion perception (except for slowly moving objects [40]). Instead of seeing a moving object, she saw the object in a series of stills, as in a movie run too slowly. This is a rather extreme case that differs qualitatively from trailing (see Table 1), but it may ultimately rely on similar mechanisms.

What May Cause the Perception of Visual Trails?

Visual trails, because of their discrete and repetitive nature, may represent the perceptual manifestation of an underlying periodic process. This periodicity could arise outside of the neural system (e.g., eye movements, motor tremor), or may be the result of faulty motion computation mechanisms (e.g., motion smear suppression), or, finally, may point to a more general, quasi-periodic sampling process that affects, among other things, the motion perception system. These three (non-exhaustive and non-exclusive) possibilities are developed below.

Do Visual Trails Arise from Abnormal Eye Movements?

Trailing may not originate in the neural pathways at all. For example, if visual trails are only perceived when subjects track a moving object with their eyes, the discrete afterimages may be caused by a defective, jerky smooth-pursuit mechanism. Saccadic suppression would operate while the eyes execute repeated corrective saccades to keep track of the moving object, shutting on and off the

Table 1. Glossary.

TRAILING		
An object in motion is followed by a discrete series of stationary images that slowly fade out along its trajectory		
	Our Definition	How It Differs from Trailing
AKINETOPSIA	An object in motion is not perceived as moving continuously; rather, it is experienced as a series of stills	The perception of motion is abolished and a single frame can be seen at any one time
POLYOPIA	Multiple images of an object (from two to 100 in extreme cases) are perceived simultaneously	Polyopia is not specific to objects in motion
VISUAL PERSEVERATION	A positive afterimage persists after removal of an object	Visual perseveration for a moving object would leave a blur in its wake, rather than discrete images
PALINOPSIA	Similar to visual perseveration; some authors describe a lag between the removal of the object and the onset of the afterimage	Same as visual perseveration, potentially with a gap between the moving object and the perseverant tail

doi:10.1371/journal.pbio.1001056.t001

visual input intermittently. Note that people having experienced trailing often recount that trails follow just one or a few moving objects at a time, rather than affecting the entire visual scene, as would be the case after eye movements (or uncontrolled tremors). This makes such a motor explanation unlikely; nonetheless, a well-controlled eye tracking experiment on individuals experiencing visual trails would be essential to evaluate this hypothesis.

Do Visual Trails Reflect a Failure of Specific Motion Perception Mechanisms?

The periodicity may arise as a disturbance of certain motion processes that are normally continuous. For example, motion streak suppression [41] is an inhibitory mechanism allowing the brain to regulate the smear that a moving object leaves in its wake, owing to visible persistence (an image normally takes about 100 ms to fade from perception, long enough to blur the trajectory of a moving object as in a long-exposure photograph). Under the effects of LSD or related drugs, the streak suppression process might fall out of its normal operating range, resulting in stronger than normal local inhibition followed by excitatory rebound. This hypothesis would predict that trailing is accompanied by abnormal oscillatory activity confined within the motion perception system, a prediction that could be verified in human or non-human primates with appropriate recordings.

Could Visual Trails Be the Manifestation of Periodicities Inherent to all Perceptual Processes?

Trailing may reflect a more widespread rhythmic activity affecting (possibly among other modules) the motion perception system. Two alternatives must be distinguished here, which could be teased apart by comparing oscillatory brain activity in the normal and drug-altered states. Either this rhythmic activity is directly produced by the drug (or at least, increased beyond perceptual threshold)—a rise of inhibition could generate prominent oscillations, turning the normally continuous processing of visual information into a series of discrete snapshots—or this periodicity is always present in the normal brain, but inaccessible to conscious perception. In the latter case, the drug may increase visible persistence, or disrupt motion streak suppression, two processes that would normally serve to hide the discrete visual trails.

Related Phenomena Pointing to Periodicities in Visual Perception

A persistent thread in the perceptual literature is that temporal binning is, indeed, common to all of perception (reviewed by [42]). One of the most striking phenomenological manifestations of the

discrete nature of perception is the so-called “continuous wagon wheel illusion”: in plain sunlight, a continuously rotating, spoked wheel can be perceived to rotate in a direction opposite to its true motion. While movie watchers are accustomed to this percept (which follows from an undersampling of the continuous motion of the wheel by the discrete frame rate of the camera and the movie projection system), perceiving reversals in conditions of continuous illumination is more challenging. Quasi-periodic sampling or binning processes within the visual system have been proposed as an explanation; this interpretation has been a subject of controversy, and many experiments have been conducted to test it. In light of these recent experiments, VanRullen and colleagues hypothesized that attentional processes may function in a quasi-periodic manner [43,44]: when attention is deployed, it samples information at discrete moments in time. The rate of sampling may be dependent on the task at hand [43,45,46]. One component of this theory purports that the attentive motion system takes discrete samples of the object in motion to compute its trajectory. Past snapshots are usually concealed from conscious perception—but some substances may interfere with their proper suppression, giving rise to the trailing phenomenon.

Towards a Solution: Measuring Visual Trails

How many “ghost images” [23] trail in the wake of the moving objects? How far apart are the “discrete and discontinuous images” [2]? How long do they persist for? These questions have, so far, not been answered. In fact, visual trails have been considered an annoyance at best, and quantifying them has not been the main concern in the various clinical settings in which they were observed. Careful quantitative investigation will be needed to shed light on the mechanisms that cause visual trails to appear as a discrete series of snapshots of the moving object.

As a preliminary step towards quantitative answers, we conducted an online survey in which we asked self-declared past LSD users to decide which of ten short movies with simulated visual trails best matched their recollected experience. We varied the interval between simultaneously perceived snapshots from 25 ms to 250 ms, keeping the number of concurrently visible repetitions to four. For over more than 210 participants, the responses were not randomly distributed across the ten choices, but followed a highly consistent pattern (chi-squared test, $p < 0.0008$) with a preference for faster snapshot periodicities; participants selected a time interval between images of 67 ms on average, corresponding to an underlying periodicity in the 15–20 Hz (beta) range. Although this is the first study that tries to quantitatively assess the frame rate of LSD-induced visual trails, it

suffers from numerous pitfalls. First, we have little control over the individuals who take the survey, as they do so anonymously from their home computer. However, this protocol follows a new trend in psychological research of using the Web for large-scale studies (see the <http://www.testmybrain.org/> Web site recently created by Harvard scientists). Second, the responses rely on the memory of the percept, which may be faulty. Last, the chemical composition and dose taken by each individual is not controlled, which is problematic considering that more than 200 types of LSD tablets have been encountered since 1969 and more than 350 blotter paper designs have been in circulation since 1975 [47] (also, some blotters sold as LSD are in fact mimics). In fact, some of the most experienced users indicated that trailing is dose dependent. For all these reasons, one must be cautious in interpreting these initial results. The movie used for the survey

as well as a results summary can be found online at <http://www.cerco.ups-tlse.fr/~rufin/lldssurvey>.

Collecting further quantitative data with individuals who experience visual trailing will be necessary to tease apart the alternative accounts—possibly in combination with computational modeling. Solving the mystery of the origins of the trailing effect might reveal something deep about the mechanisms underlying perception, challenging the way we think we perceive the world.

Acknowledgments

We wish to thank Christof Koch for useful discussions and comments on the manuscript. The Multidisciplinary Association for Psychedelic Studies advertised the survey about LSD trails on their Web site (<http://www.maps.org/participate/>) and without their help we would not have collected as many responses as we did.

References

- Asher H (1971) “Trailing” phenomenon—a long-lasting LSD side effect. *Am J Psychiatry* 127: 1233–1234.
- Anderson WH, O'Malley JE (1972) Trifluoperazine for the “trailing” phenomenon. *JAMA* 220: 1244–1245.
- Abraham HD (1983) Visual phenomenology of the LSD flashback. *Arch Gen Psychiatry* 40: 884–889.
- Kawasaki A, Purvin VA (1996) Persistent palinopsia following ingestion of lysergic acid diethylamide (LSD). *Arch Ophthalmol* 114: 47–50.
- Sacks O (2004) *In the river of consciousness*. The New York Review of Books 51.
- Kraus R (1996) Visual “trails” with nefazodone treatment. *Am J Psychiatry* 153: 1365.
- Schwartz K (1997) Nefazodone and visual side effects. *Am J Psychiatry* 154: 1038.
- Horton JC, Trobe JD (1999) Akinetopsia from nefazodone toxicity. *Am J Ophthalmol* 128: 530–531.
- Tsai PH, Mendez MF (2009) Akinetopsia in the posterior cortical variant of Alzheimer disease. *Neurology* 73: 731.
- Hughes MS, Lessell S (1990) Trazodone-induced palinopsia. *Arch Ophthalmol* 108: 399.
- Fontenelle LF (2008) Topiramate-induced palinopsia. *J Neuropsychiatry Clin Neurosci* 20: 249.
- Evans RW (2006) Reversible palinopsia and the Alice in Wonderland syndrome associated with topiramate use in migraineurs. *Headache* 46: 815–818.
- Abert B, Ilsen PF (2010) Palinopsia. *Optometry* 81: 394–404.
- Hoffer A, Osmond H (1967) *The hallucinogens*. New York: Academic Press.
- Zeki S (1991) Cerebral akinetopsia (visual motion blindness). A review. *Brain* 114(Pt 2): 811–824.
- Bender M (1945) Polyopia and monocular diplopia of cerebral origin. *Arch Neurol Psychiatry* 54: 323.
- Norton JW, Corbett JJ (2000) Visual perceptual abnormalities: hallucinations and illusions. *Semin Neurol* 20: 111–121.
- Lopez J, Adornato B, Hoyt WF (1993) ‘Entomopia’: a remarkable case of cerebral polyopia. *Neurology* 43: 2145.
- Bynke H (1984) Visual perseveration following temporal-lobe surgery - case-report. *Neuroophthalmology* 4: 47–53.
- Meadows JC, Munro SS (1977) Palinopsia. *J Neurol Neurosurg Psychiatry* 40: 5–8.
- Smith PE, Shah P, Sharpe J, Todd A, Goringe AP (2003) Palinopsia. *Lancet* 361: 1098.
- Ermentrout B (1999) The excited cortex - LSD trails, phosphenes, and other visual confections. Eighth Annual Computational Neuroscience Meeting; Pittsburgh, Pennsylvania; July 1999. Abstracts. 201 p.
- Kawasaki A, Purvin V (1996) Persistent palinopsia following ingestion of lysergic acid diethylamide (LSD). *Arch Ophthalmol* 114: 47–50.
- Nichols DE (2004) Hallucinogens. *Pharmacol Ther* 101: 131–181.
- Passie T, Halpern JH, Stichtenoth DO, Emrich HM, Hintzen A (2008) The pharmacology of lysergic acid diethylamide: a review. *CNS Neurosci Ther* 14: 295–314.
- Hintzen A, Passie T (2010) *The pharmacology of LSD: a critical review*. Oxford: OUP/Beckley Foundation Press. 240 p.
- Levi L, Miller NR (1990) Visual illusions associated with previous drug abuse. *J Clin Neuroophthalmol* 10: 103–110.
- Faber RA, Benzick JM (2000) Nefazodone-induced palinopsia. *J Clin Psychopharmacol* 20: 275–276.
- Ihde-Scholl T, Jefferson JW (2001) Mirtazapine-associated palinopsia. *J Clin Psychiatry* 62: 373.
- Lauterbach EC, Abdelhamid A, Annandale JB (2000) Posthallucinogen-like visual illusions (palinopsia) with risperidone in a patient without previous hallucinogen exposure: possible relation to serotonin 5HT_{2a} receptor blockade. *Pharmacopsychiatry* 33: 38–41.
- Boyer EW, Shamon M (2005) The serotonin syndrome. *N Engl J Med* 352: 1112–1120.
- Benedetti MS (2000) Enzyme induction and inhibition by new antiepileptic drugs: a review of human studies. *Fundam Clin Pharmacol* 14: 301–319.
- Sacks OW (1999) *Migraine*. New York: Vintage Books. 368 p.
- Klee A, Willanger R (1966) Disturbances of visual perception in migraine review of the literature and a report of eight cases. *Acta Neurologica Scandinavica* 42: 400–414.
- Podoll K (2007) Migraine Aura Foundation. Available: <http://www.migraine-aura.org>. Accessed 4 April 2011.
- McKendrick AM, Badcock DR, Badcock JC, Gurgone M (2006) Motion perception in migraineurs: abnormalities are not related to attention. *Cephalalgia* 26: 1131–1136.
- Graziera C, DaSilva AFM, Snyder J, Tuch DS, Hadjikhani N (2006) Anatomical alterations of the visual motion processing network in migraine with and without aura. *PLoS Med* 3: e402. doi:10.1371/journal.pmed.0030402.
- Mendez MF, Ghajarian M, Perryman KM (2002) Posterior cortical atrophy: clinical characteristics and differences compared to Alzheimer's disease. *Dement Geriatr Cogn Disord* 14: 33–40.
- Zihl J, Von Cramon D, Mai N (1983) Selective disturbance of movement vision after bilateral brain damage. *Brain* 106: 313.
- Hess R, Baker CL, Zihl J (1989) The “motion-blind” patient: low-level spatial and temporal filters. *J Neurosci* 9: 1628.
- Burr D (1980) Motion smear. *Nature* 284: 164–165.
- VanRullen R, Koch C (2003) Is perception discrete or continuous? *Trends Cogn Sci* 7: 207–213.
- VanRullen R, Reddy L, Koch C (2005) Attention-driven discrete sampling of motion perception. *Proc Natl Acad Sci U S A* 102: 5291.
- VanRullen R, Reddy L, Koch C (2006) The continuous wagon wheel illusion is associated with changes in electroencephalogram power at approximately 13 Hz. *J Neurosci* 26: 502–507.
- VanRullen R, Carlson T, Cavanagh P (2007) The blinking spotlight of attention. *Proc Natl Acad Sci U S A* 104: 19204.
- Busch NA, VanRullen R (2010) Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proc Natl Acad Sci U S A* 107: 16048–16053.
- Laing RR, ed. (2003) *Hallucinogens: a forensic drug handbook*. San Diego: Academic Press. 304 p.

b) An online survey among past LSD users

As briefly reported at the end of the paper, we conducted an online survey among self-declared past LSD users to get a first quantitative estimate of the time interval between successive frames in the trailing phenomenon. Anonymous surveyees were asked to watch a 2min movie, composed of ten different versions of a 15s sequence, and they subsequently voted for the sequence best matching their experience. We also collected basic demographic information, as well as free comments (a dangerous thing when you deal with people taking hallucinogenic drugs... see Table 2).

(1) Simulation of visual trails

For the purpose of this study, we bought a digital camera that could take high speed video : the casio exilim EX-FH20. This little camera can capture movies at up to 1000 frames per second (fps) – though at the unimpressive resolution of 224x56 pixels. We used the 210 fps setting, with a resolution of 480x360 pixels. This means that it takes an image every 4.76ms. We took a few different movies, trying to use situations that LSD users may have encountered during their trips: the wave of a hand, a ball describing a nice parabolic trajectory as it was passed back and forth between two volleyball players, and a ball juggled by a soccer player with his foot (with hindsight, this last scene is not adapted to a US-based population!).

The first step in automatically generating visual trails was to detect motion in the movies. We developed our own straightforward approach and matlab code to perform this analysis, which worked fine for the few movies we had to process. First, define the background for each color channel (red, green, and blue): it works fairly well to take the median value for each pixel over the course of the movie as the background value for that pixel. Next, compute the absolute value of the pixel value difference between each frame of the movie and the background. Each pixel is evaluated independently as to whether it exceeds an arbitrary threshold (30) in any of the color channels, or exceeds a arbitrary global threshold when the absolute differences in each channel are added. One ends up with a speckled binary image, containing 1's where 'motion' is detected. The motion image is smoothed – through convolution with a square of 5 pixels width, and arbitrary thresholding (17), we remove 'false positives' that don't have enough neighbors. Finally, to slightly smooth the edges of detected motion areas, we convolve the image with a gaussian of 5 pixels width and 5 pixels standard deviation... the end result could be taken as a sort of rough motion probability map.

Visual trails have usually been described as slowly decaying positive afterimages along the past trajectory of a moving object. Hence, once moving objects have been detected, it is only a matter of 'pasting' some snapshots from past frames in each frame of the movie.

We had to make some choices, which were guided by the literature and input from a subject who had first-hand experience of LSD induced trailing : the motion of the leading moving object was chosen to be perceived normally, and positive afterimages were chosen to be static – rather than moving with the leading object. The weight of a given snapshot was chosen to decay exponentially – the number of perceived positive afterimages at any point in time critically depends on the time constant of this decay, and we chose it such that four afterimages could be seen in the trails. Finally, we varied the value of the last parameter – the one that we were interested in estimating through the survey: the distance (in ms) between two positive afterimages. Once again, parameterizing this distance in time units is an assumption corresponding to our belief that trailing relies on oscillatory activity, which may not be true.

Figure 36 is an illustration of the weights given to frames of the original movie in each frame of the new movie. The white line marks a given frame of the new movie. Snapshots of the three scenes that we used are shown in Figure 37. You will notice some artefacts, which we judged acceptable.




Figure 36 An illustration of the weights given to past frames in our simulations of visual trails.



Figure 37 Examples frames from the simulated movies.

(2) Experimental procedure

Our survey was advertised through the website of the Multidisciplinary Association for Psychedelic Studies ([MAPS](#)). It linked to a [webpage](#) with the movie. Surveyees could watch the movie and then fill out a survey questionnaire (Figure 38) powered by [KwikSurveys](#).

We varied a single parameter: the time between successive snapshots. We fixed the number of simultaneously perceived positive afterimages to 4, following suggestions from a LSD user. The times that we used were: 23.8ms (5 frames @210fps), 52.4ms, 76.2ms, 100ms, 123.8ms, 152.4ms, 176.2ms, 200ms, 223.8ms, and 252.4ms. The 10 different sequences were presented in a fixed random order for the first 236 subjects (223.8|176.2|152.4|76.2|252.4|200|123.8|52.4|23.8|100), then the order of presentation was reversed for the following surveyees.

LSD INDUCED VISUAL TRAILS ..

Please answer the following questions to the best of your ability.

* Where are you from?

* What year were you born?

* Gender
 Male Female
[Reset](#)

* Highest level of education completed :

* How many times have you taken LSD?
 0 1 2 to 5 6 to 10 11 to 20 more than 21
[Reset](#)

* How many times have you experienced LSD induced visual trails? (during the trip)
 0 1 2 to 5 6 to 10 11 to 20 more than 21
[Reset](#)

* In the movie you watched, which sequence was the BEST match to your experience of visual trails?
 1 2 3 4 5 6 7 8 9 10
[Reset](#)

* Rate from 0 star (not at all like your experience) to 10 stars (exactly like your experience) how well the best sequence matched your experience.

[Reset](#)
 (OPTIONAL) How can it be improved?

(OPTIONAL)
 Do you have any further comments regarding LSD induced visual trails?

[Finish Survey](#)

Figure 38 The first version of the survey form which we posted on the KwikSurveys website.

(3) Results

Figure 39 shows the results from the merged datasets, as there was no noticeable influence of the order of presentation on surveyees' preferences (or demography).

It can be seen that we've been receiving responses at an approximately steady rate of one per day (Figure 39a). Most surveyees are from the US (Figure 39b), and, quite surprisingly, participants are mostly under 30 (Figure 39c) – LSD was most prevalent at the end of the 60s, but our data may reflect a sampling bias due to the medium we used to recruit participants (the World Wide Web). There were more than twice as many males as females among surveyees (Figure 39d), and most surveyees were fairly educated (Figure 39e; given their age, some of them who answered "some high school" are still likely to go on to college...). It appears in Figure 39f that the perception of visual trails happens on a great majority of LSD trips (strong diagonal on the plot). To the left of the diagonal are the users who have experienced visual trails on only a fraction of their trips. To the right of the diagonal are the users who have experienced visual trails more often than they have taken LSD: this corresponds to "flashbacks", or Hallucinogen Persistent Perceptual Disorders. It should be noted that our results may not accurately reflect the proportion of LSD users experiencing visual trails; indeed, one of the instructions

to the participants was to complete the survey if they had experienced visual trails... so we are missing the users who have never experienced visual trails.

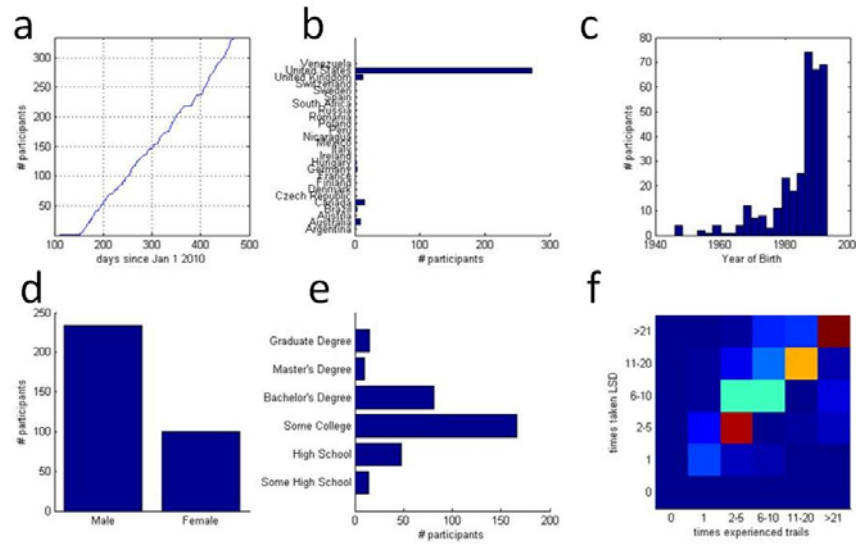


Figure 39 a cumulated number of participants since the start of the survey. b country of origin of participants. c year of birth of participants. d gender of participants. e education of participants. f occurrence of visual trails as a function of LSD use.

Pooling all responses and looking at their distribution, we get the distribution depicted in Figure 40.

Responses are not randomly distributed, and there is a consistent pattern towards choosing faster “snapshot” periodicities. The mean and median of the distribution were computed, and are displayed on the figure. The first conclusion one may draw from this figure is that, had a sequence with even tighter snapshots been presented, it would have received many votes. We will address this concern in the discussion. It may be a good idea to look at the results for a subset of the population – for example, the most experienced LSD users. Why? because they are the ones most likely to know what a “real” trail looks like. We selected the participants who reportedly took LSD more than 21 times and experienced visual trails more than 21 times as well. This subset represents about 20% of surveyees. In Figure 41, you can see that the preference for these subjects tends to be for slightly longer durations between individual images (as reflected by the median of the distribution).

(4) Discussion

This study is far from being perfect – when we designed it, we were under the impression that everyone who takes LSD experiences strobe-like tracers, a view that we had gathered from case reports in the literature and talking to a few people who had taken LSD in the past. Hence, we set out to measure the time interval between successive snapshots of the moving objects. Instead of directly asking questions such as “how many images do you see trailing behind a moving object?” “how distant are the discrete positive afterimages from each other?” “do the trailing images move, following the real object, or are they static?”, we decided it would be more objective to fix a few parameters and see if there would be a consensus among past LSD users. There is an obvious confound : we did not state explicitly that the only parameter that was being varied was the interval between successive snapshots – in retrospect, this may have been a mistake. As we fixed

the number of afterimages to 4 in all sequences, the length of the trails was directly proportional to the interval between snapshots. Hence, surveyees may have based their decision on the length of the trails rather than the spacing between successive afterimages. This is indeed apparent from many of the comments that we collected.




Figure 40 Overall distribution of surveyees' choices. The goodness of the match between the video sequence and actual trails experienced by the subject is color-coded, dark blue meaning a very poor match and dark red a perfect match.




Figure 41 Distribution of surveyees' choices for a subset of the population, the users who reportedly took LSD more than 21 times.

Perhaps the most interesting results in this study come from reading the surveyees' comments. Going through the more than 350 completed surveys, there are a few points that came back quite often in the comments of the participants:

- Trailing is very dependent on the dose and the quality of the LSD taken : summarizing what we read in many of the comments, low doses lead to smaller 'trails', which are better described as motion blur or long-exposure photographs rather than discrete positive afterimages or multiple-exposure photographs. Only the highest doses lead to the trailing effect as described in case reports in the literature (e.g. (W. H. Anderson & O'Malley, 1972)), and as simulated in our video sequences. This explains the large number of participants who preferred the shortest intervals between frames; it is expected that the majority of users will have taken low doses of LSD – a

shortcoming of the survey was to not ask specifically for this information, although it is hard to compare dosage given the variability in the chemical composition of black market blotters. In fact, about one third of the surveyees who chose the 23.8ms or 47.6ms sequences commented that the videos were too “choppy”, that they should be more “blurry”, “smeared-out”, “fluid”. Some representative quotes : “like speed lines in a cartoon”, “more like the object was stretched”, “the trail is never separated into individual images”. It is thus clear many of these participants did not experience trailing in the way we defined it in our review. In a way, their experience is better captured by the word “tracers”, which is also part of LSD users’ vernacular. At this point we need to emphasize that trailing as we defined it was indeed experienced by some of the participants surveyed. It is difficult to provide an accurate estimate of the proportion of surveyees who indeed perceived trailing – our cues rely on the snapshot interval chosen, the quality of the match and the comments they provided. Among these cues, the judgement of the quality of the match may in fact rely on other factors (e.g. color and other visual effects), and the snapshot interval, being our dependent variable, is not an honest discriminator. Remain the comments. Luckily, comments were provided by more than 75% of the surveyees. From these comments, we can undoubtedly say about 5-10% of the participants claim that they perceived discrete positive afterimages under LSD; this is a lower bound. Some representative quotes: “For a very strong dose, trails could be up to 10 or 15 afterimages, with many of them nearly as strong as the first one, and not fading quickly”, “almost like a camera shutter going off several times really quickly to capture some snapshots in a series. Like I said a sort of strobe effect”, “if I were to move my hand in front of my face and focus on an empty space between I’d see more hands”, “it is not usually a smooth trail, but choppy as if I am seeing a series of snapshots”, “the ‘snapshots’ of the afterimages occurred according to a consistent clock; a slow hand sweep might create 4 to 8 afterimages with short spacing, a longer hand motion would create the same 4-8 images, more widely spaced by the higher speed, but still perfectly equidistant one from the other”, “repeated objects are spaced a little further away especially when the object is moving faster”, “the later trails blur together into a smear of color, whereas the trail near the object is more of the stutter-repeat-frame effect you show in your examples”, “sequence 3 (*author comment: 150ms*) and 9 (*author comment: 25ms*) were most like the trails I had seen but sequence 3 only on 2 very intense and memorable occasions”.

- Some participants pointed out that tracers are more readily experienced with light moving objects against a dark background, referring to glow-sticks and cell phone screens in the context of a rave... once again, we believe that the perceptual experience that these participants had was not what we have defined as the trailing phenomenon, but increased visual persistence.
- Another major criticism that kept coming back was about the opacity of the trails. Many participants commented that the trails are more “transparent”, “see-through”. We also got a lot of comments on the colors and other “fractalization” and “breathing” that my videos were lacking – which made it difficult for some subjects to identify the video sequences with their hallucinogenic experience.

It clearly appears that our sample of participants should be divided in two categories: the ones who have experienced true trailing, of the kind described in case reports in the literature and evidenced by the sample of comments that we report; and the ones who have simply seen tracers, or motion blur, likely through some increase of visual persistence. However, in the data that we collected, there is no clear bi-modality in the distribution over all participants, or variable that would allow us to clearly distinguish the two populations – comments that clearly identify trailing are difficult to select but in a few very obvious cases. Perhaps the closest we can get to separating the populations in our data is by considering only the most experienced users, who have taken LSD more than 21 times, and have likely been exposed to higher doses than less

experienced users; indeed, the second graph in the results section (Figure 41) shows a less pronounced dominance of the very short intervals, to the benefit of slightly longer periods (75-125ms).

To further drive the point that some LSD users experience trailing as a series of discrete positive afterimages of moving objects, Figure 42 shows two pictures that were mentioned as close depictions of trailing by two separate participants.




Figure 42 : a) Queen – Bohemian Rhapsody video (from www.youtube.com). This visual effect actually corresponds to the description given by more than one participant – in the immediate vicinity of the moving object, discrete snapshots can be individualized, but they merge together in a blur as one moves further along the trail. b) Thomas Eakins – *Motion Study: George Reynolds nude, pole-vaulting to left, 1885* (from www.pafa.org). This is very similar to Etienne-Jules Marey’s chronophotographs already mentioned by O. Sacks (Sacks 2004) as a good depiction of the trailing phenomenon.

I should mention an extremely interesting comment from one participant, which pertains to my initial attempts to find auditory evidence for aliasing in section II.B (page 40).

I do not believe that the process that causes a person to experience trailing is limited to sight alone. I know that I have experienced aural sensations in which sounds will sound as if they are echoing, which for the sense of hearing is akin to seeing an object ‘trail’ through your field of vision as time progresses

What have we learnt from this study? First, that not all LSD users experience trailing as we defined it in the paper – a very high dose of LSD may be necessary to perceive such extreme effects. Second, with all the precautions that we pointed out here and in the paper, for people who experience trailing (as opposed to tracers), snapshots may be taken at intervals of around 75ms on average. This result is a preliminary quantitative estimate of snapshot periodicity in the trailing phenomenon, and much better controlled empirical research is needed to address the question – and distinguish the various speculative accounts of the trailing phenomenon that we have formulated.

As I was poking around and spending time on various websites to research LSD and its effects, I came across a device coined the “LSD flight simulator” (Figure 43). Using this device recreates some of the visual hallucinations that LSD intake triggers. I of course purchased it to try it out. All you have to do is blow through a small hole while looking at the sun, eyes closed; your breath will make a blade rotate, which periodically hides the sun. Basically, a powerful flickering light is shone on your eyelids. When the flickering rate is right, you start experiencing some weird geometrical, colorful patterns – visual hallucinations. It seems that these hallucinations are thus induced by photic stimulation, probably through entrainment of

an oscillation in early visual areas. If some effects of LSD can be reproduced by this mechanism, it may mean that visual trailing is in fact also a by-product of abnormal oscillatory activity rather than a case in which an intrinsic sampling mechanism would break through to perception. But, this is very speculative – the question remains open and testable.



Figure 43 The LSD flight simulator. Downloaded from http://www.collegian.com/index.php/article/2006/11/daydreams_can_come_true

I have proposed some experiments that should be performed on people experiencing visual trailing. As I did not readily have access to such patients, I could not perform these experiments myself...

Insightful?	<p><i>isnt tracers really simply a blur effect in your frontal cortex from the lsd acting as a neurochemical filter between your synapses and rna receptors?</i></p> <p><i>The visual trails may just be energy</i></p>
Delirious?	<p><i>People who you are more emotionally connected with than others will emit different colours around the edges of the trails than those you are less familiar with [...] The colours are different for every person.</i></p> <p><i>if you were to do an around the world on a yoyo, the trails would make it look like the yoyo isn't spinning, but in multiple places at the same time. As one yoyo disappears, another fades in.</i></p> <p><i>include the colors and more important is the fibonacci-fractal pattern that encompasses existence</i></p> <p><i>The effect of LSD on the mind is very abstract, as I'm sure you know</i></p>
Sobering	<p><i>The doors of perception can not be explained, they must be walked through</i></p> <p><i>There is nothing beneficial to learn from this study</i></p>
Suggestive	<p><i>drop some f... lsd and see for yourselves, I promise you won't regret it</i></p>

Table 2 Best of : some selected quotes from the comments section of the survey.

3. Summary and Discussion

In this chapter, I explored the realm of pathological disorders, and presented cases that may reveal periodic mechanisms at work in healthy subjects. I advocated a quantitative approach to studying temporal aspects of perception in patients. In the case of visual trailing (or akinetopsia), simple experiments (implementing for instance an adjustment method based on our

simulations of visual trails, in which subjects could vary different parameters) could yield estimates of the duration between snapshots of moving objects, and brain/eye movement recordings could help constrain the speculations that flourish when we imaginative scientists have too little data to rely upon...

E. PAPER 3 : VanRullen and Dubois (2011) – Frontiers in Psychology

We wrote a review paper (currently under revision) which can serve as a summary for this section on the psychophysical and behavioral evidence for rhythms in perception.



The psychophysics of brain rhythms

Rufin VanRullen^{1,2*} and Julien Dubois^{1,2,3}

¹ Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, Université de Toulouse, Toulouse, France

² Faculté de Médecine de Purpan, CNRS, UMR 5549, Toulouse, France

³ Division of Biology, California Institute of Technology, Pasadena, CA, USA

Edited by:

Gregor Thut, University of Glasgow, UK

Reviewed by:

Hinze Hogendoorn, Utrecht University, Netherlands

Ayelet Nina Landau, Ernst Strüngmann Institute in Cooperation with Max Planck Society, Germany

*Correspondence:

Rufin VanRullen, Centre de Recherche Cerveau et Cognition, Pavillon Baudot, Hôpital Purpan, Place du Dr. Baylac, 31052 Toulouse, France.
e-mail: rufin.vanrullen@cerco.ups-tlse.fr

It is becoming increasingly apparent that brain oscillations in various frequency bands play important roles in perceptual and attentional processes. Understandably, most of the associated experimental evidence comes from human or animal electrophysiological studies, allowing direct access to the oscillatory activities. However, such periodicities in perception and attention should, in theory, also be observable using the proper psychophysical tools. Here, we review a number of psychophysical techniques that have been used by us and other authors, in successful and sometimes unsuccessful attempts, to reveal the rhythmic nature of perceptual and attentional processes. We argue that the two existing and largely distinct debates about discrete vs. continuous perception and parallel vs. sequential attention should in fact be regarded as two facets of the same question: how do brain rhythms shape the psychological operations of perception and attention?

Keywords: oscillation, perception, attention, psychophysics, discrete perception, sequential attention

INTRODUCTION

Neurons convey information by means of electrical signals. Due to intrinsic properties of neuronal networks (e.g., conduction delays, balance between excitation and inhibition, membrane time constants), these electrical pulses give rise to large-scale periodic fluctuations of the background electric potential, which constitute the brain “rhythms” and oscillations (Buzsáki, 2006). Some oscillations – like the “alpha” rhythm at 8–13 Hz can be seen with the naked eye on an electro-encephalographic trace (Berger, 1929), while others require sophisticated analysis methods or recordings with a higher signal-to-noise ratio (using intra-cerebral probes in animals and, more rarely, in humans). There are many theories implicating brain oscillations in the performance of particular cognitive functions such as perception (Eckhorn et al., 1988; Gray et al., 1989; Engel et al., 1991; Singer and Gray, 1995; von der Malsburg, 1995), attention (Niebur et al., 1993; Fell et al., 2003; Womelsdorf and Fries, 2007), consciousness (Koch and Braun, 1996; Gold, 1999; Engel and Singer, 2001), and memory (Lisman and Idiart, 1995; Klimesch, 1999; Kahana et al., 2001). There are also flurries of experimental studies supporting (and sometimes, invalidating) these theories based on electrophysiological measurements of brain activity (Revonsuo et al., 1997; Tallon-Baudry et al., 1997, 1998; Fries et al., 2001; Jensen et al., 2002; Gail et al., 2004; Ray and Maunsell, 2010). It is somewhat less ordinary, on the other hand, to investigate the consequences of brain oscillations using psychophysical techniques. Yet one major prediction of the above-mentioned theories is directly amenable to psychophysical experimentation: indeed, an oscillatory implementation at the neuronal level should imply that the relevant cognitive function fluctuates periodically, and such fluctuations should be measurable with standard (or slightly more sophisticated) experimental psychology techniques.

The purpose of this article is to review some of the psychophysical techniques that have been applied recently to the study of brain oscillations. In so doing, we will touch upon two classical debates

in perception research. First, scientists have long theorized that oscillations could divide the continuous sequence of inputs feeding into our perceptual systems into a series of discrete cycles or “snapshots” (Pitts and McCulloch, 1947; Stroud, 1956; Harter, 1967; Allport, 1968; Sanford, 1971), but this idea is far from mainstream nowadays; we refer to this debate as “discrete vs. continuous perception” (VanRullen and Koch, 2003). Second, there is another long-standing debate known as “parallel vs. sequential attention”: does attention concentrate its resources simultaneously or sequentially when there are multiple targets to focus on? Though this discussion is generally disconnected from the topic of brain oscillations, we will argue that it is in fact germane to the previous question. The sequential attention idea – which has traditionally been the dominant one – originated with the assumption that high-level vision cannot process more than one object at a time, and must therefore shift periodically between the stimuli (Eriksen and Spencer, 1969; Treisman, 1969; Kahneman, 1973), just like our gaze must shift around because our fovea cannot fixate on multiple objects simultaneously. Interestingly, sequential attention theories require a (possibly irregular) cyclic process for disengaging attention at the current target location and engaging it anew. It is easy to see – although little noticed in the literature – that discrete perception is merely an extension of this idea, obtained by assuming that the “periodic engine” keeps running, even when there is only one stimulus to process. This connection between the two theories will be a recurring thread in the present narrative.

Of course, the primary source of evidence about these topics is (and should remain) based on neurophysiological recordings, which provide direct measurements of the neuronal oscillations. For example, we have recently reviewed our past EEG work on the perceptual correlates of ongoing oscillations, with a focus on linking these oscillations to the notion of discrete perception (VanRullen et al., 2011). Another example is a recent study in macaque monkeys revealing that oscillatory neuronal activity in the frontal eye field

(FEF, a region involved in attention and saccade programming) reflects the successive cycles of a clearly sequential attentional exploration process during visual search (Buschman and Miller, 2009). In this review we focus on purely psychophysical methods, not because they are better than direct neurophysiological measurements, but because they also inform us about psychological and perceptual consequences of the postulated periodicities. However, due to the inherent temporal limitations of most psychophysical methods, it should be kept in mind that in practice this approach is probably restricted to the lower end of the frequency spectrum, i.e., oscillations in the delta (0–4 Hz), theta (4–8 Hz), alpha (8–14 Hz), and possibly beta (14–30 Hz) bands. Higher-frequency oscillations (e.g., gamma: 30–80 Hz) may still play a role in sensory processing, but they are generally less amenable to direct psychophysical observation.

There exist many psychophysical paradigms designed to test the temporal limits of sensory systems but that do not directly implicate periodic perception or attention, because their results can also be explained by temporal “smoothing” or integration, in the context of a strictly continuous model of perception (Di Lollo and Wilson, 1978). These paradigms are nonetheless useful to the discrete argument because they constrain the range of plausible periodicities: for example, temporal numerosity judgments or simultaneity judgments indicate that the temporal limit for individuating visual events is only around 10 events per second (White et al., 1953; Lichtenstein, 1961; White, 1963; Holcombe, 2009), suggesting a potential oscillatory correlate in the alpha band (Harter, 1967). We will not develop these results further here, focusing instead on paradigms that unequivocally indicate periodic sampling of perception or attention. Similarly, although some psychophysical studies have demonstrated that perception and attention can be entrained to a low-frequency rhythmic structure in the stimulus sequence (Large and Jones, 1999; Mathewson et al., 2010), we will only concentrate here on studies implicating *intrinsic* perceptual and attentional rhythms (i.e., rhythms that are not present in the stimulus). We will see that progress can be made on the two debates of discrete vs. continuous perception and sequential vs. parallel attention by addressing them together rather than separately.

RHYTHMIC SAMPLING OF A SINGLE STIMULUS: DISCRETE VS. CONTINUOUS PERCEPTION

Suppose that a new stimulus suddenly appears in your visual field, say a red light at the traffic intersection. For such a transient onset, a sequence of visual processing mechanisms from your retina to your high-level visual cortex will automatically come into play, allowing you after a more or less fixed latency to “perceive” this stimulus, i.e., experience it as part of the world in front of you. Hopefully you should then stop at the intersection. What happens next? For as long as the stimulus remains in the visual field, you will continue to experience it. But how do you know it is still there? You might argue that if it were gone, the same process as previously would now signal the transient offset (together with the onset of the green light), and you would then recognize that the red light is gone. But in-between those two moments, you did experience the red light as present – did you only fill in the mental contents of this intervening period after the green light appeared? This sounds unlikely, at least if your traffic lights last as long as they do around here. Maybe the

different stages of your visual system were constantly processing their (unchanged) inputs and feeding their (unchanged) outputs to the next stage, just in case the stimulus might happen to change right then – a costly but plausible strategy. An intermediate alternative would consist in sampling the external world periodically to verify, and potentially update, its contents; the period could be chosen to minimize metabolic effort, while maximizing the chances of detecting any changes within an ecologically useful delay (e.g., to avoid honking from impatient drivers behind you when you take too long to notice the green light). These last two strategies are respectively known as continuous and discrete perception.

The specific logic of the above example may have urged you to favor discrete perception, but the scientific community traditionally sides with the continuous idea. It has not always been so, however. In particular, the first observations of EEG oscillations in the early twentieth century (Berger, 1929), together with the simultaneous popularization of the cinema, prompted many post-war scientists to propose that the role of brain oscillations could be to chunk sensory information into unitary events or “snapshots,” similar to what happens in the movies (Pitts and McCulloch, 1947; Stroud, 1956; Harter, 1967). Much experimental research ensued, which we have already reviewed elsewhere (VanRullen and Koch, 2003). The question was never fully decided, however, and the community’s interest eventually faded. The experimental efforts that we describe in this section all result from an attempt to follow up on this past work and revive the scientific appeal of the discrete perception theory.

PERIODICITIES IN REACTION TIME DISTRIBUTIONS

Some authors have reasoned that if the visual system samples the external world discretely, the time it would take an observer to react after the light turns green would depend on the precise moment at which this event occurred, relative to the ongoing samples: if the stimulus is not detected within one given sample then the response will be delayed at least until the next sampling period. This relation may be visible in histograms of reaction time (RT). Indeed, multiple peaks separated by a more or less constant period are often apparent in RT histograms: these multimodal distributions have been reported with a period of approximately 100 ms for verbal choice responses (Venables, 1960), 10–40 ms for auditory and visual discrimination responses (Dehaene, 1993), 10–15 ms for saccadic responses (Latour, 1967), 30 ms for smooth pursuit eye movement initiation responses (Poppel and Logothetis, 1986). It must be emphasized, however, that an oscillation can only be found in a histogram of post-stimulus RTs if each stimulus either evokes a novel oscillation, or resets an existing one. Otherwise (and assuming that the experiment is properly designed, i.e., with unpredictable stimulus onsets), the moment of periodic sampling will always occur at a random time with respect to the stimulus onset; thus, the peaks of response probability corresponding to the recurring sampling moments will average out, when the histogram is computed over many trials. In other words, even though these periodicities in RT distributions are intriguing, they do not unambiguously demonstrate that perception samples the world periodically – for example, it could just be that each stimulus onset triggers an oscillation in the motor system that will subsequently constrain the response generation process. In the following sections,

we present other psychophysical methods that can reveal perceptual periodicities within *ongoing* brain activity, i.e., without assuming a post-stimulus phase reset.

DOUBLE-DETECTION FUNCTIONS

As illustrated in the previous section, there is an inherent difficulty in studying the perceptual consequences of ongoing oscillations: even if the pre-stimulus oscillatory phase modulates the sensory processing of the stimulus, this pre-stimulus phase will be different on successive repetitions of the experimental trial, and the average performance over many trials will show no signs of the modulation. Obviously, this problem can be overcome if the phase on each trial can be precisely estimated, for example using EEG recordings (VanRullen et al., 2011). With purely psychophysical methods, however, the problem is a real challenge.

An elegant way to get around this challenge has been proposed by Latour (1967). With this method, he showed preliminary evidence that visual detection thresholds could fluctuate along with ongoing oscillations in the gamma range (30–80 Hz). The idea is to present two stimuli on each trial, with a variable delay between them, and measure the observer’s performance for detecting (or discriminating, recognizing, etc.) both stimuli: even if each stimulus’s absolute relation to an ongoing oscillatory phase cannot be estimated, the probability of double-detection should oscillate as a function of the inter-stimulus delay (Figure 1). In plain English, the logic is that when the inter-stimulus delay is a multiple of the oscillatory period, the observer will be very likely to detect both stimuli (if they both fall at the optimal phase of the oscillation) or to miss both stimuli altogether (if they both fall at the opposite phase); on the other hand, if the delay is chosen in-between two multiples of the oscillatory period, then the observer will be very likely to detect only one of the two stimuli (if the first stimulus occurs at the optimal phase, the other will fall at the opposite, and vice-versa).

More formally, let us assume that the probability of measuring our psychological variable ψ (e.g., target detection, discrimination, recognition, etc.) depends periodically (with period $2\pi/\omega$) on the time of presentation of the stimulation s ; to a first approximation this can be noted:

$$p(\psi = 1 | s = 1) = p_0 \cdot (1 + a \cdot \sin(\omega \cdot t)) \tag{1}$$

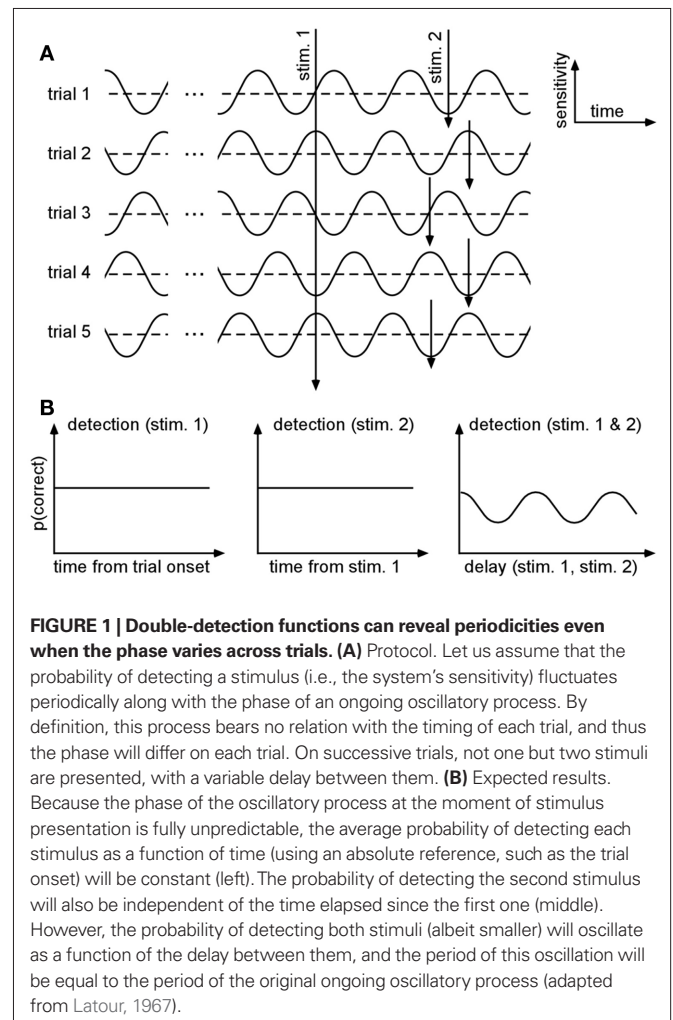
where p_0 is the average expected measurement probability, and a is the amplitude of the periodic modulation. Since the time t of stimulation (with respect to the ongoing oscillation) may change for different repetitions of the measurement, only p_0 can be measured with classical trial averaging methods (i.e., the “sine” term will average out to a mean value of zero). However, if two identical stimulations are presented, separated by an interval δt , the conditional probability of measuring our psychological variable *twice* can be shown to be (there is no room here, unfortunately, for the corresponding mathematical demonstration):

$$p(\psi = 2 | s = 2) = p_0^2 \cdot (1 + ((a^2)/2) \cdot \cos(\omega \cdot \delta t)) \tag{2}$$

The resulting probability only depends on the interval δt (chosen by the experimenter), and thus does not require knowledge of the exact oscillatory phase on every trial. This means that, using double

stimulations and double-detection functions, one can derive *psychophysically* the rate ω of the periodic process, and its modulation amplitude a (Figure 1).

In practice, unfortunately, this method is not as easy to apply as it sounds. One important caveat was already mentioned by Latour: the inter-stimulus delay must be chosen to be long enough to avoid direct interactions between the two stimuli (e.g., masking, apparent motion, etc.). This is because the mathematical derivation of Eq. 2 assumes independence between the detection probabilities for the two stimuli. To ensure that this condition is satisfied, the stimuli should be separated by a few 100 ms (corresponding to the integration period for masking or apparent motion); on the other hand, this implies that several oscillatory cycles will occur between the two stimuli, and many external factors (e.g., phase slip, reset) can thus interfere and decrease the measured oscillation. This in turn suggests that the method may be more appropriate for revealing low-frequency oscillations than high-frequency ones (e.g., gamma). Another important limitation is that the magnitude of the measured oscillation in the double-detection function (2) is squared, compared to the magnitude of the original perceptual oscillation. Although this is not a problem if the perceptual oscillation is strong (i.e., the square of a number close to 1 is also close to 1), it can



become troublesome when the original perceptual oscillation is already subtle (e.g., for a 20% modulation of the visual threshold, one can only expect a 4% modulation in the double-detection function). Altogether, these limitations may explain why Latour's results have, so far, not been replicated or extended.


TEMPORAL ALIASING: THE WAGON WHEEL ILLUSION

Engineers know that any signal sampled by a discrete or periodic system is subject to potential "aliasing" artifacts (**Figure 2**): when the sampling resolution is lower than a critical limit (the "Nyquist rate") the signal can be interpreted erroneously. This is true, for instance, when a signal is sampled in the temporal domain (**Figure 2A**). When this signal is a periodic visual pattern in motion, aliasing produces a phenomenon called the "wagon wheel illusion" (**Figure 2B**): the pattern appears to move in the wrong direction. This is often observed in movies or on television, due to the discrete sampling of video cameras (generally around 24 frames per second). Interestingly, a similar perceptual effect has also been reported under continuous conditions of illumination, e.g., day-

light (Schouten, 1967; Purves et al., 1996; VanRullen et al., 2005b). In this case, however, because no artificial device is imposing a periodic sampling of the stimulus, the logical conclusion is that the illusion must be caused by aliasing within the visual system itself. Thus, this "continuous version of the wagon wheel illusion" (or "c-WWI") has been interpreted as evidence that the visual system samples motion information periodically (Purves et al., 1996; Andrews et al., 2005; Simpson et al., 2005; VanRullen et al., 2005b).

There are many arguments in favor of this "discrete" interpretation of the c-WWI. First, the illusion occurs in a very specific range of stimulus temporal frequencies, compatible with a discrete sampling rate of approximately 13 Hz (Purves et al., 1996; Simpson et al., 2005; VanRullen et al., 2005b). As expected according to the discrete sampling idea, this critical frequency remains unchanged when manipulating the spatial frequency of the stimulus (Simpson et al., 2005; VanRullen et al., 2005b) or the type of motion employed, i.e., rotation vs. translation motion, or first-order vs. second-order motion (VanRullen et al., 2005b). EEG correlates of the perceived illusion confirm these psychophysical findings and point to an oscillation in the same frequency range around 13 Hz (VanRullen et al., 2006; Piantoni et al., 2010). Altogether, these data suggest that (at least part of) the motion perception system proceeds by sampling its inputs periodically, at a rate of 13 samples per second.

There are, of course, alternative accounts of the phenomenon. First, it is noteworthy that the illusion is not instantaneous, and does not last indefinitely, but it is instead a bistable phenomenon, which comes and goes with stochastic dynamics; such a process implies the existence of a competition between neural mechanisms supporting the veridical and the erroneous motion directions (Blake and Logothetis, 2002). Within this context, the debate centers around the source of the erroneous signals: some authors have argued that they arise not from periodic sampling and aliasing, but from spurious activation in low-level motion detectors (Kline et al., 2004; Holcombe et al., 2005) or from motion adaptation processes that would momentarily prevail over the steady input (Holcombe and Seizova-Cajic, 2008; Kline and Eagleman, 2008). We find these accounts unsatisfactory, because they do not seem compatible with the following experimental observations: (i) the illusion is always maximal around the same temporal frequency, whereas the temporal frequency tuning of low-level motion detectors differs widely between first and second-order motion (Hutchinson and Ledgeway, 2006); (ii) not only the magnitude of the illusion, but also its spatial extent and its optimal temporal frequency – which we take as a reflection of the system's periodic sampling rate – are all affected by attentional manipulations (VanRullen et al., 2005b; VanRullen, 2006; Macdonald et al., under review); in contrast, the amount of motion adaptation could be assumed to vary with attentional load (Chaudhuri, 1990; Rezec et al., 2004), but probably not the frequency tuning of low-level motion detectors; (iii) motion adaptation itself can be dissociated from the wagon wheel illusion using appropriate stimulus manipulations; for example, varying stimulus contrast or eccentricity can make the motion aftereffects (both static and dynamic versions) decrease while the c-WWI magnitude increases, and vice-versa (VanRullen, 2007); (iv) finally, the brain regions responsible for the c-WWI effect, repeatedly identified in the right parietal lobe (VanRullen et al., 2006, 2008; Reddy et al., 2011), point to a higher-level cause than the mere adaptation of low-level motion detectors.



Disentangling the neural mechanisms of the continuous wagon wheel illusion could be (and actually, is) the topic of an entirely separate review (VanRullen et al., 2010). To summarize, our current view is that the reversed motion signals most likely originate as a form of aliasing due to periodic temporal sampling by attention-based motion perception systems, at a rate of ~13 Hz; the bistability of the illusion is due to the simultaneous encoding of the veridical motion direction by other (low-level, or “first-order”) motion perception systems. The debate, however, is as yet far from settled. At any rate, this phenomenon illustrates the potential value of temporal aliasing as a paradigm to probe the discrete nature of sensory perception.

OTHER FORMS OF TEMPORAL ALIASING

The sampling frequency evidenced with the c-WWI paradigm may be specific to attention-based motion perception mechanisms. It is only natural to try and extend the temporal aliasing methodology to perception of other types of motion, to perception of visual features other than motion or to perception in sensory modalities other than vision. If evidence for temporal aliasing could be found in these cases, the corresponding sampling frequencies may then be compared to one another and further inform our understanding of discrete perception. Is there a single rhythm, a central (attentional) clock that samples all sensory inputs? Or is information from any single channel of sensory information read out periodically at its own rate, independently from other channels? While the first proposition reflects the understanding that most have of the theory of discrete perception (Kline and Eagleman, 2008), the latter may be a much more faithful description of reality; additionally, the sampling rate for a given channel may vary depending on task demands and attentional state, further blurring intrinsic periodicities.

The simple generic paradigm which we advocate to probe the brain for temporal aliasing is as follows. Human observers are presented with a periodic time-varying input which physically evolves in an unambiguously defined direction; they are asked to make a two-alternative forced choice judgment on the direction of evolution of this input, whose frequency is systematically varied by the experimenter across trials. A consistent report of the wrong direction at a given input frequency may be taken as a behavioral correlate of temporal aliasing, and the frequencies at which this occurs inform the experimenter about the underlying sampling frequency of the brain for this input.

Two main hurdles may be encountered in applying this paradigm. The first one lies in what should be considered a “consistent” report of the wrong direction. Clearly, for an engineered sampling system, one can find input frequencies at which the system will *always* output the wrong direction. For a human observer, however, several factors could be expected to lower the tendency to report the wrong direction, even at frequencies that are subject to aliasing: measurement noise, the potential variability of the hypothetical sampling frequency over the duration of the experiment, and most importantly, the potential presence of alternate sources of information (as in the c-WWI example, where competition occurs between low-level and attention-based motion systems). In the end, even if aliasing occurs, it may not manifest as a clear and reliable percept of the erroneous direction, but rather as a subtle increase of the probability of reporting the wrong direction at certain frequencies.

Recently, we proposed a method to evaluate the presence of aliasing in psychometric functions, based on model fitting (Dubois and VanRullen, 2009). (A write-up of this method and associated findings can be accessed at <http://www.cerco.ups-tlse.fr/~rufin/assc09/>). Results of a 2-AFC motion discrimination experiment were well explained by considering two motion sensing systems, one that functions continuously and one that takes periodic samples of position to infer motion. These two systems each give rise to predictable psychometric functions with few parameters, whose respective contributions to performance can be inferred by model fitting. Evidence for a significant contribution of a discrete process sampling at 13 Hz was found – thus confirming our previous conclusions from the c-WWI phenomenon. Furthermore, the discrete process contributed more strongly to the perceptual outcome when motion was presented inter-ocularly, than binocularly; this is compatible with our postulate that discrete sampling in the c-WWI is a high-level effect, since inter-ocular motion perception depends on higher-level motion perception systems (Lu and Sperling, 2001).

The second pitfall is that the temporal resolution for discriminating the direction of the time-varying input under consideration should be at least as good as the hypothesized sampling frequency. If the psychometric function is already at chance at the frequency where aliasing is expected to take place, this aliasing will simply not be observed – whether the perceptual process relies on periodic sampling or not. Our lab learned this the hard way: many of the features that we experimented with so far, besides luminance and contrast-defined motion, can only be discriminated at low-temporal frequencies – they belong to Holcombe’s “seeing slow” category (Holcombe, 2009). For example, we hypothesized that motion stimuli designed to be invisible to the first-order motion perception system, such as stereo-defined motion (Tseng et al., 2006), would yield maximal aliasing as there is no other motion perception system offering competing information. Unfortunately, these stimuli do not yield a clear percept at temporal frequencies beyond 3–4 Hz, meaning that any aliasing occurring at higher frequencies would have escaped our notice. The “motion standstill” phenomenon reported by Lu and colleagues (Lu et al., 1999; Tseng et al. 2006) with similar stimuli at frequencies around 5 Hz remains a potential manifestation of temporal aliasing, although we have not satisfactorily replicated it in our lab yet. We also hypothesized that binding of spatially distinct feature conjunctions, such as color and motion, could rely on sequential attentional sampling of the two features (Moutoussis and Zeki, 1997), and should thus be subject to aliasing. Again, we were disappointed to find that performance was at chance level at presentation rates higher than 3–4 Hz (Holcombe, 2009), precluding further analysis. We also attempted to adapt the wagon wheel phenomenon to the auditory modality. Here, perception of sound source motion (e.g., a sound rotating around the listener) also appeared limited to about 3 Hz (Feron et al., 2010). We then reasoned that frequency, rather than spatial position, was the primary feature for auditory perception, and designed periodic stimuli that moved in particular directions in the frequency domain – so-called Shepard or Risset sequences (Shepard, 1964). Again, we found that the direction of these periodic frequency sweeps could not be identified when the temporal frequency of presentation was increased beyond 3–4 Hz.

In sum, although temporal aliasing is, in principle, a choice paradigm to probe the rhythms of perception, our attempts so far at applying this technique to other perceptual domains than motion (the c-WWI) have been foiled by the strict temporal limits of the corresponding sensory systems. What we can safely conclude is that, if discrete sampling exists in any of these other perceptual domains, it will be at a sampling rate above 3–4 Hz. We have not exhausted all possible stimuli and encourage others to conduct their own experiments. There are two faces to the challenge: finding stimuli that the brain “sees fast” enough, and using an appropriate model to infer the contribution of periodic sampling to the psychometric performance (in case other sources of information and sizeable variability across trials should blur the influence of discrete processes).

RHYTHMIC SAMPLING OF MULTIPLE STIMULI: SEQUENTIAL VS. PARALLEL ATTENTION

Let us return to our previous hypothetical situation. Now you have passed the traffic lights and driven home, and you turned on the TV to find out today’s lottery numbers. There are a handful of channels that can provide this information at this hour, so you go to “multi-channel” mode to monitor them simultaneously. The lottery results are not on, so you will wait until any channel shows them. How will you know which one? You try to process all channels at once, but their contents collide and confuse you. What if one of them shows the numbers but you notice it too late? By focusing on a single channel you would be sure not to miss the first numbers, but then what if you picked the wrong channel? In such a situation, it is likely that you will switch your attention rapidly between the different candidate channels until you see one that provides the required information. Your brain often faces the same problem when multiple objects are present in the visual field and their properties must be identified, monitored or compared.

A LONG-STANDING DEBATE

Visual search

Whether your brain simultaneously and continuously shares its attentional resources (i.e., in “parallel”) between candidate target objects, or switches rapidly and sequentially between them, has been the subject of intense debate in the literature. We refer to this debate as “parallel vs. sequential attention.” Originally, attention was assumed to be a unitary, indivisible resource, and consequently the sequential model was favored, often implicitly (Eriksen and Spencer, 1969; Treisman, 1969; Kahneman, 1973). The first two decades of studies using the visual search paradigm were heavily biased toward this assumption (Treisman and Gelade, 1980; Wolfe, 1998): when a target had to be identified among a varying number of distractors and the observer’s RT was found to increase steadily with the number of items, it was assumed that the additional time needed for each item reflected the duration of engaging, sampling and disengaging attention (hence the term “serial search slope”). It was only in the 1990s that this assumption was seriously challenged by proponents of an alternate model, according to which attention is always distributed in parallel among items, and the increase in RT with increasing item number simply reflects the increasing task difficulty or decreasing

“signal-to-noise ratio” (Palmer, 1995; Carrasco and Yeshurun, 1998; Eckstein, 1998; McElree and Carrasco, 1999; Eckstein et al., 2000). Both models are still contemplated today – and indeed, they are extremely difficult to distinguish experimentally (Townsend, 1990).

Multi-object tracking

The same difficulty also plagues paradigms other than visual search. Multi-object tracking, for instance, corresponds to a situation in which several target objects are constantly and randomly moving around the visual field, often embedded among similarly moving distractors (Pylyshyn and Storm, 1988). Sometimes, the objects are moving in feature space (i.e., changing their color or their orientation) rather than in physical space (Blaser et al., 2000). The common finding that up to four – or sometimes more (Cavanagh and Alvarez, 2005) objects can be efficiently tracked at the same time has been taken as evidence that attention must be divided in parallel among the targets (Pylyshyn and Storm, 1988). However, in the limit where attention would be assumed to move at lightning speed, it is obvious that this simultaneous tracking capacity could be explained equally well by sequential shifts of a single attention spotlight, than by divided or parallel attention. Indeed, at least some of the existing data have been found compatible with a sequential process (Howard and Holcombe, 2008; Oksama and Hyona, 2008). Since there is no general agreement concerning the actual speed of attention (Duncan et al., 1994; Moore et al., 1996; Hogendoorn et al., 2007), the question remains open.

Simultaneous/sequential paradigm

Other paradigms have been designed with the explicit aim of teasing apart the parallel and serial attention models. In the simultaneous/sequential paradigm, the capacity of attention to process multiple items simultaneously is assessed by presenting the relevant information for a limited time in each display cycle. In one condition (simultaneous) this information is delivered at once for all items; in the other condition (sequential) each item’s information is revealed independently, at different times. In both conditions the critical information is thus shown for the same total amount of time, such that a parallel model of attentional allocation would predict comparable performance; however, a serial attentional model would suffer more in the simultaneous condition, because attention would necessarily miss the relevant information in one stimulus while it is sampling the other(s), and vice-versa (Eriksen and Spencer, 1969; Shiffrin and Gardner, 1972). The paradigm was recently applied to the problem of multiple-object tracking (Howe et al., 2010), and the data were deemed incompatible with serial attention sampling. A major source of confounds in this paradigm, however, is that, depending on stimulus arrangement parameters, certain factors (e.g., grouping, crowding) can improve or decrease performance in the simultaneous condition independently of attention; similarly, other factors (e.g., apparent motion, masking) can improve or decrease performance in the sequential condition. It is unclear in the end how to tease apart the effects of attention from the potentially combined effects of all these extraneous factors.

Split spotlight studies

To finish, there is yet another class of experiments that were intended to address a distinct albeit related question: when attention is divided among multiple objects, does the focus simply expand its size to include all of the targets, or does it split into several individual spotlights? To test this, these paradigms generally measure an improvement of performance due to attention at two separate locations concurrently; the critical test is then whether a similar improvement can also be observed at an intervening spatial location: if yes, the spotlight may have been simply enlarged, if not it may have been broken down into independent spotlights. Psychophysical studies tend to support the multiple spotlights account (Bichot et al., 1999; Awh and Pashler, 2000). The same idea has also been applied to physiological measurements of the spotlight, demonstrating that EEG or fMRI brain activations can be enhanced by attention at two concurrent locations, without any enhancement at intervening locations (Muller et al., 2003; McMains and Somers, 2004). Now, how do these results on the spatial deployment of attention pertain to our original question about the temporal dynamics of attention? Inherent in the logic of this paradigm is the assumption that attentional resources are divided constantly over time; multiple spotlights are implicitly assumed to operate simultaneously, rather than as a single, rapidly shifting focus. To support this assumption, authors often use limited presentation times (so attention does not have time to shift between targets), and speculate on the speed of attention. As mentioned before, since this speed is largely unknown, a lot of the data remain open to interpretation. In fact, our recent results in a very similar paradigm (in which we varied the delay between stimulus onset and the subsequent measurement of attentional deployment) showed that multiple simultaneous spotlights can in fact be observed, but are short-lived; when several target locations need to be monitored for extended periods of time, the attentional system quickly settles into a single-spotlight mode (Dubois et al., 2009). In another related study we found that attention could not simultaneously access information from two locations, but instead relied on rapid sequential allocation (Hogendoorn et al., 2010).

The conclusion from studies that have used this kind of paradigm is also fairly representative of the current status of the entire “sequential vs. parallel attention” debate, which we have briefly reviewed here. As summarized in a recent (and more thorough) review by Jans et al. (2010), most of the so-called demonstrations of multiple parallel attention spotlights rely on strong – and often unsubstantiated – assumptions about the temporal dynamics of attention. In sum, parallel attention has by no means won the prize.

TEMPORAL ALIASING RETURNS

Could aliasing (see Temporal Aliasing: the Wagon Wheel Illusion and Other Forms of Temporal Aliasing) provide a way of resolving the “sequential vs. parallel attention” debate? If attention focuses on each target sequentially rather than continuously, the target information will be sampled more or less periodically, and should thus be subject to aliasing artifacts; furthermore, the rate of sampling for each target should be inversely related to the number of targets to sample (i.e., the “set size”). On the other hand, parallel attention models have no reason to predict aliasing; and, even if aliasing were to occur, no reason to predict a change of aliasing frequency as a function of set size. We recently tested this idea using a variant of

the continuous wagon wheel illusion phenomenon (Macdonald et al., under review). On each 40 s trial, we showed one, two, three or four wheels rotating in the same direction; the frequency of rotation was varied between trials. From time to time, a subset of the wheels briefly reversed their direction, and the subjects’ task was to count and report how many of these reversal events had occurred in each trial. We reasoned that any aliasing would be manifested as an overestimation of the number of reversals. As expected from our experiments with the c-WWI, we found significant aliasing in a specific range of rotation frequencies. Most importantly, the frequency of maximum aliasing significantly decreased as set size was increased, as predicted by the “sequential attention” idea. Although the magnitude of this decrease was lower than expected (the sampling frequency was approximately divided by 2, from ~13 Hz down to ~7 Hz, while the set size was multiplied by 4), this finding poses a very serious challenge to the “parallel attention” theory.

THE BLINKING SPOTLIGHT OF ATTENTION

The crucial difficulty in distinguishing parallel and sequential accounts of attention is a theoretical one: any “set size effect” that can be explained by a division of attentional resources in time can, in principle, be explained equally well using a spatial division of the same resources (Townsend, 1990; Jans et al., 2010). There is a form of equivalence between the temporal and spatial domains, a sort of Heisenberg uncertainty principle, that precludes most attempts at jointly determining the spatial and temporal distributions of attention. In a recent experiment, however, we tried to break down this equivalence by measuring set size effects following a *temporal* manipulation of the stimuli – namely, after varying their effective duration on each trial (VanRullen et al., 2007). Thus we could tell, for example, how well three stimuli were processed when they were presented for 300 ms, and we could compare this to the performance obtained for one stimulus presented 100, or 300 ms (or any other combination of set size and duration). The interest of this procedure was that different models of attentional division (e.g., sequential and parallel models) would make different predictions about how the psychometric function for processing one stimulus as a function of its duration should translate into corresponding psychometric functions for larger set sizes. To simplify, the space–time equivalence was broken, because performance for sequentially sampling two (or three, or four) stimuli each for a fixed period could be predicted exactly, by knowing the corresponding performance for a single stimulus that lasted for a duration equivalent to the sampling period; a simple parallel model, of course, could be designed to explain the change in performance from one to two (or three, or four) stimuli, but if the model was wrong it would then do a poor job at explaining performance obtained at other set sizes.

We compared three distinct models of attention allocation, each with a single free parameter. In the “parallel” model, all stimuli were processed simultaneously, and only the efficiency of this processing varied as a function of attentional load (i.e., set size); the cost in efficiency was manipulated by the model’s free parameter. The second model, coined “sample-when-divided,” corresponded to the classic idea of a switching spotlight: when more than one stimulus was present, the otherwise constant attentional resource was forced to sequentially sample the stimulus locations; the model’s free parameter was its sampling period, which affected its ability

to process multiple stimuli. Finally, we decided to consider a third model, termed “sample-always,” similar to the previous one except for the fact that it still collected and integrated successive attentional samples even when a single stimulus was present (see **Figure 3** for an illustration); this model’s behavior was also governed by its sampling period. Our strategy was, then, to compare the three models’ ability to emulate the actual psychometric functions of human observers.

Compatible with existing findings (Palmer, 1995), we revealed that the parallel model could outperform the classic version of the sequential model – the “sample-when-divided” one, which considers that the attentional spotlight shifts around sequentially, but only when attention must be divided. However, the truly optimal model to explain human psychometric functions was the other variant of the sequential idea, a model in which attention always samples information periodically, regardless of set size. The rate of sampling was found to be ~ 7 Hz. When attention is divided, successive samples naturally focus on different stimuli, but when it is concentrated on a single target, the samples continue to occur repeatedly every ~ 150 ms, simply accumulating evidence for this one stimulus. In other words, our findings supported a “blinking spotlight” of attention (VanRullen et al., 2007) over the sequentially “switching spotlight” (and over the multiple “parallel spotlights”).




FIGURE 3 | Relating discrete perception with sequential attention. (A) A sensory process that samples a single visual input periodically illustrates the concept of discrete perception. **(B)** A sensory process that serially samples three simultaneously presented visual stimuli demonstrates the notion of a sequential attention spotlight. Since many of our findings implicate attention in the periodic sampling processes displayed in **(A)**, we propose that both types of periodic psychological operations **(A,B)** actually reflect a common oscillatory neuronal process. According to this view, the spotlight of attention is intrinsically rhythmic, which gives it both the ability to rapidly scan multiple objects, and to discretely sample a single source. (The yellow balls linked by red lines illustrate successive attentional samples.)

REFERENCES

- Allport, D. A. (1968). Phenomenal simultaneity and the perceptual moment hypothesis. *Br. J. Psychol.* 59, 395–406.
- Andrews, T., Purves, D., Simpson, W. A., and VanRullen, R. (2005). The wheel keeps turning: reply to Holcombe et al. *Trends Cogn. Sci. (Regul. Ed.)* 9, 561.
- Awh, E., and Pashler, H. (2000). Evidence for split attentional foci. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 834–846.
- Berger, H. (1929). Über das Elektroenkephalogramm des Menschen. *Arch. Psychiatr. Nervenkr.* 87, 527–570.
- Bichot, N. P., Cave, K. R., and Pashler, H. (1999). Visual selection mediated by location: feature-based selection of noncontiguous locations. *Percept. Psychophys.* 61, 403–423.
- Blake, R., and Logothetis, N. K. (2002). Visual competition. *Nat. Rev. Neurosci.* 3, 13–21.
- Blaser, E., Pylyshyn, Z. W., and Holcombe, A. O. (2000). Tracking an object through feature space. *Nature* 408, 196–199.
- Buschman, T. J., and Miller, E. K. (2009). Serial, covert shifts of attention during visual search are reflected by the frontal eye fields and correlated with population oscillations. *Neuron* 63, 386–396.

CONCLUSION

The notion of a blinking spotlight illustrates a fundamental point: that discrete sampling and sequential attention could be two facets of a single process (**Figure 3**). Proponents of the discrete sampling theory should ask themselves: what happens when there are more than one relevant stimuli in the visual field? Can they all be processed in a single “snapshot”? Advocates of sequential attention should ponder about the behavior of attention when it has only one target to monitor: is it useful – or even possible – for attention to pause its exploratory dynamics?

The simple theory that we propose is that periodic “covert” attentional sampling may have evolved from “overt” exploratory behavior (i.e., eye movements), as a means to quickly and effortlessly scan internal representations of the environment (VanRullen et al., 2005a; Uchida et al., 2006). Just as eye movements continue to occur even when there is only one object in the scene – lest the object quickly fade from awareness (Ditchburn and Ginsborg, 1952; Coppola and Purves, 1996; Martinez-Conde et al., 2006) – it is sensible to posit that attentional sampling takes place regardless of the number of objects to sample. Perception can then be said to be “discrete” or “periodic,” insofar as a very significant portion of its inputs (those depending on attentional mechanisms) are delivered periodically. For example, the ~ 13 Hz discrete sampling responsible for the continuous wagon wheel illusion was found to be driven by attention (VanRullen et al., 2005b; VanRullen, 2006; Macdonald et al., under review). The frequency of this sampling progressively decreased to ~ 7 Hz when two, three, and finally four “wagon wheel” stimuli had to be simultaneously monitored (Macdonald et al., under review). Interestingly, this ~ 7 Hz periodicity was also the one indicated by our model of the “blinking spotlight” of attention (VanRullen et al., 2007). Altogether, our data raise the intriguing suggestion that attention creates discrete samples of the visual world with a periodicity of approximately one tenth of a second.

To conclude, we argue that it is constructive to unite the two separate psychophysical debates about discrete vs. continuous perception (see Rhythmic Sampling of a Single Stimulus: Discrete vs. Continuous Perception) and sequential vs. parallel attention (see Rhythmic Sampling of Multiple Stimuli: Sequential vs. Parallel Attention). Discrete perception and sequential attention may represent perceptual and psychological manifestations of a single class of periodic neuronal mechanisms. Therefore, psychophysical progress in solving those debates could, ultimately, contribute to uncovering the role of low-frequency brain rhythms in perception and attention.

ACKNOWLEDGMENT

This research was funded by a EURYI Award and an ANR grant 06JCJC-0154 to RV.

- Buzsaki, G. (2006). *Rhythms of the Brain*. New York: Oxford University Press.
- Carrasco, M., and Yeshurun, Y. (1998). The contribution of covert attention to the set-size and eccentricity effects in visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 24, 673–692.
- Cavanagh, P., and Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends Cogn. Sci. (Regul. Ed.)* 9, 349–354.
- Chaudhuri, A. (1990). Modulation of the motion aftereffect by selective attention. *Nature* 344, 60–62.
- Coppola, D., and Purves, D. (1996). The extraordinarily rapid disappearance of entopic images. *Proc. Natl. Acad. Sci. U.S.A.* 93, 8001–8004.
- Dehaene, S. (1993). Temporal oscillations in human perception. *Psychol. Sci.* 4, 264–270.
- Di Lollo, V., and Wilson, A. E. (1978). Iconic persistence and perceptual moment as determinants of temporal integration in vision. *Vision Res.* 18, 1607–1610.
- Ditchburn, R. W., and Ginsborg, B. L. (1952). Vision with a stabilized retinal image. *Nature* 170, 36–37.
- Dubois, J., Hamker, F. H., and VanRullen, R. (2009). Attentional selection of noncontiguous locations: the spotlight is only transiently “split.” *J. Vis.* 9, 1–11.
- Dubois, J., and VanRullen, R. (2009). Evaluating the contribution of discrete perceptual mechanisms to psychometric performance. *Paper presented at the Association for the Scientific Study of Consciousness*, Berlin.
- Duncan, J., Ward, R., and Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature* 369, 313–315.
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol. Cybern.* 60, 121–130.
- Eckstein, M. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing. *Psychol. Sci.* 9, 111–118.
- Eckstein, M., Thomas, J. P., Palmer, J., and Shimozaki, S. S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Percept. Psychophys.* 62, 425–451.
- Engel, A. K., Kreiter, A. K., Konig, P., and Singer, W. (1991). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proc. Natl. Acad. Sci. U.S.A.* 88, 6048–6052.
- Engel, A. K., and Singer, W. (2001). Temporal binding and the neural correlates of sensory awareness. *Trends Cogn. Sci. (Regul. Ed.)* 5, 16–25.
- Eriksen, C. W., and Spencer, T. (1969). Rate of information processing in visual perception: some results and methodological considerations. *J. Exp. Psychol.* 79, 1–16.
- Fell, J., Fernandez, G., Klaver, P., Elger, C. E., and Fries, P. (2003). Is synchronized neuronal gamma activity relevant for selective attention? *Brain Res. Brain Res. Rev.* 42, 265–272.
- Feron, F. X., Frissen, I., Boissinot, J., and Guastavino, C. (2010). Upper limits of auditory rotational motion perception. *J. Acoust. Soc. Am.* 128, 3703–3714.
- Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563.
- Gail, A., Brinksmeyer, H. J., and Eckhorn, R. (2004). Perception-related modulations of local field potential power and coherence in primary visual cortex of awake monkey during binocular rivalry. *Cereb. Cortex* 14, 300–313.
- Gold, I. (1999). Does 40-Hz oscillation play a role in visual consciousness? *Conscious Cogn.* 8, 186–195.
- Gray, C. M., Konig, P., Engel, A. K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338, 334–337.
- Harter, M. R. (1967). Excitability cycles and cortical scanning: a review of two hypotheses of central intermissivity in perception. *Psychol. Bull.* 68, 47–58.
- Hogendoorn, H., Carlson, T. A., VanRullen, R., and Verstraten, F. A. (2010). Timing divided attention. *Atten. Percept. Psychophys.* 72, 2059–2068.
- Hogendoorn, H., Carlson, T. A., and Verstraten, F. A. (2007). The time course of attentive tracking. *J. Vis.* 7, 1–10.
- Holcombe, A. O. (2009). Seeing slow and seeing fast: two limits on perception. *Trends Cogn. Sci. (Regul. Ed.)* 13, 216–221.
- Holcombe, A. O., Clifford, C. W., Eagleman, D. M., and Pakarian, P. (2005). Illusory motion reversal in tune with motion detectors. *Trends Cogn. Sci. (Regul. Ed.)* 9, 559–560.
- Holcombe, A. O., and Seizova-Cajic, T. (2008). Illusory motion reversals from unambiguous motion with visual, proprioceptive, and tactile stimuli. *Vision Res.* 48, 1743–1757.
- Howard, C. J., and Holcombe, A. O. (2008). Tracking the changing features of multiple objects: progressively poorer perceptual precision and progressively greater perceptual lag. *Vision Res.* 48, 1164–1180.
- Howe, P. D., Cohen, M. A., Pinto, Y., and Horowitz, T. S. (2010). Distinguishing between parallel and serial accounts of multiple object tracking. *J. Vis.* 10, 11.
- Hutchinson, C. V., and Ledgeway, T. (2006). Sensitivity to spatial and temporal modulations of first-order and second-order motion. *Vision Res.* 46, 324–335.
- Jans, B., Peters, J. C., and De Weerd, P. (2010). Visual spatial attention to multiple locations at once: the jury is still out. *Psychol. Rev.* 117, 637–684.
- Jensen, O., Gelfand, J., Kounios, J., and Lisman, J. E. (2002). Oscillations in the alpha band (9–12 Hz) increase with memory load during retention in a short-term memory task. *Cereb. Cortex* 12, 877–882.
- Kahana, M. J., Seelig, D., and Madsen, J. R. (2001). Theta returns. *Curr. Opin. Neurobiol.* 11, 739–744.
- Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall.
- Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Res. Brain Res. Rev.* 29, 169–195.
- Kline, K., and Eagleman, D. M. (2008). Evidence against the temporal subsampling account of illusory motion reversal. *J. Vis.* 8, 11–15.
- Kline, K., Holcombe, A. O., and Eagleman, D. M. (2004). Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field. *Vision Res.* 44, 2653–2658.
- Koch, C., and Braun, J. (1996). Towards the neuronal correlate of visual awareness. *Curr. Opin. Neurobiol.* 6, 158–164.
- Large, E. W., and Jones, M. R. (1999). The dynamics of attending: how people track time-varying events. *Psychol. Rev.* 106, 119–159.
- Latour, P. L. (1967). Evidence of internal clocks in the human operator. *Acta Psychol. (Amst)* 27, 341–348.
- Lichtenstein, M. (1961). Phenomenal simultaneity with irregular timing of components of the visual stimulus. *Percept. Mot. Skills* 12, 47–60.
- Lisman, J. E., and Idiart, M. A. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267, 1512–1515.
- Lu, Z. L., Lesmes, L. A., and Sperling, G. (1999). Perceptual motion standstill in rapidly moving chromatic displays. *Proc. Natl. Acad. Sci. U.S.A.* 96, 15374–15379.
- Lu, Z. L., and Sperling, G. (2001). Three-systems theory of human visual motion perception: review and update. *J. Opt. Soc. Am. A. Opt. Image Sci. Vis.* 18, 2331–2370.
- Martinez-Conde, S., Macknik, S. L., Troncoso, X. G., and Dyar, T. A. (2006). Microsaccades counteract visual fading during fixation. *Neuron* 49, 297–305.
- Mathewson, K. E., Fabiani, M., Gratton, G., Beck, D. M., and Lleras, A. (2010). Rescuing stimuli from invisibility: inducing a momentary release from visual masking with pre-target entrainment. *Cognition* 115, 186–191.
- McElree, B., and Carrasco, M. (1999). The temporal dynamics of visual search: evidence for parallel processing in feature and conjunction searches. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 1517–1539.
- McMains, S. A., and Somers, D. C. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron* 42, 677–686.
- Moore, C. M., Egeth, H., Berglan, L. R., and Luck, S. J. (1996). Are attentional dwell times inconsistent with serial visual search? *Psychon. Bull. Rev.* 3, 360–365.
- Moutoussis, K., and Zeki, S. (1997). A direct demonstration of perceptual asynchrony in vision. *Proc. R. Soc. Lond. B Biol. Sci.* 264, 393–399.
- Muller, M. M., Malinowski, P., Gruber, T., and Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature* 424, 309–312.
- Niebur, E., Koch, C., and Rosin, C. (1993). An oscillation-based model for the neuronal basis of attention. *Vision Res.* 33, 2789–2802.
- Oksama, L., and Hyona, J. (2008). Dynamic binding of identity and location information: a serial model of multiple identity tracking. *Cogn. Psychol.* 56, 237–283.
- Palmer, J. (1995). Attention in visual search: distinguishing four causes of a set-size effect. *Curr. Dir. Psychol. Sci.* 4, 118–123.
- Piantoni, G., Kline, K. A., and Eagleman, D. M. (2010). Beta oscillations correlate with the probability of perceiving rivalrous visual stimuli. *J. Vis.* 10, 18.
- Pitts, W., and McCulloch, W. S. (1947). How we know universals: the perception of auditory and visual forms. *Bull. Math. Biophys.* 9, 127–147.
- Poppel, E., and Logothetis, N. (1986). Neuronal oscillations in the human brain. Discontinuous initiations of pursuit eye movements indicate a 30-Hz temporal framework for visual information processing. *Naturwissenschaften* 73, 267–268.
- Purves, D., Paydarfar, J. A., and Andrews, T. J. (1996). The wagon wheel illusion

- in movies and reality. *Proc. Natl. Acad. Sci. U.S.A.* 93, 3693–3697.
- Pylyshyn, Z. W., and Storm, R. W. (1988). Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spat. Vis.* 3, 179–197.
- Ray, S., and Maunsell, J. H. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* 67, 885–896.
- Reddy, L., Remy, F., Vayssiere, N., and VanRullen, R. (2011). Neural correlates of the continuous wagon wheel illusion: a functional MRI study. *Hum. Brain Mapp.* 32, 163–170.
- Revsuo, A., Wilenius-Emet, M., Kuusela, J., and Lehto, M. (1997). The neural generation of a unified illusion in human vision. *Neuroreport* 8, 3867–3870.
- Rezec, A., Krekelberg, B., and Dobkins, K. R. (2004). Attention enhances adaptability: evidence from motion adaptation experiments. *Vision Res.* 44, 3035–3044.
- Sanford, J. A. (1971). “A periodic basis for perception and action,” in *Biological Rhythms and Human Performance*, ed. W. Colquhoun (New York: Academic Press), 179–209.
- Schouten, J. F. (1967). “Subjective stroboscopy and a model of visual movement detectors,” in *Models for the Perception of Speech and Visual Form I*, ed. Wathen-Dunn (Cambridge, MA: MIT Press), 44–45.
- Shepard, R. N. (1964). Circularity in judgments of relative pitch. *J. Acoust. Soc. Am.* 36, 2346–2353.
- Shiffrin, R. M., and Gardner, G. T. (1972). Visual processing capacity and attentional control. *J. Exp. Psychol.* 93, 72–82.
- Simpson, W. A., Shahani, U., and Manahilov, V. (2005). Illusory percepts of moving patterns due to discrete temporal sampling. *Neurosci. Lett.* 375, 23–27.
- Singer, W., and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annu. Rev. Neurosci.* 18, 555–586.
- Stroud, J. M. (1956). “The fine structure of psychological time,” in *Information Theory in Psychology*, ed. H. Quastler (Chicago, IL: Free Press), 174–205.
- Tallon-Baudry, C., Bertrand, O., Delpuech, C., and Pernier, J. (1997). Oscillatory gamma-band (30–70 Hz) activity induced by a visual search task in humans. *J. Neurosci.* 17, 722–734.
- Tallon-Baudry, C., Bertrand, O., Peronnet, F., and Pernier, J. (1998). Induced gamma-band activity during the delay of a visual short-term memory task in humans. *J. Neurosci.* 18, 4244–4254.
- Townsend, J. (1990). Serial vs. parallel processing: sometimes they look like tweedledum and tweedledee but they can (and should) be distinguished. *Psychol. Sci.* 1, 46–54.
- Treisman, A. (1969). Strategies and models of selective attention. *Psychol. Rev.* 76, 282–299.
- Treisman, A., and Gelade, G. (1980). A feature-integration theory of attention. *Cogn. Psychol.* 12, 97–136.
- Tseng, C. H., Gobell, J. L., Lu, Z. L., and Sperling, G. (2006). When motion appears stopped: stereo motion standstill. *Proc. Natl. Acad. Sci. U.S.A.* 103, 14953–14958.
- Uchida, N., Kepecs, A., and Mainen, Z. F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat. Rev. Neurosci.* 7, 485–491.
- VanRullen, R. (2006). The continuous wagon wheel illusion is object-based. *Vision Res.* 46, 4091–4095.
- VanRullen, R. (2007). The continuous wagon wheel illusion depends on, but is not identical to neuronal adaptation. *Vision Res.* 47, 2143–2149.
- VanRullen, R., Busch, N. A., Drewes, J., and Dubois, J. (2011). Ongoing EEG phase as a trial-by-trial predictor of perceptual and attentional variability. *Front. Psychology* 2:60. doi: 10.3389/fpsyg.2011.00060
- VanRullen, R., Carlson, T., and Cavanagh, P. (2007). The blinking spotlight of attention. *Proc. Natl. Acad. Sci. U.S.A.* 104, 19204–19209.
- VanRullen, R., Guyonneau, R., and Thorpe, S. J. (2005a). Spike times make sense. *Trends Neurosci.* 28, 1–4.
- VanRullen, R., Reddy, L., and Koch, C. (2005b). Attention-driven discrete sampling of motion perception. *Proc. Natl. Acad. Sci. U.S.A.* 102, 5291–5296.
- VanRullen, R., and Koch, C. (2003). Is perception discrete or continuous? *Trends Cogn. Sci.* 7, 207–213.
- VanRullen, R., Pascual-Leone, A., and Battelli, L. (2008). The continuous wagon wheel illusion and the ‘when’ pathway of the right parietal lobe: a repetitive transcranial magnetic stimulation study. *PLoS ONE* 3, e2911. doi: 10.1371/journal.pone.0002911
- VanRullen, R., Reddy, L., and Koch, C. (2006). The continuous wagon wheel illusion is associated with changes in electroencephalogram power at approximately 13 Hz. *J. Neurosci.* 26, 502–507.
- VanRullen, R., Reddy, L., and Koch, C. (2010). “A motion illusion revealing the temporally discrete nature of awareness,” in *Space and Time in Perception and Action*, ed. R. Nijhawan (Cambridge: Cambridge University Press), 521–535.
- Venables, P. H. (1960). Periodicity in reaction time. *Br. J. Psychol.* 51, 37–43.
- von der Malsburg, C. (1995). Binding in models of perception and brain function. *Curr. Opin. Neurobiol.* 5, 520–526.
- White, C. (1963). Temporal numerosity and the psychological unit of duration. *Psychol. Monogr.* 77, 1–37.
- White, C. T., Cheatham, P. G., and Armington, J. C. (1953). Temporal numerosity: II. Evidence for central factors influencing perceived number. *J. Exp. Psychol.* 46, 283–287.
- Wolfe, J. M. (1998). “Visual Search” in *Attention*, ed. H. Pashler (London: University College London Press), 13–73.
- Womelsdorf, T., and Fries, P. (2007). The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* 17, 154–160.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 07 March 2011; accepted: 09 August 2011; published online: 29 August 2011.

Citation: VanRullen R and Dubois J (2011) The psychophysics of brain rhythms. *Front. Psychology* 2: 203. doi: 10.3389/fpsyg.2011.00203

This article was submitted to *Frontiers in Perception Science*, a specialty of *Frontiers in Psychology*.

Copyright © VanRullen and Dubois. This is an open-access article subject to a non-exclusive license between the authors and Frontiers Media SA, which permits use, distribution and reproduction in other forums, provided the original authors and source are credited and other Frontiers conditions are complied with.

III. EEG evidence of rhythms in attention and perception

In the general introduction of this thesis, we motivated our quest for periodic mechanisms in perception with the observation of rhythmic activity throughout the brain when electrodes are stuck to the scalp. Yet, since then, we haven't looked at any recording of brain activity... We spent three chapters trying to probe behavior for signs of intrinsic rhythms in the workings of attention and perception. Many more clever (and cleverer) behavioral experiments can (and will) be designed... However, as we argued previously, it would be tremendously helpful to be able to determine the state of the brain at any point in time, so that we can relate perceptual outcome to this instantaneous state – this is the critical test to distinguish the discrete perception hypothesis from a mere integration window. Hence, as we enter the second part of this thesis, it is high time brain activity entered the picture.

The working hypothesis throughout this thesis is that attention is a periodic process, and takes samples which are the building blocks of perception. Testing this hypothesis involves performing analyses that significantly depart from the classical approach in cognitive neuroscience. Since the field started, the main approach has been to show a few different stimuli to the brain, perform many repetitions of the same condition, and look at the average “responses”, i.e. brain activity evoked by the different stimuli. This angle of attack has been justified to get rid of the pervasive trial-to-trial variability, which was hindering any serious scientific progress. Neuroscientists decided that the brain was an extremely noisy device, and that evidencing its key functionalities could only be done by discarding the “noise”. This approach has been extremely successful and taught us a lot: receptive fields of neurons along the visual pathway, modularity, attentional effects,... basically everything that you can find in textbooks and that I succinctly described in the introduction of this thesis. But it is quite disturbing that we routinely throw out 95-99% of the recorded activity in order to study these tiny modulations in detail (the magnitude of a typical ERP or BOLD signal change represents roughly a 1-5% modulation of the ongoing activity). Is the brain really wasting this much energy for nothing? Evidently not. Studying the state of the brain prior to stimulus presentation constitutes a major paradigm shift, but it is getting increasingly popular. A review of the entirety of the corresponding literature is not within the scope of this part of the thesis; we will focus on a subset of the existing studies in section B (page 103). The approach that we advocate here consists in studying pre-stimulus oscillatory activity and using it to predict perceptual outcome. In a given frequency band, oscillatory activity may be studied in one of two ways : either one focuses on the amplitude of the oscillation (a measure of the degree of synchronization of the underlying neuronal population), or one focuses on the phase of the oscillation (a measure of the instantaneous state of the synchronized population, i.e. more or less depolarized). The degree of synchronization inferred from the amplitude of the oscillation is of great importance : if there is no synchronization at all, there is no sense in looking at the exact state of the synchronized activity... However, for our purpose of uncovering periodicities in perceptual processes, we are interested in the exact timing of stimulation relative to the instantaneous state of the synchronized neuronal population, and how this may affect perceptual outcome. We will thus focus on studying the phase of ongoing oscillations in relation with perception. We recently wrote a methodological review summarizing our approach, which you can find at the end of this section (page 150). For now, I propose to reformulate some key points discussed in that review, as both section B (page 103) and section C (page 118) rely heavily on these methods. A complementary approach to support our hypothesis

consists in finding electrophysiological evidence for the periodic nature of attention in a divided attention paradigm. section D (page 131) reports on an attempt at finding real time electrophysiological correlates of the position of the attentional spotlight, in order to follow its behavior in the context of divided attention. In essence, we come back to the debates that we have already addressed in the first part of the thesis, with the additional tools afforded by looking “under the hood”.

A. Methods for linking ongoing EEG phase to trial-by-trial variability

In a typical EEG experiment (this is also readily applicable to magnetoencephalography, MEG, or recordings of local field potentials, LFP), the experimenter will collect two parallel datasets : on the one hand, a continuous recording of brain activity (sampled at 500 to 1000hz); on the other hand, behavioral data – e.g. typically a button press made by the subject at the end of each trial, or the time it takes to make that response, ... In other lines of research, the term “behavioral” data could be extended to encompass other kinds of measurements, such as autonomic responses (skin conductance, pupil dilation). Our goal is to find a pre-stimulus oscillation whose phase predicts the behavioral variable – or, put slightly more reasonably (when a prediction is right only 55% of the time, does it really deserve to be called a prediction?), whose phase explains away some of the variance of the behavioral variable : for a given trial, knowing the phase of this oscillation at the time when the stimulus is presented (or slightly before) provides information about the outcome of the behavioral variable.

The behavioral variable may be discrete, taking one of a few possible states on each trial. This would for instance be the case for the report of a simple perceptual decision (yes/no, left/right, ...). It may however also be continuous, as in the case of a reaction time. Some continuous behavioral variables can easily be discretized : for a reaction times, this can e.g. be done by grouping values that fall within a given quintile and assigning them a label.

Let’s start with a simple scenario. Suppose that you have a fairly good hypothesis as to which oscillation will have an effect on the behavioral variable. For instance, you expect to see an effect of the phase of the alpha rhythm (8-12hz) at the onset of stimulus presentation on the perception of a masked stimulus (Mathewson, Gratton, Fabiani, D. M. Beck, & Ro, 2009). Because the alpha rhythm is such a prominent rhythm in surface EEG, a simple approach will work : band-pass filter your EEG data – between 0.01Hz and 25Hz in (Mathewson et al., 2009), then do stimulus-onset locked averages, grouping trials that have the same behavioral outcome (this is for a discrete behavioral variable). If you can see a prestimulus oscillation with a period of 100ms in the averages of each condition, which you don’t see in the average over all trials, and if the phase of this oscillation differs between the different behavioral outcomes, you have shown – implicitly – that the phase of the ongoing alpha rhythm is informative on the behavioral outcome. As an illustration, we reproduced figure 4f from (Mathewson et al., 2009) (Figure 44). A simple parametric test, e.g. a Student’s t-test, can be performed at each time point to assess the significance of the results in this straightforward approach (appropriately taking multiple comparisons into account when setting the threshold).

It is critical, in the case when prestimulus effects are of interest, to enforce a variable inter-trial interval. Indeed, if the inter-trial interval is kept constant, the brain will learn when to expect the onset of each trial, and this has been shown to affect the distribution of prestimulus phases in such a way that it won’t be uniform any more (Barry et al., 2004) – not all phase values will be represented pre-stimulus due to expectation/attentional effects. It is a key aspect of our approach that the phase of the oscillation of interest should be uniformly distributed when all trials are taken into account.




Figure 44 Grand average ERP at channel Pz in the 100 ms preceding and after the onset of the target. Note the opposite phase for the undetected (dashed) and detected (solid) targets. When collapsed across all targets, no phase-locking is apparent (gray). Figure and caption reproduced from (Mathewson et al., 2009)

Of course, the simple approach that we described – averaging per condition in the time domain, is most likely to yield results if the signal has first been filtered in the right frequency band due to a strong prior hypothesis, or if the rhythm of interest is a very prominent rhythm in the signal. The approach we advocate in other situations where the effects are likely to be more subtle and difficult to see with the naked eye consists in a more explicit consideration of phase values. In a first step, the temporal signal should be decomposed into its different frequency components, and time-resolved phase and amplitude information should be extracted for each frequency band. A few approaches have been applied to get to this result; arguably, the two approaches that give the best time-frequency decomposition in terms of optimizing time-frequency resolution are: i) using an analytical tool, the Hilbert transform, in each frequency band (e.g. (Tass et al., 1998)) or ii) convolving the signal with a complex wavelet in each frequency band (e.g. (J P Lachaux, Rodriguez, Martinerie, & F J Varela, 1999)). Though a good understanding of signal processing is necessary for anyone who wants to perform such analyses, it is beyond the scope of this thesis to give a detailed account of these two methods – for our purposes, it suffices to know that they are pretty much equivalent (Le Van Quyen et al., 2001). The output is a complex number at each time and frequency point that were computed. The angle of the complex number at a given time frequency point (T,F) corresponds to the phase of the oscillation of center frequency F at time T. Our question is still: do the values of the phase of a given oscillation at a given time before stimulus presentation tend to cluster around a mean value in one condition, while clustering around a different mean value in another condition – and be randomly distributed over all trials? A Phase Locking Factor (PLF (C Tallon-Baudry, O Bertrand, Delpuech, & Permier, 1997)) can be computed to quantify the clustering of phase values across trials at each (T,F) : it corresponds to the length of the vector resulting from averaging the values of the phase at (T,F) across trials in the complex domain. In the limit that the phase at (T,F) is always exactly the same for all trials, the PLF would be exactly 1. In the limit that the phase is completely random, the PLF would be close to 0 (never exactly zero due to finite number of trials). An image being worth a thousand words – this saying is especially true for science! – here is an illustration of what the PLF is (Figure 45).

So, we end up with a few (T,F) maps of PLF, each of them calculated for a value of the behavioral outcome (considering that the behavioral variable takes a few possible discrete values). Some (T,F) in the pre-stimulus oscillatory phase may have a high PLF for some of the behavioral outcomes... this is essentially what we are looking for. And this is where the messy business starts – we need to decide whether the PLF is significantly higher than it would be by chance, anywhere in the (T,F) area of interest. There are two main hurdles to overcome statistically :




Figure 45 The Phase Locking Factor (a.k.a Phase Locking Value) reflects how much the distribution of phase angles, e.g. at a given time point across trials, differs from uniformity. On the top left, each blue cross shows the phase value for a given trial, plotted on the unit circle. The phases from the different trials can be averaged (the corresponding unit vectors, of which 3 are drawn, are averaged in 2D), and the length of the average vector corresponds to the PLF. A PLF of 0 means that the phases are randomly distributed while a PLF of 1 means that the phase is always the same from one trial to the next. From left to right and top to bottom, the distributions differ more and more from uniformity and the PLF increases.

- Due to finite sampling, the null hypothesis is not a PLF value of zero. This, and the added complication of dealing with circular variables, leads us to advocate the use of permutation or bootstrapping tests (Vinck, Lima, et al., 2010; Vinck, van Wingerden, Womelsdorf, Fries, & Pennartz, 2010) in the assessment of significance for PLF values. In a permutation test, the idea is to randomly shuffle the behavioral labels of all trials – hence reassigning the EEG data of a given trial to a behavioral outcome which may not be the actual outcome of that trial. The surrogate dataset thus constructed is run through the same analysis as the real dataset (calculation of PLF for each behavioral outcome, average over electrodes and subjects if this step was performed in analyzing the real data), and a map of PLF in the (T,F) area of interest is generated. The shuffling is repeated many times. We end up with a null distribution of the PLF in the (T,F) area of interest for our dataset, which the PLF computed with the real data can be compared to.
- Casting a wide net in terms of the (T,F) area of interest leads to a severe multiple comparison situation, which has to be dealt with appropriately. We have used the False Discovery Rate (FDR) correction (Benjamini & Hochberg, 1995), which controls the rate of falsely rejected null hypotheses, while being less stringent than the conservative Bonferroni correction. It is straightforward to implement : it simply relies on the distribution of p-values across the multiple performed comparisons. It should be noted that there is yet another multiple comparison issue, stemming from the number of channels that are analysed in parallel (32, 64, 128, sometimes up to 256 in EEG studies); our approach has been to perform statistics on the mean of all channels, which allows

us to draw safe conclusions about statistical significance before delving into the topography of the effects, and avoids the ominous pitfall of circular analysis.

Once a significant effect of the phase of an ongoing oscillation has been evidenced with this “brute-force” approach, one can get an idea of the size of the effect by trying to do trial-wise predictions of the behavioral outcome based on the phase of the oscillation, etc... So, let’s apply this approach and show that perception relies on oscillatory brain activity!

B. The rhythm of attentional sampling for visual detection

1. Phase of ongoing oscillations and perception

Oscillations measured in the EEG are a manifestation of synchronous fluctuations of the membrane potentials of large ensembles of neurons. This means that a large proportion of the membrane potentials of the underlying neurons go from a hyperpolarized state to a depolarized state, synchronously. Thus, the probability of firing of the underlying neurons will vary periodically: if a given input pattern arrives at a time when the population is more depolarized, the neurons will be more likely to reach threshold and fire than if the same input pattern arrives in a hyperpolarized state. Electrophysiological studies have evidenced the influence of the phase of ongoing rhythms on neural firing (Bishop, 1932; György Buzsáki & Draguhn, 2004; Fries, Nikolić, & Singer, 2007; Peter Lakatos, George Karmos, Mehta, Ulbert, & C. E. Schroeder, 2008; Peter Lakatos et al., 2005; Rajkai et al., 2008; A. Sirota et al., 2008).

It is by no means a revolutionary idea to hypothesize a link between the phase of ongoing brain oscillations and perceptual processes. Our contribution was simply to show it empirically... Bishop (Bishop, 1932) already suggested in 1932 that the cortex exhibits cyclic changes in excitability. Lindsley (Lindsley, 1952) proposed that these oscillations in excitability might be indexed by the phase of alpha. More recently, evidence of a relationship between alpha phase and the amplitude of event-related potentials (Barry et al., 2004; Haig & Gordon, 1998; Jansen & M. Brandt, 1991; Kruglikov & Schiff, 2003; Trimble & Potts, 1975) has been put forward. Behaviorally, little has been done : the phase of the alpha rhythm has been shown to e.g. influence the speed of manual responses (Callaway & Yeager, 1960; Dustman & E. Beck, 1965; Peter Lakatos et al., 2008) and the perception of simultaneity ((Gho & Francisco J Varela, 1988; F J Varela, Toro, John, & Schwartz, 1981)– see page 118).

The simplest possible test of the “cycles of excitability”, that of linking visual detection probability to the phase of ongoing rhythms, had however seldom been performed. Phase effects of infraslow oscillations (0.01-0.1hz) on perception of somatosensory stimulation at threshold have been evidenced by Monto et al (Monto, S. Palva, Juha Voipio, & J. M. Palva, 2008), using full-band EEG recordings (Vanhatalo et al., 2002); typically, frequencies below 0.5hz are not analysed in surface EEG, but the use of special amplifiers allows to investigate slower oscillations. They report a correlation of task performance with the phase of infraslow fluctuations (ISF), the stimulus detection probability being 55% larger in the rising phase than in the falling phase of the ISF cycle. The amplitude of higher frequency oscillations (1-40hz) in the EEG was also found to correlate with the phase of the ISF. It has been known for some time that changes in general attentional/arousal state are reflected in the power spectrum of the EEG; this was for example studied in an auditory vigilance task (Makeig & T. P. Jung, 1996) : target detection was linked to some 15-20s-scale tonic decreases in theta band power and tonic increases in gamma band power. These tonic changes at rather long time scales are thought to index the activity of thalamo-cortical circuits involved in modulating arousal and sensory processing in sleep (Mircea Steriade, D. A. McCormick, & Sejnowski, 1993). Another line of evidence which relates to these slow intrinsic fluctuations is found in the literature on the influence of prestimulus baseline activity on perception; notably, fMRI studies by Andreas Kleinschmidt and his students have demonstrated how the perception of ambiguous visual and auditory stimuli can be predicted from baseline shifts (Hesselmann, Kell, & Kleinschmidt, 2008; Hesselmann, Kell, Eger, & Kleinschmidt, 2008; Sadaghiani, Hesselmann, &

Kleinschmidt, 2009). The ISF likely reflects fluctuations of arousal, accompanied by tonic shifts in cortical excitability, which affect perception; it is quite distinct from the sampling rhythms that we have been discussing throughout this thesis, which should be linked to higher frequency brain rhythms. The experiment that we performed, and describe in the next section, was a timely contribution to the literature, as evidenced by the almost simultaneous publication of a very similar study (Mathewson et al., 2009).

2. PAPER 4 : Busch, Dubois and VanRullen (2009) – The Journal of Neuroscience.

The Phase of Ongoing EEG Oscillations Predicts Visual Perception

Niko A. Busch,^{1,2} Julien Dubois,^{1,2,3} and Rufin VanRullen^{1,2}

¹Université de Toulouse, Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, France, ²Centre National de la Recherche Scientifique, UMR5549, Faculté de Médecine de Rangueil, 31062 Toulouse, France, and ³California Institute of Technology, Pasadena, California 91125

Oscillations are ubiquitous in electrical recordings of brain activity. While the amplitude of ongoing oscillatory activity is known to correlate with various aspects of perception, the influence of oscillatory phase on perception remains unknown. In particular, since phase varies on a much faster timescale than the more sluggish amplitude fluctuations, phase effects could reveal the fine-grained neural mechanisms underlying perception. We presented brief flashes of light at the individual luminance threshold while EEG was recorded. Although the stimulus on each trial was identical, subjects detected approximately half of the flashes (hits) and entirely missed the other half (misses). Phase distributions across trials were compared between hits and misses. We found that shortly before stimulus onset, each of the two distributions exhibited significant phase concentration, but at different phase angles. This effect was strongest in the theta and alpha frequency bands. In this time–frequency range, oscillatory phase accounted for at least 16% of variability in detection performance and allowed the prediction of performance on the single-trial level. This finding indicates that the visual detection threshold fluctuates over time along with the phase of ongoing EEG activity. The results support the notion that ongoing oscillations shape our perception, possibly by providing a temporal reference frame for neural codes that rely on precise spike timing.

Introduction

Sensory systems are incessantly confronted with a continuous stream of information. Can this information be processed in an equally continuous manner at each moment in time, or does perception fluctuate between favorable and less favorable periods (VanRullen and Koch, 2003)? Ongoing oscillations may contribute to this temporal modulation of information processing, since they affect the local electrical field and the intrinsic excitability of neuronal populations (Buzsáki and Draguhn, 2004; Fries et al., 2007; Sirota et al., 2008). Indeed, tonic shifts in the power of spontaneous brain rhythms in certain frequency bands, particularly the alpha (8–12 Hz) and gamma (>30 Hz) bands, are known to accompany changes of neural response amplitude (Başar et al., 1998), attentional state (Worden et al., 2000; Sauseng et al., 2005; Thut et al., 2006; Fries et al., 2008), and perceptual performance (Ergenoglu et al., 2004; Gonzalez Andino et al., 2005; Babiloni et al., 2006; Thut et al., 2006; Womelsdorf et al., 2006; Hanslmayr et al., 2007; Romei et al., 2008; van Dijk et al., 2008). Yet oscillations are not only characterized by their power but also by their instantaneous phase. Since oscillatory phase at a given frequency reflects cyclic fluctuations of a network's excitability that occur on much shorter timescales than variations in

oscillatory power at the same frequency (Bishop, 1932; Buzsáki and Draguhn, 2004; Lakatos et al., 2005; Klimesch et al., 2007; Lakatos et al., 2008; Montemurro et al., 2008; Rajkai et al., 2008; Sirota et al., 2008), phase effects may provide deeper insight into the fine-grained coding of sensory information. While a relationship between phase of spontaneous EEG oscillations and the amplitude of subsequent event-related potentials (ERPs) (Varela et al., 1981; Jansen and Brandt, 1991; Haig and Gordon, 1998; Makeig et al., 2002; Kruglikov and Schiff, 2003; Barry et al., 2004) or the speed of manual responses (Callaway and Yeager, 1960; Dustman and Beck, 1965; Lakatos et al., 2008) has been demonstrated, far less is known about how phase of spontaneous oscillations affects perception itself. While an influence of phase on perception could be expected based on theoretical considerations (VanRullen and Koch, 2003; Schroeder and Lakatos, 2009), EEG experiments (Monto et al., 2008) and psychophysical studies (Jones et al., 2002) have so far demonstrated these effects only for slow frequencies (i.e., frequencies <2 Hz). To understand the role of brain oscillations in perception, it is critical to assess these effects on a temporal scale that is more compatible with the temporal resolution of our visual experience (VanRullen and Koch, 2003; Holcombe, 2009). Specifically, the phase of ongoing oscillations may represent an indicator of perceptual cycles, such that a stimulus appearing at the optimal phase would be optimally registered and perceived, while at another phase it might be entirely missed.

We tested this hypothesis by investigating the influence of prestimulus oscillations of the human EEG on visual perception. We used a signal detection experiment in which visual stimuli were presented at threshold, so that on average only half of the stimuli were perceived, even though all stimuli were identical. We

Received Jan. 8, 2009; revised April 2, 2009; accepted May 19, 2009.

We thank Leila Reddy and James Macdonald for their useful comments. This work was supported by grants from the Agence Nationale de la Recherche (Project ANR 06JCJC-0154), the Fyssen Foundation, and the European Young Investigator Awards to R.V. and the "European Platform for Life Sciences, Mind Sciences, and the Humanities" grant by the Volkswagen Stiftung to N.A.B.

Correspondence should be addressed to Niko A. Busch, Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, 31062 Toulouse Cédex 9, France. E-mail: busch@cerco.ups-tlse.fr.

DOI:10.1523/JNEUROSCI.0113-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/297869-08\$15.00/0

predicted that this trial-to-trial variability of perception could be systematically linked to the phase of ongoing EEG oscillations at or just preceding stimulus onset. We present a single-trial time–frequency analysis demonstrating that the phase of EEG theta (4–8 Hz) and alpha (8–12 Hz) ongoing oscillations strongly influences whether or not a stimulus is perceived.

Materials and Methods

Subjects. Fourteen participants volunteered after giving written informed consent. One participant was excluded from analysis due to excessive artifacts in the EEG data, which contaminated >50% of the trials. Another subject was excluded due to unstable behavioral performance (hit rate out of the range of $50 \pm 25\%$ on four of six blocks). Twelve subjects remained in the sample (four female; mean age: 27.8 years; eight right handed). All participants had normal or corrected to normal vision. The experimental protocol was approved by the Centre National de la Recherche Scientifique ethical committee.

Stimuli and procedure. Stimuli were presented on a black background on a 160 Hz cathode ray tube monitor. The experiment was written in Matlab using the Psychophysics Toolbox (Brainard, 1997). Subjects performed a demanding visual detection task. Each trial started with the presentation of a central fixation cross and two peripheral markers, one above and one below the position where the target stimulus was to be presented (eccentricity: 7° visual angle to the right of the fixation cross). Subjects were instructed to always maintain central fixation but to pay covert attention to the site indicated by the markers. After a variable delay (range: 1000–2000 ms), a target was presented on 80% of the trials (the remaining 20% were target-absent “catch” trials introduced to estimate the false alarm rate of our participants). The target was a very small point of light (diameter: $7'$ visual angle) presented for 6 ms. The target’s luminance was determined for each subject before the start of the EEG session with a staircase procedure, which found the individual luminance threshold at which 50% of the stimuli were detected. After a delay of 1500 ms following target offset, the fixation cross turned into a question mark. Subjects were instructed to report whether or not they had perceived a target by pressing one of two buttons. No instruction for response speed was given. The button press initiated the start of the next trial. The experiment consisted of six blocks of 250 trials.

EEG acquisition and analysis. A SynAmps amplifier system (Neuroscan) was used to record EEG from 32 electrodes mounted in an elastic cap. The electrode layout was modified from the 10–20 system with an additional row of occipital electrodes and a linked-ears reference. Data were recorded in the frequency range from DC to 300 Hz with a sampling rate of 1000 Hz. Electrode impedances were kept at $<5 \text{ k}\Omega$. Data were downsampled off-line to 500 Hz and epoched from -1500 ms before to 1500 ms after stimulus onset. The horizontal electrooculogram (HEOG) was computed as the difference potential between two lateral frontal channels (F7 – F8). An automatic artifact rejection excluded epochs in which the signal exceeded $\pm 75 \mu\text{V}$, and the remaining data were screened manually for residual artifacts. Event-related potentials (ERPs) were computed as the average across all trials per condition. ERPs were baseline corrected by subtracting the average of the 800 ms prestimulus baseline.

The analysis focused on the comparison of EEG spectral power and phases between hits and misses. Phase and power were computed by means of a continuous wavelet transform of single-trial data for the frequency range from 3 to 100 Hz. For increased visibility, results are only plotted in the frequency range between 3 and 50 Hz, since no significant effects were found beyond 50 Hz. The length of the wavelets increased linearly from 3 cycles at 3 Hz to 8 cycles at 100 Hz. This modified wavelet transform was selected to optimize the trade-off between temporal resolution at lower frequencies and stability at higher frequencies. At each time t and frequency f , the result of the wavelet transform for trial k is a complex number in which A represents the amplitude of the signal and φ its phase:

$$A_{k(t,f)} e^{i\varphi_{k(t,f)}}.$$

The extent of phase concentration across trials is quantified by the inter-trial coherence (ITC; also called phase-locking factor or phase-locking value).

$$\text{ITC}_{(t,f)} = \frac{1}{n} \sum_{n=1}^k e^{i(\varphi_{k(t,f)} - 2\pi)}$$

The ITC measure takes values between 0 and 1. A value of 0 represents absence of synchronization across trials between EEG data and the time-locking events, and a value of 1 indicates perfect synchronization. ITC is computed by normalizing the lengths of the complex vectors (representing amplitude and phase) to 1 for all trials and then computing their complex average. Thus, only the information about the phase of the spectral estimate of each trial is taken into account.

We hypothesized that hits and misses were each associated with a particular phase of spontaneous EEG oscillations just before stimulus onset. Since stimulus onset is randomized and unpredictable, however, the phases of EEG oscillations can be assumed to be randomly distributed across all trials (i.e., hits and misses combined). Thus, while the distribution of phases across the whole set of trials should be random, the phase distribution of hits and misses should exhibit a stronger phase concentration. We quantified the differences in phase distributions for hits and misses by computing a phase bifurcation index (Φ). Φ is thus computed by comparing the ITC of hits and misses against the ITC of all trials:

$$\Phi_{t,f} = (\text{ITC}_{\text{hits}(t,f)} - \text{ITC}_{\text{all}(t,f)}) \times (\text{ITC}_{\text{misses}(t,f)} - \text{ITC}_{\text{all}(t,f)}).$$

When phases are locked to different phase angles for hits and misses, Φ will take a positive value (Fig. 1). The upper bound of the phase bifurcation is 1, indicating perfect phase locking in both conditions ($\text{ITC} = 1$), but at exactly opposite phases (hence, a combined phase locking of $\text{ITC} = 0$). The null hypothesis (random phase distributions for hits and misses) predicts a Φ close to zero. When only one condition exhibits phase locking (e.g., when ERPs are evoked only for hits after stimulus presentation), Φ takes negative values (see Fig. 1 for illustrations of these scenarios).

We computed Φ for each point in the time–frequency plane from -800 to 800 ms and from 3 to 100 Hz and statistically evaluated its magnitude with a resampling test. In a first step, data from all trials, regardless of condition, were pooled. Two sets of trials (corresponding to “hits” and “misses”) were then drawn randomly from this pool, and the phase bifurcation index was computed. This procedure was repeated 500 times per subject, thus producing for each subject a distribution of Φ s based on shuffled data under the null hypothesis. In a second step, one of these pseudo- Φ s was drawn at random from the null distribution of each subject, and their grand average was computed. This procedure was repeated 100,000 times. For each point in the time–frequency plane, a p value was computed as the proportion of these pseudo-grand averages that exceeded the observed grand average. This p value thus indicates at which time and frequency the observed phase distributions for hits and misses are more divergent than what is expected for random data (see Fig. 3). The same procedure was used accordingly to evaluate the power differences between hits and misses. To correct for multiple comparisons, we analyzed the resulting distributions of p values with the false discovery rate (FDR) procedure (Benjamini and Hochberg, 1995) to compute a p threshold that set the expected rate of falsely rejected null hypotheses to 5%.

Results

On average, subjects detected half of the targets (mean hit rate 53.7%, SE 4.1), while the false alarm rate on target-absent trials (20% of all trials) was very low (mean: 3.0%, SE 0.8).

ERP amplitudes differed largely between perceptual conditions (see supplemental Fig. S1, available at www.jneurosci.org as supplemental material). For hits, two main ERP components were found: a negative component with a maximum at 240 ms at parietal electrodes contralateral to stimulation and a positive component peaking at 450 ms with a central topography. These

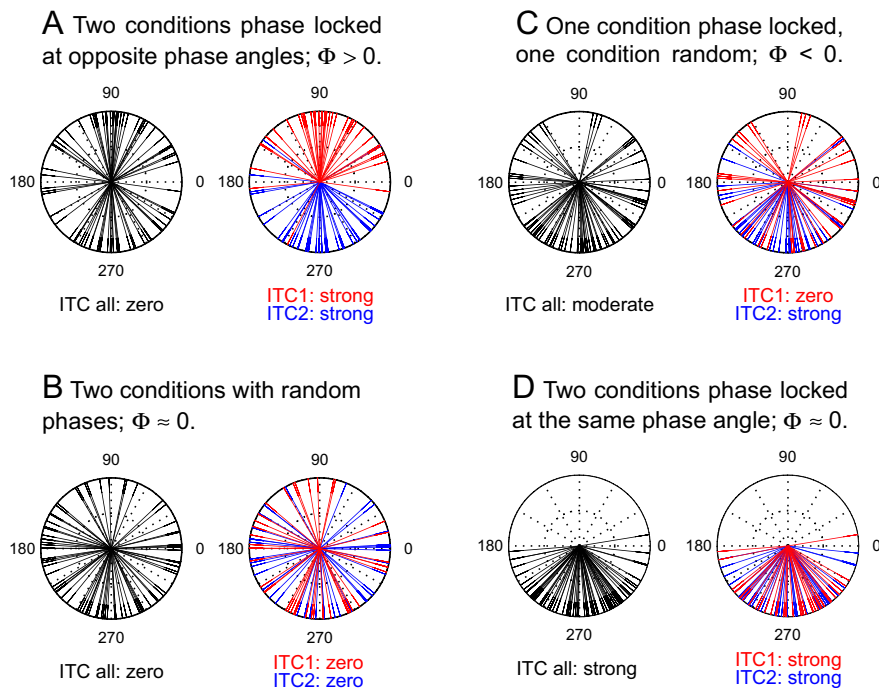


Figure 1. Illustration of the phase bifurcation index with hypothetical data. The index is computed as $\Phi = (ITC_1 - ITC_{all}) \times (ITC_2 - ITC_{all})$, where $ITC_{1,2,all}$ is the intertrial coherence in conditions 1 and 2 and in both conditions combined. Each line segment represents a trial, with its angle representing phase at the frequency and time point of interest. The left circles illustrate global phase distributions, while the right circles show the two experimental conditions separately. **A**, Both conditions are phase locked at opposite angles. As a result, the ITC of both conditions combined is close to zero, resulting in a strong positive Φ . This is equivalent to the hypothesis that hits and misses are each associated with different phase angles. **B**, The null hypothesis states that both conditions are randomly distributed, resulting in a Φ close to zero. **C**, Only one condition is phase locked, and the other is randomly distributed. The resulting ITC of both conditions is stronger than the ITC of condition 1, but smaller than the ITC of condition 2, resulting in a negative Φ . This situation is expected during the ERP time range, where an ERP is present only for detected stimuli. **D**, Both conditions are strongly phase locked at the same phase angle; the resulting Φ is close to zero.

ERPs were virtually absent when the target was not perceived. The differences were confirmed by a paired *t* test computed on the mean amplitude from 230 to 250 ms at channel T5 ($t_{(11)} = -3.07$; $p = 0.01$) and on the mean amplitude from 430 to 470 ms at channel Cz ($t_{(11)} = 9.71$; $p < 0.001$).

We investigated whether ongoing EEG activity before stimulus presentation influences detection by analyzing spectral power and phase in the prestimulus time window. In particular, we (1) analyzed power differences between hit and miss trials and (2) computed the phase bifurcation index (see Materials and Methods and Fig. 1); significantly positive phase bifurcation indicates that oscillations are locked to different phase angles on hit and miss trials. While a power difference in the prestimulus period is expected based on previous studies, an effect of prestimulus phase on subsequent perception would constitute a major conceptual advance.

We started out by computing spectra of power effects and phase bifurcation as averages across electrodes and time points in the 800 ms prestimulus interval (Fig. 2). These pooled measures allowed us to determine frequency bands of interest, on which more fine-grained analyses were then performed. We found that spectral power was significantly stronger for misses than for hits in the 6–12 Hz frequency range (Fig. 2A, left inset). In this frequency range, the effect was strongest in the –600 to –300 ms time window preceding stimulus onset (Fig. 2A, bottom inset). In this time–frequency window, a widely distributed topography was observed with a maximum at the frontocentral channel Fz. Subsequent power analyses were thus focused on this channel. We found that spectral power was significantly stronger preced-

ing missed targets than detected targets (Fig. 3A), and this effect reached a maximum at 8.2 Hz and –492 ms preceding stimulus onset. To statistically confirm these effects, we used a resampling test and applied the FDR method (Benjamini and Hochberg, 1995) to correct for the effect of multiple comparisons across time points and frequencies. Effects that satisfied a 5% FDR criterion are outlined in white in Figure 3A. The impact of spectral power on detection performance was assessed by binning single trials according to the spectral power at that time–frequency point and computing the average detection rate of our observers in each of 10 bins. To minimize the impact of individual variations in performance, detection rate was standardized for each observer by dividing the hit rate in each bin by the average hit rate of that observer. Standardized hit rates were highest in the bin corresponding to lowest spectral power, and lowest in the bin with the highest power. The difference between these two bins amounted to a 12% change in performance, and was confirmed statistically with a repeated measurements ANOVA with “bin” as a factor (10 levels): $F_{(9,99)} = 2.404$; $p = 0.0165$.

A similar analysis was performed for the phase bifurcation index Φ —a measure of the divergence of two phase distributions (Fig. 1). The phase bifurcation spectrum (averaged across channels and time points in the prestimulus interval) was strongest in a frequency range from 6 to 10 Hz. Here, effects exceeded the 95% confidence limit that was based on 100,000 synthetic datasets with random phase distributions (Fig. 2B, left inset). Strongest phase bifurcation was observed in the –300 to –50 ms time window preceding stimulus onset (Fig. 2B, bottom inset), indicating significant phase concentration for hits and misses, but at different phase angles. We ascertained that phases were uniformly distributed across all trials (i.e., when hits and misses were combined) for this time–frequency point with a Rayleigh test (cf. Fisher, 1995). When *p* values were Bonferroni-corrected for multiple comparisons, none of the subjects had a significant deviation from uniformity, but without correction for multiple comparisons the deviation was statistically significant for 1 out of the 12 subjects. To assess the significance of this finding, we performed a randomization test by computing a distribution of random phases for each subject with the same number of “trials” as in the actual data, and the Rayleigh test was computed for these random phases. This procedure was repeated 100,000 times for the group of 12 subjects. The probability of observing at least 1 out of 12 significant deviations from uniformity was 0.46, even though phases were now drawn from a random distribution. Thus, observing only one significant deviation from uniformity in our dataset is highly expected even for randomly distributed phases. We thus conclude that it is safe to assume that phases were indeed uniformly distributed in our data.

The phase bifurcation effect in this time–frequency range had a frontocentral topography with a maximum at channel Fz. We

evaluated the statistical significance of this effect at channel Fz with a resampling test and found the strongest effect at 7.1 Hz and 120 ms preceding stimulus onset (Fig. 3B). We used the FDR method to correct for multiple comparisons; white outlines in the inset of Figure 3B indicate the 55 time–frequency points that passed the 5% FDR criterion (meaning that only 3 points out of these 55 are expected to be false positives due to multiple comparisons). To verify that these results were not due to a single outlier in our subject pool, significance maps were recomputed 12 times, each time leaving out a different subject; the region of maximum significance (using an FDR of 0.1) always counted at least nine time–frequency points, every time including the point of interest at 7.1 Hz and –120 ms preceding stimulus onset.

Phases at this time–frequency point were pooled into 11 bins. The effect of phase bifurcation was significant across subjects, and the mean phase values for hits and for misses were approximately similar across subjects (Fig. 3B, top inset). However, the exact phase at which performance was highest could vary slightly between subjects. We thus adjusted each subject's phase distributions such that for each subject the phase at which performance was best was aligned to a phase angle of zero. Thus, a trivial feature of the resulting distribution is a peak exceeding the standardized average hit rate at a phase of zero. However, we also found that performance decreased monotonically to a minimum at the opposite phase—a nontrivial property confirming our hypothesis that hits and misses are associated with opposite phase angles. Comparing the average of the two bins adjacent to the zero bin with the average of the two bins at the opposite phase revealed that phase accounted for a difference of 16% of performance (standardized performance: 1.05 vs 0.89), i.e., even stronger than the corresponding effect observed for power (Fig. 3A). This effect was confirmed by an ANOVA with a 10-level “bin” factor (leaving out the bin centered on zero): $F_{(9,99)} = 2.696$; $p = 0.0076$.

The influence of EEG phase on performance was independently assessed using a linear classifier approach: a support vector machine classifier was trained to discriminate hits versus miss trials based on instantaneous phase at electrode Fz. Classifier output on a distinct set of test trials was maximally correlated with observers' performance in the time–frequency window around 7 Hz, –120 ms (supplemental Fig. S2, available at www.jneurosci.org as supplemental material).

These analyses relied on a time–frequency transform that convolves the signal with a wavelet function, which is extended in time and the spread of which is inversely related to the frequency being analyzed. Thus, it is conceivable that some of the reported




Figure 2. Raw effects of oscillatory power and phase. **A**, Difference in spectral power in decibels between hits and misses, averaged across channels and subjects. Negative values indicate stronger power for misses. Left inset shows power difference averaged across time points in the prestimulus window (shaded areas: SEM). The main difference is found in the 6–12 Hz frequency range. Bottom inset, Power difference averaged across frequencies in this range, with a maximally negative difference in the prestimulus time range between –600 and –300 ms (gray shaded area). The topography shows the distribution of the power difference from 6 to 12 Hz and from –600 to –300 ms preceding stimulus onset. **B**, Phase bifurcation index (Φ), averaged across all channels and subjects. Positive values indicate that phase distributions are locked to different phase angles for hits and misses (e.g., in the prestimulus time range), while negative values indicate that only one condition is phase locked (e.g., phase locking exclusively for hits in the ERP time range). Left inset shows Φ averaged across all time points in the prestimulus window (vertical lines represent the 95% confidence interval). Bottom inset, Φ averaged across frequencies between 6 and 10 Hz (shaded area: SE). Phase bifurcation is strongest from –300 to –50 ms preceding stimulus onset (gray shaded area). The topography shows the distribution of Φ from 6 to 10 Hz and from –300 to –50 ms preceding stimulus onset.

prestimulus effects were affected by EEG data collected after stimulus onset. However, we repeated the analysis with much shorter wavelet functions (limited to 1 cycle instead of 3 or more cycles) and confirmed that the prestimulus effects of phase and power precede the time range that can potentially be affected by poststimulus data (compare supplemental Fig. S3, available at www.jneurosci.org as supplemental material).

We also ascertained that power and phase effects were not due to any eye movement artifacts before stimulus onset that survived the artifact rejection procedure. For example, differences in power and phase on hit and miss trials could have simply been caused by differences in the number or the timing of eye blinks before stimulus onset. Eye blinks are associated with high amplitude signals at frontal channels. However, we found no evidence

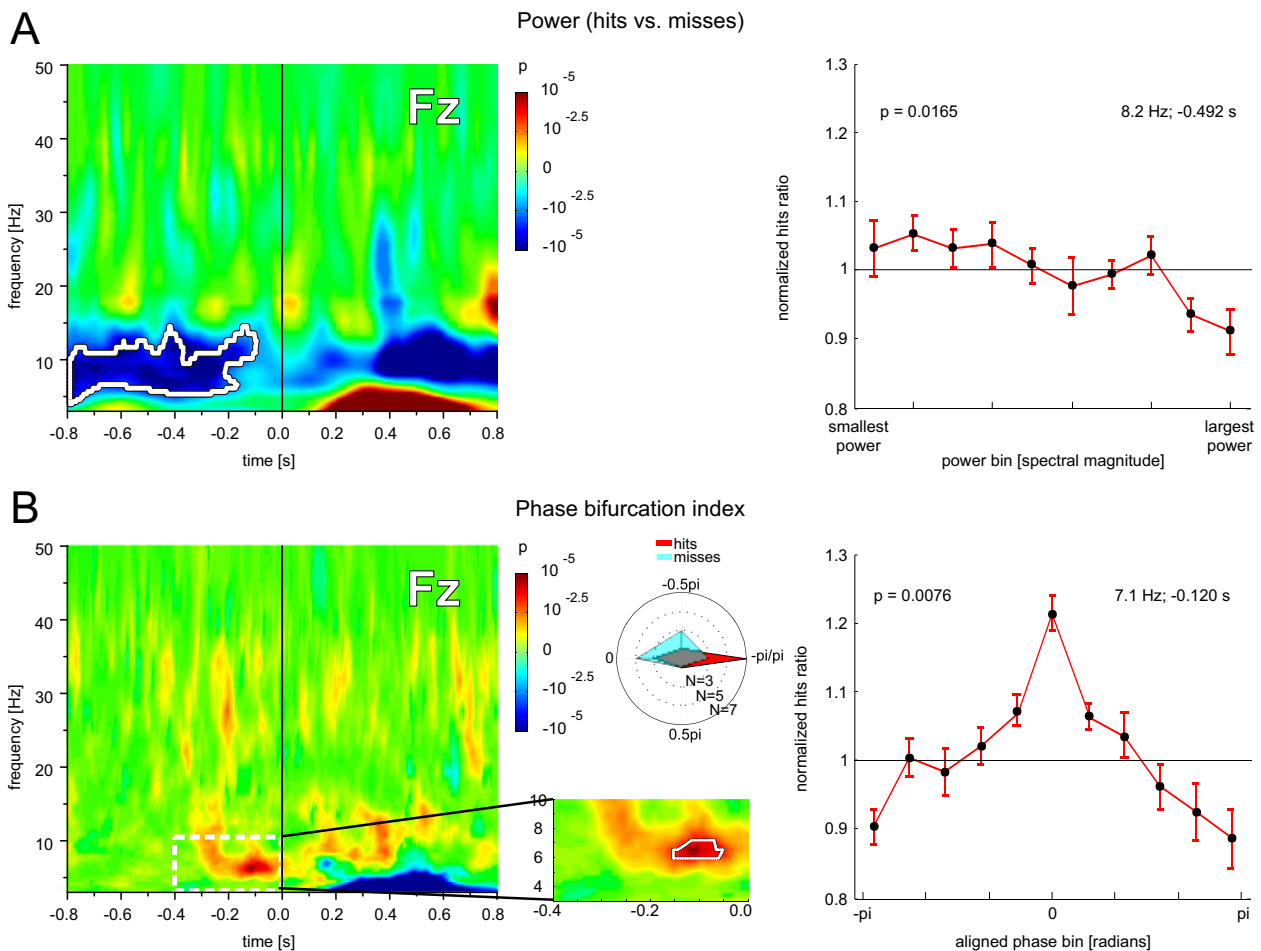


Figure 3. Relation between perception and oscillatory activity. **A**, Statistical significance of the power difference at channel Fz between hits and misses (left panel), evaluated with a resampling test. Alpha power (6–12 Hz) preceding stimulus onset was significantly stronger for misses than for hits. The color map represents uncorrected *p* values and the white outline delimits significant effects corresponding to an FDR of 5%. The right panel shows the relationship between spectral power in the alpha band (8.2 Hz) at –492 ms and performance after single trials were pooled in 10 power bins. The horizontal line indicates average performance (standardized to 1) across all bins. Error bars represent SEM. Performance is superior on trials with lowest alpha power (1-way ANOVA, $p < 0.0165$). **B**, Statistical significance of phase bifurcation between hits and misses at channel Fz (left panel); axes and conventions are as in **A**. At ~7 Hz and –120 ms prestimulus, hits and misses are associated with different phase angles. Top inset, The circular histograms of mean phase angles at 7.1 Hz and 120 ms for hits and misses across participants. Phases were pooled into four phase bins. The distance from the origin indicates the number of subjects falling within a bin, and the angles indicate the lower bound of each phase bin. Bottom inset, Close-up on the 55 time–frequency points with significant phase bifurcation that satisfy an FDR of 5%. The right panel shows the relationship between phase (at 7.1 Hz; –120 ms) and standardized performance after phases were aligned for each subject so that the optimal phase corresponds to a zero phase angle. Performance declines to a minimum at the opposite phase angle (1-way ANOVA, $p < 0.01$).

for eye blinks on single trials or in the average ERP, either at Fz (where power and phase bifurcation effects were strongest) or in the HEOG (compare supplemental Fig. S4, available at www.jneurosci.org as supplemental material), for hits or misses. Alternatively, hits and misses might have been associated with horizontal saccades toward or away from the target, respectively. If so, these saccades might have introduced an electrical artifact with opposite polarity (i.e., opposite phase angles) for hits and misses. However, inspection of the HEOG on single trials and in the averaged ERP did not reveal any signs of systematic prestimulus saccades, least of all in opposite directions (compare supplemental Fig. S4, available at www.jneurosci.org as supplemental material). In sum, these results indicate that power and phase bifurcation effects in the main analysis were not brought about by ocular artifacts.

Discussion

The traditional approach in cognitive neuroscience has been to investigate neural activity in response to (that is, succeeding) an experimental event. With this reasoning, it is usually assumed that the brain state preceding the event does not play a meaning-

ful role in how the event will be processed, and that trial-to-trial variations in the response to identical stimuli reflect random noise. In recent years however, there has been a growing interest in the role of prestimulus brain states in many different domains of research. Examples include prestimulus shifts of firing rates and synchronization (Fries et al., 2001; Womelsdorf et al., 2006), anticipatory baseline changes of blood oxygenation level-dependent activity in functional magnetic resonance imaging (McMains et al., 2007; Silver et al., 2007; Sestieri et al., 2008), spatiotemporal activity patterns (Arieli et al., 1996), and spontaneous EEG activity (Başar et al., 1998; Ergenoglu et al., 2004; Lakatos et al., 2005; Thut et al., 2006; Hanslmayr et al., 2007; Monto et al., 2008; van Dijk et al., 2008; Schroeder and Lakatos, 2009). We investigated the influence of prestimulus oscillations on visual perception. We presented visual stimuli near threshold, and found that their detection probability was strongly dependent on the phase of spontaneous EEG oscillations in the low alpha and theta bands just before stimulus onset. This result suggests that perception may operate in successive periodic cycles,

alternating between phases of optimal excitability where threshold stimuli are consciously perceived, and phases associated with stronger inhibition at which the same stimuli are more likely to escape detection.

Our analysis also confirmed previous studies reporting that stimuli preceded by strong alpha power are less likely to be detected (Ergenoglu et al., 2004; Babiloni et al., 2006; Thut et al., 2006; Hanslmayr et al., 2007). However, the influence of phase on visual detection performance, accounting for a minimum of 16% variability in performance (Fig. 3*B*), was in fact stronger than the previously characterized effect of alpha power, which in our study accounted for 12% of performance variability. While these numbers may appear small, it is important to remember that our estimates rely on single-trial EEG data, which in general are deemed to be inherently too noisy to be analyzed (Picton et al., 2000). The measured effects of phase and power were largely independent, as indicated by the different topographies and time courses: while alpha power effects prevailed throughout the prestimulus time window and were strongest ~500 ms before stimulus onset, the effect of phase on perception was restricted to the interval immediately preceding stimulus onset. In line with this distinction, it should be noted that the effects of prestimulus power (Ergenoglu et al., 2004) or interelectrode synchrony (Hanslmayr et al., 2007) on performance are related to the subject's tonic state of attention or arousal, which usually fluctuates on a much longer timescale than the rhythmic modulations of excitability associated with oscillatory phase (Makeig and Jung, 1996; Klimesch et al., 2007). Thus, while spontaneous power of neural oscillations undoubtedly influences visual perception, the effects of prestimulus phase may be more closely linked to the actual coding and processing of visual information.

The existence of a relationship between spontaneous EEG alpha phase and the amplitude of subsequent ERPs (Varela et al., 1981; Jansen and Brandt, 1991; Haig and Gordon, 1998; Makeig et al., 2002; Kruglikov and Schiff, 2003; Barry et al., 2004) or the speed of manual responses (Callaway and Yeager, 1960; Dustman and Beck, 1965) has been postulated for decades. However, this is rarely taken into account in theories of visual perception, perhaps because it could be attributed to a mere facilitation at the biophysical level, without direct relevance to conscious perception. In contrast, our study directly demonstrates that prestimulus phase of ongoing neural oscillations can affect visual perception itself. Note that in our study the major effect was found at ~7 Hz, i.e., at the intersection between theta and alpha frequency bands, whereas most of the previous studies mentioned above reported their effects in the alpha band—the frequency range from 8 to 12 Hz. This apparent divergence may be explained by the wide range of frequencies that were analyzed without restriction in our study, while previous studies tended to limit their analysis to a narrow frequency range by application of a preselected bandpass filter (usually centered around 10 Hz) (Dustman and Beck, 1965; Varela et al., 1981; Jansen and Brandt, 1991; Haig and Gordon, 1998). In addition, the absence of significant phase effects in the gamma band (>30 Hz) is worth mentioning, since such effects might have been expected on theoretical grounds (Fries et al., 2007). However, this absence may not rule out a role for gamma oscillations, but could instead reflect an intrinsic limitation of surface EEG recordings: small shifts of conduction delays, synaptic transmission delays, and other biophysical parameters on the order of a few milliseconds exert a maximally disruptive influence on the phase measured at higher frequencies. For example, a mere 8 ms shift actually corresponds to a full phase reversal for an oscillation at 60 Hz. Moreover, the signal-to-noise ratio in sur-

face EEG recordings decreases at higher frequencies. In addition, amplitude at high frequencies is often coupled to the phase of lower frequencies (Lakatos et al., 2005), making gamma activity detectable only in certain time windows. Therefore, the gamma phase recorded at the scalp surface may not always faithfully reflect the relevant neuronal oscillation. Similarly, we could not assess effects of slow frequencies (i.e., slower than 3 Hz), which would have required much longer intertrial intervals. It has been demonstrated recently that detection of somatosensory stimuli is strongly correlated with the phase of frequencies below 1 Hz (Monto et al., 2008), and future studies should address this phenomenon in the visual domain. It may be speculated that such effects are strongest in experiments with inherent temporal structure and predictable stimulus onsets, which would allow a sensory system to use a low frequency oscillation to prepare the excitability state of the system for the approximate time of stimulus arrival [for similar ideas, see Jones et al. (2002), Lakatos et al. (2008), and Schroeder and Lakatos (2009)].

What could be the source of the effect of EEG phase on visual perception reported here? Given its frontocentral topography and its frequency range in the theta and low alpha band, one might conjecture a source within frontal midline structures. These areas have previously been discussed as generators of EEG theta rhythms (Debener et al., 2005; Onton et al., 2005; Tsujimoto et al., 2006), which are involved in a variety of cognitive tasks (for review, see Mitchell et al., 2008) requiring sustained attention (Sauseng et al., 2007) or working memory (Gevins et al., 1997; Jensen and Tesche, 2002), and are correlated with manual reaction times (Gonzalez Andino et al., 2005). Therefore, it may not be surprising that frontal theta rhythms could also participate in the dynamic shaping of perceptual experience. However, it should be emphasized that the topography in Figure 2*B* corresponds to a map of the effects' statistical significance, without providing information about the polarity of the underlying electrical signals. Such polarity information is essential for localizing neuroelectric sources. Moreover, the present data were recorded from only 32 channels, while source localization requires a substantially larger number of channels. Thus, any conjecture as to the cortical generators of the phase effects has to remain speculative at present. Future studies should apply similar phase analyses in source space using denser electrode montages.

To conclude, the finding that threshold stimuli are sometimes perceived and sometimes missed depending on the phase of spontaneous brain oscillations supports the assumption that these oscillations reflect cyclic variations of neural excitability (Bishop, 1932; Harter, 1967; Buzsáki and Draguhn, 2004; VanRullen et al., 2005; Fries et al., 2007; Rajkai et al., 2008; Sirota et al., 2008). Such excitability cycles may instantiate a discrete processing mode, framing visual perception into discrete epochs or "perceptual moments" (Stroud, 1967; Varela et al., 1981; Purves et al., 1996; VanRullen and Koch, 2003; Smith et al., 2006). The advantage of a discrete processing scheme could be to transform stimulus information into a temporal code. Indeed, in a network whose excitability is modulated by spontaneous oscillations, the relative strength of input stimulus features is converted into a phase code: the stronger the feature value, the earlier the corresponding units will respond within a given oscillation cycle (O'Keefe and Recce, 1993; Hopfield, 1995; Mehta et al., 2002; Margrie and Schaefer, 2003; Buzsáki and Draguhn, 2004; Lisman, 2005; Fries et al., 2007). The hippocampus heavily relies on this principle (coined "theta phase precession") to encode spatial location as the animal navigates through the spatially restricted firing field of so-called "place cells" (O'Keefe and Recce, 1993;

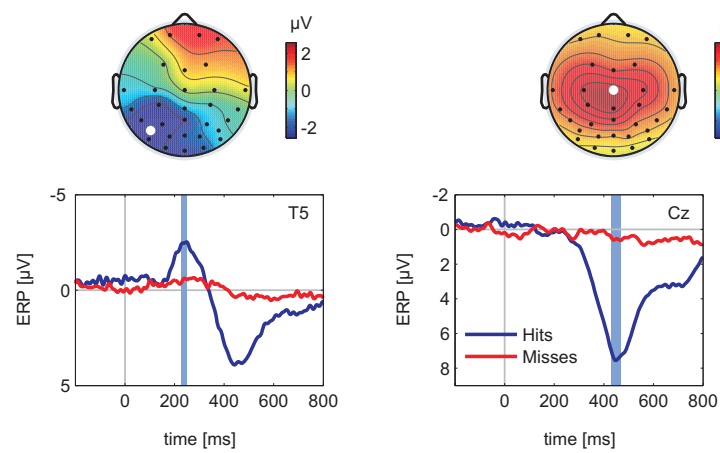
Dragoi and Buzsáki, 2006). Similarly, visual attributes of a stimulus could be represented in each oscillation cycle by the relative phases of firing across the neuronal population (Thorpe et al., 2001; VanRullen et al., 2005; Fries et al., 2007). While further electrophysiological evidence will be needed to support this speculative proposal, the present demonstration of periodic fluctuations of visual perceptual abilities could constitute an important first step in this direction.

References

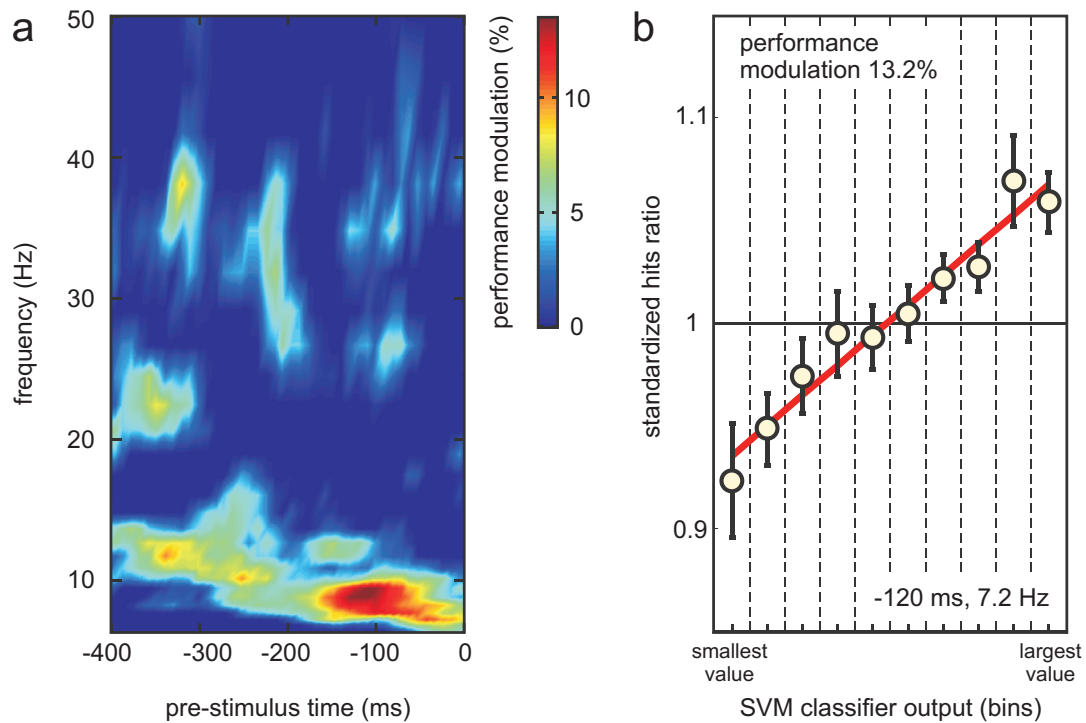
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871.
- Babiloni C, Brancucci A, Del Percio C, Capotosto P, Arendt-Nielsen L, Chen ACN, Rossini PM (2006) Anticipatory electroencephalography alpha rhythm predicts subjective perception of pain intensity. *J Pain* 7:709–717.
- Barry RJ, Rushby JA, Johnstone SJ, Clarke AR, Croft RJ, Lawrence CA (2004) Event-related potentials in the auditory oddball as a function of EEG alpha phase at stimulus onset. *Clin Neurophysiol* 115:2593–2601.
- Başar E, Rahm E, Demiralp T, Schürmann M (1998) Spontaneous EEG theta activity controls frontal visual evoked potential amplitudes. *Electroencephalogr Clin Neurophysiol* 108:101–109.
- Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B* 57:289–300.
- Bishop GH (1932) Cyclic changes in excitability of the optic pathway of the rabbit. *Am J Physiol* 103:213–224.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Callaway E 3rd, Yeager CL (1960) Relationship between reaction time and electroencephalographic alpha phase. *Science* 132:1765–1766.
- Debener S, Ullsperger M, Siegel M, Fiehler K, von Cramon DY, Engel AK (2005) Trial-by-trial coupling of concurrent electroencephalogram and functional magnetic resonance imaging identifies the dynamics of performance monitoring. *J Neurosci* 25:11730–11737.
- Dragoi G, Buzsáki G (2006) Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* 50:145–157.
- Dustman RE, Beck EC (1965) Phase of alpha brain waves, reaction time and visually evoked potentials. *Electroencephalogr Clin Neurophysiol* 18:433–440.
- Ergenoglu T, Demiralp T, Bayraktaroglu Z, Ergen M, Beydagi H, Uresin Y (2004) Alpha rhythm of the EEG modulates visual detection performance in humans. *Brain Res Cogn Brain Res* 20:376–383.
- Fisher NI (1995) *Statistical analysis of circular data*. Cambridge, UK: Cambridge UP.
- Fries P, Neuenschwander S, Engel AK, Goebel R, Singer W (2001) Rapid feature selective neuronal synchronization through correlated latency shifting. *Nat Neurosci* 4:194–200.
- Fries P, Nikolić D, Singer W (2007) The gamma cycle. *Trends Neurosci* 30:309–316.
- Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* 28:4823–4835.
- Gevins A, Smith ME, McEvoy L, Yu D (1997) High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. *Cereb Cortex* 7:374–385.
- Gonzalez Andino SL, Michel CM, Thut G, Landis T, Grave de Peralta R (2005) Prediction of response speed by anticipatory high-frequency (gamma band) oscillations in the human brain. *Hum Brain Mapp* 24:50–58.
- Haig AR, Gordon E (1998) EEG alpha phase at stimulus onset significantly affects the amplitude of the P3 ERP component. *Int J Neurosci* 93:101–115.
- Hanslmayr S, Aslan A, Staudigl T, Klimesch W, Herrmann CS, Bäuml KH (2007) Prestimulus oscillations predict visual perception performance between and within subjects. *Neuroimage* 37:1465–1473.
- Harter MR (1967) Excitability cycles and cortical scanning: a review of two hypotheses of central intermittency in perception. *Psychol Bull* 68:47–58.
- Holcombe AO (2009) Seeing slow and seeing fast: two limits on perception. *Trends Cogn Sci* 13:216–221.
- Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36.
- Jansen BH, Brandt ME (1991) The effect of the phase of prestimulus alpha activity on the averaged visual evoked response. *Electroencephalogr Clin Neurophysiol* 80:241–250.
- Jensen O, Tesche CD (2002) Frontal theta activity in humans increases with memory load in a working memory task. *Eur J Neurosci* 15:1395–1399.
- Jones MR, Moynihan H, MacKenzie N, Puente J (2002) Temporal aspects of stimulus-driven attending in dynamic arrays. *Psychol Sci* 13:313–319.
- Klimesch W, Sauseng P, Hanslmayr S (2007) EEG alpha oscillations: the inhibition-timing hypothesis. *Brain Res Rev* 53:63–88.
- Kruglikov SY, Schiff SJ (2003) Interplay of electroencephalogram phase and auditory-evoked neural activity. *J Neurosci* 23:10122–10127.
- Lakatos P, Shah AS, Knuth KH, Ulbert I, Karmos G, Schroeder CE (2005) An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *J Neurophysiol* 94:1904–1911.
- Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320:110–113.
- Lisman J (2005) The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus* 15:913–922.
- Makeig S, Jung TP (1996) Tonic, phasic, and transient EEG correlates of auditory awareness in drowsiness. *Brain Res Cogn Brain Res* 4:15–25.
- Makeig S, Westerfield M, Jung TP, Enghoff S, Townsend J, Courchesne E, Sejnowski TJ (2002) Dynamic brain sources of visual evoked responses. *Science* 295:690–694.
- Margrie TW, Schaefer AT (2003) Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol* 546:363–374.
- McMains SA, Fehd HM, Emmanouil TA, Kastner S (2007) Mechanisms of feature- and space-based attention: response modulation and baseline increases. *J Neurophysiol* 98:2110–2121.
- Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–746.
- Mitchell DJ, McNaughton N, Flanagan D, Kirk IJ (2008) A review of frontal-midline theta from the perspective of hippocampal “theta”. *Prog Neurobiol* 86:156–185.
- Montemurro MA, Rasch MJ, Murayama Y, Logothetis NK, Panzeri S (2008) Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr Biol* 18:375–380.
- Monto S, Palva S, Voipio J, Palva JM (2008) Very slow EEG fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. *J Neurosci* 28:8268–8272.
- O’Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- Onton J, Delorme A, Makeig S (2005) Frontal midline EEG dynamics during working memory. *Neuroimage* 27:341–356.
- Picton TW, Bentin S, Berg P, Donchin E, Hillyard SA, Johnson R Jr, Miller GA, Ritter W, Ruchkin DS, Rugg MD, Taylor MJ (2000) Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology* 37:127–152.
- Purves D, Paydarfar JA, Andrews TJ (1996) The wagon wheel illusion in movies and reality. *Proc Natl Acad Sci U S A* 93:3693–3697.
- Rajkai C, Lakatos P, Chen CM, Pincze Z, Karmos G, Schroeder CE (2008) Transient cortical excitation at the onset of visual fixation. *Cereb Cortex* 18:200–209.
- Romei V, Brodbeck V, Michel C, Amedi A, Pascual-Leone A, Thut G (2008) Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. *Cereb Cortex* 18:2010–2018.
- Sauseng P, Klimesch W, Stadler W, Schabus M, Doppelmayr M, Hanslmayr S, Gruber WR, Birbaumer N (2005) A shift of visual spatial attention is selectively associated with human EEG alpha activity. *Eur J Neurosci* 22:2917–2926.
- Sauseng P, Hoppe J, Klimesch W, Gerloff C, Hummel FC (2007) Dissociation of sustained attention from central executive functions: local activity and interregional connectivity in the theta range. *Eur J Neurosci* 25:587–593.
- Schroeder CE, Lakatos P (2009) Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci* 32:9–18.
- Sestieri C, Sylvester CM, Jack AI, d’Avossa G, Shulman GL, Corbetta M

- (2008) Independence of anticipatory signals for spatial attention from number of nontarget stimuli in the visual field. *J Neurophysiol* 100:829–838.
- Silver MA, Ress D, Heeger DJ (2007) Neural correlates of sustained spatial attention in human early visual cortex. *J Neurophysiol* 97:229–237.
- Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* 60:683–697.
- Smith ML, Gosselin F, Schyns PG (2006) Perceptual moments of conscious visual experience inferred from oscillatory brain activity. *Proc Natl Acad Sci U S A* 103:5626–5631.
- Stroud JN (1967) The fine structure of psychological time. *Ann N Y Acad Sci* 138:623–631.
- Thorpe S, Delorme A, Van Rullen R (2001) Spike-based strategies for rapid processing. *Neural Netw* 14:715–725.
- Thut G, Nietzel A, Brandt SA, Pascual-Leone A (2006) Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J Neurosci* 26:9494–9502.
- Tsujimoto T, Shimazu H, Isomura Y (2006) Direct recording of theta oscillations in primate prefrontal and anterior cingulate cortices. *J Neurophysiol* 95:2987–3000.
- van Dijk H, Schoffelen JM, Oostenveld R, Jensen O (2008) Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *J Neurosci* 28:1816–1823.
- VanRullen R, Koch C (2003) Is perception discrete or continuous? *Trends Cogn Sci* 7:207–213.
- VanRullen R, Guyonneau R, Thorpe SJ (2005) Spike times make sense. *Trends Neurosci* 28:1–4.
- Varela FJ, Toro A, John ER, Schwartz EL (1981) Perceptual framing and cortical alpha rhythm. *Neuropsychologia* 19:675–686.
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736.
- Worden MS, Foxe JJ, Wang N, Simpson GV (2000) Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *J Neurosci* 20:RC63.

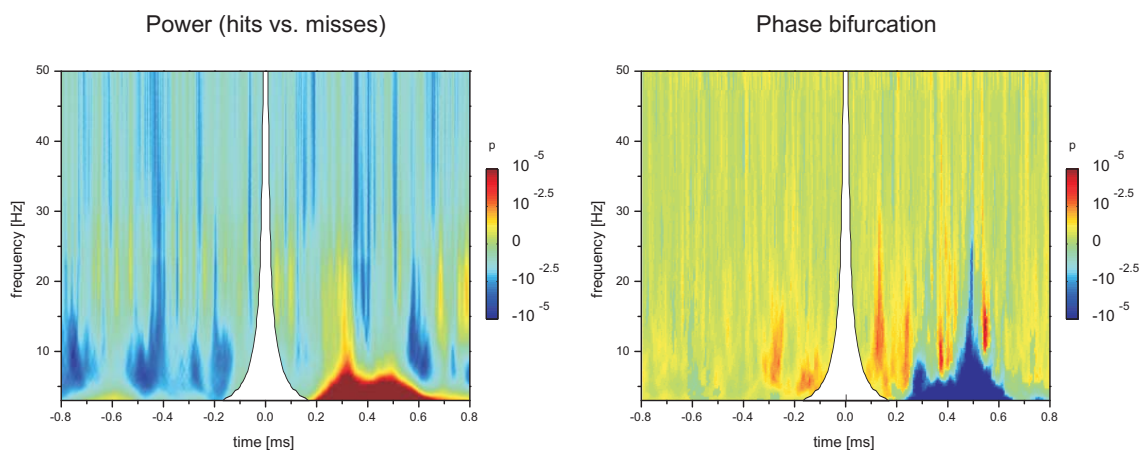
Supplementary information




Supplementary Figure 1: Event-related potentials (ERPs; grand average over 12 subjects). Left: ERP time course at electrode T5 and topography in the time range shaded in blue (230 – 250 ms). Right: ERP time course at electrode Cz and topography in the time range 430 – 470 ms.



Supplementary Figure 2: For each subject, we repeated the following procedure 20 times. At each time-frequency point within the range of 4 – 50 Hz and -400 – 0 ms pre-stimulus, we trained a linear Support Vector Machine classifier to discriminate hits and misses based on the corresponding phase at electrode Fz (represented as a unit-length vector in the 2D complex plane). Label -1 was assigned to miss trials and label +1 to hit trials. Training was performed on a randomly selected subset of 80% of trials and repeated 5 times, each time leaving out another 20% of trials, resulting in a full cross-validation procedure. The classifier was then tested on the remaining 20% of trials of the dataset: for each trial, it would return a decision value (a number usually between -1 and 1) and would use the sign of that decision value to predict a label for the test trial. If phase information does allow the classifier to reliably generalize from the training set to the test set, then the subject's hit rate should be correlated with these decision values. We thus organized decision values into 10 bins, and for each bin, computed the standardized hit rate of the subject over the corresponding trials. We quantified the influence of phase information on detection performance as the slope of the best-fitting line linking the binned classifier output decision value to the standardized hit rate. This slope was averaged over the 5 repetitions of the 5-fold cross-validation and over the 20 repetitions of the entire procedure, and the result finally averaged over subjects. An example is shown in panel b for the time-frequency point giving the maximum phase effect in our study: 7 Hz and -120ms. The grand-average map of performance modulation by phase information at each time-frequency point is shown in panel a. This independent analysis thus confirms the main result of our study (Fig 3b in the main manuscript), i.e. that phase information around 7 Hz shortly before stimulus onset can affect the subsequent detection performance (here with a maximum modulation amplitude of about 14%).



Supplementary Figure 3: Statistical significance of power differences (left) and phase bifurcation (right) between hits and misses at electrode Fz. Analyses similar to those displayed in Figure 2 were performed, but using a wavelet with a length of only one cycle for each frequency. The white filled area masks the time range in which prestimulus and poststimulus effects could have been potentially confounded due to the wavelet transform. The main results from the previous analysis using a longer wavelet are confirmed here, indicating that effects in the prestimulus time window were not due to contamination with poststimulus data.



Supplementary Figure 4: Results at channel Fz (left) and the horizontal electrooculogram (HEOG; right). In the top four panels, each colored horizontal trace represents a single-trial (low-pass filtered at 100 Hz) whose potential variations are color-coded (data are plotted for all subjects). A moving average across 100 adjacent single trials was applied vertically to highlight trial-to-trial consistencies. The bottom two panels show the grand-averaged event-related potential (ERP). We investigated whether prestimulus effects were possibly affected by ocular artifacts. For example, differences in power and phase on hit and miss trials could have simply been caused by differences in the number or the timing of eye blinks prior to stimulus onset. Eye blinks are associated with high amplitude signals at frontal channels. However, we found no evidence for eye blinks on single trials or in the average ERP, either at Fz (where power and phase bifurcation effects were strongest) or in the HEOG (cf. supplementary Figure S3), for hits or misses. Alternatively, hits and misses might have been associated with horizontal saccades towards or away from the target, respectively. If so, these saccades might have introduced an electrical artifact with opposite polarity (i.e. opposite phase angles) for hits and misses. However, inspection of the HEOG on single trials and in the averaged ERP did not reveal any signs of systematic prestimulus saccades, least of all in opposite directions (cf. supplementary Figure S4). The HEOG is sensitive to macro-saccades down to a few degrees visual angle (Hillyard and Galambos, 1970; Lins et al., 1993). We cannot exclude, however, the possibility that the data might be contaminated by artifacts caused by a number of smaller saccades, i.e. so-called micro-saccades, defined as very short saccades of less than 1° (Martinez-Conde et al., 2004). However, such artifacts have generally been reported to affect mostly higher frequencies (i.e. above 20 Hz; Trujillo et al., 2005; Yuval-Greenberg et al., 2008). In sum, these results indicate that power and phase bifurcation effects in the main analysis were not brought about by ocular artifacts.

References

- Hillyard SA, Galambos R (1970) Eye movement artifact in the CNV. *Electroencephalogr Clin Neurophysiol* 28:173–182.
- Lins OG, Picton TW, Berg P, Scherg M (1993) Ocular artifacts in EEG and event-related potentials. I: Scalp topography. *Brain Topogr* 6:51–63.
- Martinez-Conde S, Macknik SL, Hubel DH (2004) The role of fixational eye movements in visual perception. *Nat Rev Neurosci* 5:229–240.
- Trujillo LT, Peterson MA, Kaszniak AW, Allen JJB (2005) EEG phase synchrony differences across visual perception conditions may depend on recording and analysis methods. *Clin Neurophysiol* 116:172–189.
- Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY (2008) Transient induced gamma-band response in EEG as a manifestation of miniature saccades. *Neuron* 58:429–441.

C. A brain rhythm for perceptual framing?

1. “Perceptual Framing and Cortical Alpha Rhythm”

In 1981, Francisco Varela and colleagues published a paper entitled “Perceptual Framing and Cortical Alpha Rhythm” in *Neuropsychologia* (F J Varela et al., 1981). This study has a special meaning for me – it is the first neuroscience research paper I have ever read, and it is while trying to replicate its findings that I learnt the basics of modern neuroscience analytical techniques, from mastering Matlab to computing statistics, via time frequency analysis, machine learning, you name it.

Varela states in the introduction:

it is natural to assume a relationship between temporal framing and the periodic, rhythmic activity detectable in the electroencephalogram (EEG).

He chooses apparent motion as a behavioral task to evidence the temporal framing phenomenon. The subject is seated in the dark and looking at two LEDs, which at each trial briefly light up (6ms duration) one after the other, with a variable Stimulus Onset Asynchrony (SOA). Depending on the SOA, the subject may experience one of three percepts : simultaneity (for shorter SOAs), apparent motion (for intermediate SOAs) or sequentiality (for longer SOAs). For illustrative purposes, Figure 46 is a plot of data collected during our replication attempts, for one subject.

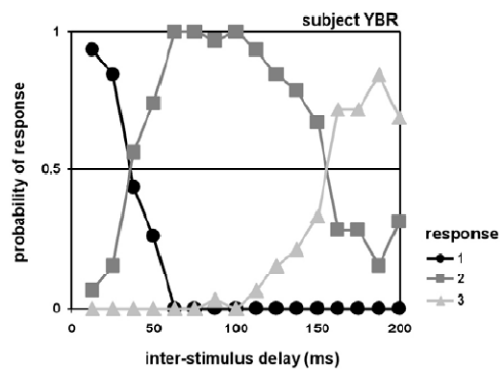


Figure 46 Psychometric curves for a representative subject presented with two flashes with a varying inter-stimulus delay (x-axis) and performing a 3-AFC task (1 : simultaneous, 2: apparent motion, 3: successive). Note the two crossings (between 1 and 2, and between 2 and 3) which define ambiguous inter-stimulus delays.

You can see that at a delay of about 35ms, the subject whose performance is depicted in Figure 46 answered “simultaneous” in half of the trials, and “apparent motion” in the other half. This value of the inter-stimulus delay could be defined as the threshold for perception of apparent movement, or as the Just Noticeable Difference (JND) between the times at which the two flashes are presented, to use the vernacular of classical psychophysics. The physical input is exactly the same, yet the reported percept varies; at 35ms, the ambiguity is maximum. Such situations, in which the only source of variability that can explain the variability of the report is an internal one, are dear to vision scientists who use them to

isolate brain processes specific of one or the other percept. In this case, what is the brain process that can explain perceiving simultaneity VS sequentiality (in the form of apparent motion)? The most obvious interpretation of behavior relies on the perceptual moment hypothesis: if the two flashes fall within the same perceptual epoch, they should be perceived as simultaneous; whereas if they fall into successive perceptual epochs, they should be perceived as sequential (apparent motion). This interpretation suggests that the percept should depend on an ongoing brain rhythm, the rhythm which organizes perceptual epochs.

Varela hypothesizes that the occipital alpha rhythm is the rhythm of perceptual epochs. Varela reports the first ambiguous ISI (between simultaneous and apparent motion) to be 57/67/77ms respectively in his three subjects. In the perceptual moment framework, the first ambiguous SOA corresponds to half the length of a perceptual moment (Figure 47) : this yields an estimate of 114/134/154ms for a moment in Varela's subjects (for our subject, closer to 70ms). This period indeed seems to roughly fall within the range of the occipital alpha rhythm (frequency 7-13hz, i.e. period 77-143ms), so Varela's hypothesis is not completely random. To test it, he set up an apparatus that allows him to start the LED sequence at a precise phase of the ongoing alpha rhythm.




Figure 47 A very schematic explanation of why the first ambiguous SOA corresponds to half the duration of a perceptual epoch. A sinusoidal brain rhythm, filled in light blue, instantiates perceptual epochs (vertical lines). Two flashes (yellow dots) are presented successively, with a fixed SOA (red line) equal to half of the duration of a perceptual epoch. In such conditions, the two flashes fall within the same epoch in half of the cases, and in successive epochs in the other half, leading to a simultaneous percept in half the cases, and a successive percept in the other half.

While I was initially impressed by the whole online apparatus, I realized that it was probably the best way for Varela and colleagues to do this at the time – they didn't have the computational resources we have now to digitally filter data offline, and doing it electronically was the way to go. In this context, they just had to make a note of the phase at which they started each trial and could later analyze behavior accordingly. They performed about 100 trials per subject alternating randomly between three conditions: triggering the LED sequence at the peak of the alpha cycle (as measured from electrode O1), triggering the sequence at the trough of the alpha cycle, or triggering the sequence at random phases of the alpha cycle. The results they presented show a stunning effect of the phase of presentation on perception, and I reproduce here some figures from their paper (Figure 48).




Figure 48 **a** Histogram of the probability of perceiving simultaneity for three subjects, whose first ambiguous SOA is indicated to the right, together with the total number N of trials. The central histograms are responses in the absence of any correlation with the alpha phase. When the lights are triggered either with the positive peak of alpha (left) or the negative peak of alpha (right) extracted from the occipital derivation (O1), a change in probability is observed. In these experiments, the overall probability difference between positive and negative was 0.5. **b** Probability of perceiving sequentiality for two subjects. **c** Probability of perceiving simultaneity for three subjects, in the same format as **a**. In this case, triggering was obtained from a parietal derivation (Pz), not an occipital one (O1) as in **a**. Adapted from (F J Varela et al., 1981)

In essence, the phase of the occipital alpha rhythm at which the first flash is triggered can account for most of the variance in the reported subjective experience; in the case of testing at the first ambiguous SOA, when the negative phase is used to trigger the flash sequence, the percept of the subject is almost deterministic (simultaneous)! (Figure 48a) An effect of phase of the alpha rhythm is found on two subjects at the second ambiguous SOA, when subjects are as likely to perceive either motion or sequentiality (Figure 48b). Also, and critically, the phase of alpha at Pz is not as predictive of perceptual outcome as the phase of alpha at O1 (Figure 48c)... Clearly, the phase of the ongoing alpha rhythm at electrode O1 is highly correlated with perception of temporal order. Figure 49 is the summary figure that Rufin published in 2003 with Christof Koch (Rufin VanRullen & Christof Koch, 2003), to illustrate the findings of (F J Varela et al., 1981).

There was a follow up publication from Varela's group: "Quantitative assessment of the dependency of the visual perceptual frame on the cortical alpha rhythm" by Gho and Varela in a 1989 issue of the Journal de Physiologie (Paris) (Gho & Francisco J Varela, 1988). Gho set out to better understand the relationship between the alpha rhythm and perceptual frames. He used a slightly different display : a single LED showed the two flashes, hence the percept was either one flash or two flashes (notice that subjects may experience three flashes in some conditions (Bowen, 1989), a phenomenon that is under investigation in our group but is not reported by Gho and Varela). Gho sampled different SOAs for each phase of the alpha rhythm at which he triggered the stimulus, in order to derive part of the psychometric curve for the perception of simultaneity in each case. For sake of completeness, he also added two conditions to the triggering: not only did he start stimuli at the peak and the trough (as in the previous study), he also tested the descending and ascending zero-crossings of the alpha rhythm. The results are reproduced in Figure 50.

Gho and Varela found an effect in only 3 of 6 tested subjects; the three subjects that did not show the effect were excluded, based on their lack of alpha blocking upon stimulus presentation – the authors take it as the sign of a weak correlation between visual stimuli and cortical state as revealed by EEG. For the three remaining subjects, Gho and Varela use an estimate of the lateral shift of the psychometric curve to report the effect size. This makes it rather difficult to

compare the effect with the previous study. The one subject whose raw performance they plot does show a sizeable (comparable to 1981) and reproducible (across sessions) effect, as you can see in Figure 50b. But it is their best subject, and they report a much smaller effect (in terms of lateral shift) for the other two. In fact, if you believe in linear estimates of “psychometric curves” based on three points, the lateral shift between extreme conditions is much less than you would expect even for their best subject (16ms, while you may expect a shift of up to half of a cycle, i.e. ~50ms). Finally, the extremes are not found at opposite phases of the alpha cycle, as you would naively expect. Clearly, these observations severely weaken the initial conclusion that the alpha rhythm is a direct reflection of perceptual frames. Given the exclusion of half the subjects, a skeptical observer may go as far as to reject any link... However, it is also quite a different paradigm, because the two flashes were presented in the same location. Can we replicate the effect reported in 1981?




Figure 49 Perception and alpha phase. The perception of a given physical event can be influenced by the phase of the ongoing alpha EEG cycle at which this event takes place. For phases separated by 180 deg, the perception of two successive flashes goes from ‘sequential’ to ‘simultaneous’. Reproduced from (Rufin VanRullen & Christof Koch, 2003)




Figure 50 a Response curve for the probability of perceived simultaneity for subject JG as a function of interstimulus interval. Continuous line drawn by hand. **b** Response curves for subject JG obtained by triggering the onset of the light stimuli with the subject's own alpha rhythm in different phases, indicated in the inset. Only the linear segment of the response curve is shown here; control data from a is reproduced for comparison. Regression lines show a displacement to opposite sides of the control line of variable magnitude. Adapted from (Gho & Francisco J Varela, 1988)

2. Replication attempt

In (Rufin VanRullen & Christof Koch, 2003), there is a footnote corresponding to the description of Varela's (Gho & Francisco J Varela, 1988; F J Varela et al., 1981) results:

Note added in proof: The authors would like to indicate that their own recent attempt to replicate this result in 10 subjects using at least 600 trials per subject was unsuccessful: no significant correlation was found between the phase of alpha activity and stimulus perception on a trial-by-trial basis. A similar negative result was obtained independently by D. Eagleman (personal communication).

When I started my thesis, one of my first projects was to scrutinize the data that Rufin collected in 2003 (which the footnote is based on) to see if I could find any phase effects, casting a wider net (in terms of electrodes and frequencies). I describe these efforts here. They did not meet a clear success – however, they set the stage for further replication attempts, which I did not implement but which I discuss here. They will be implemented by a current PhD student in the lab.

a) Stimulus, procedure and EEG acquisition

Our first replication attempt did not explicitly manipulate the phase at which the first flash was presented. Instead, we enforced random inter-trial intervals, which lead to a more or less uniform distribution of phases. A posteriori offline analysis allows us to see if there is a lawful relationship between the percept and the phase of some prestimulus oscillation, for a given ambiguous interstimulus delay. We also posited that if there was a rhythm on which perceptual framing depended, the position of the flashes should not matter too much. Hence, we had different possible combinations of flashes positions and directions, and initially expected to find an effect independently of this variable.

The experiment was divided into two parts :

- PART 1 : behavior

we presented 2-flash stimuli (4 locations x 2 directions – see Figure 51) with variable and random inter-stimulus delays, from 12.5 ms to 200 ms by steps of 12.5 ms. The subject reported whether she perceived the flashes as 'simultaneous' (1), 'in motion' (2) or 'distinct events' (3). This allowed us to plot a psychometric curve (percept VS delay) and determine two 'ambiguous delays' for each subject, i.e. respectively the delay for which she answered 50% 'simultaneous' and 50% 'in motion' (short ambiguous ISI or delay1) and the delay for which she answered 50% 'in motion' and 50% 'distinct events' (long ambiguous ISI or delay2).

- PART 2 : EEG

we presented 2-flash stimuli (same layout as previously) with 2 ambiguous delays & 3 unambiguous delays in at least 12 blocks of 120 trials:

- Delay = 12.5 ms on 8 trials




Figure 51 Stimulus display in our replication attempt. The dotted contours mark the positions at which flashes can appear, while subjects fixate on the central cross.

- Delay = delay1 on 48 trials
- Delay = (delay1 + delay2)/2 on 8 trials
- Delay = delay2 on 48 trials
- Delay = 200 ms on 8 trials

Subjects kept reporting their perception, as in the behavioral part of the experiment; they were not informed of the manipulation of inter-stimulus delay distribution.

A 32-channel EEG (Neuroscan, El Paso, TX) was recorded continuously (sampling rate, 1000 Hz). The electrode layout was modified from the 10–20 system with an additional row of occipital electrodes, and recorded with electrode Cz as reference. The data was re-referenced to an average reference offline. A hardware notch filter was applied at 50 Hz (European electrical standards) to discard ambient electrical noise.

b) Behavioral results

Each subject performed 4 blocks of 128 trials, yielding 32 trials at each SOA (all conditions together). There were 8 conditions: top counterclockwise/clockwise, bottom counterclockwise/clockwise, left counterclockwise/clockwise, right counterclockwise/clockwise. Hence, 4 trials per SOA were performed for each condition. We initially assumed that the psychometric curves would not depend on the condition, and determined the ambiguous SOAs by considering all conditions together. Figure 52 depicts the psychometric curves of two representative subjects, and the ambiguous SOAs, DELAY1 and DELAY2, determined for each subject, are reported in Table 3.

Subjectively, the perceptual difference between “apparent motion”(response 2) and “sequential” (response 3) was not very clear, and subjects report often guessing rather than making a perceptual decision. The subjective difference between “simultaneous” and “apparent motion” was clearer. Furthermore, the interpretation in terms of perceptual cycles is easier to understand in the case of a judgement of simultaneity than of a judgement of the difference between apparent motion and sequentiality. In the following, we therefore focus on DELAY1 exclusively.




Figure 52 Psychometric curves for two subjects, in the behavioral experiment to determine the ambiguous SOAs DELAY1 and DELAY2. Subjects performed a 3-AFC : red is for “simultaneous”, green for “apparent motion” and blue for “sequential”.

	SUBJECT1	SUBJECT2	SUBJECT3	SUBJECT4	SUBJECT5	SUBJECT6	SUBJECT7	SUBJECT8	SUBJECT9
DELAY1(ms)	31.25	18.75	31.25	31.25	18.75	31.25	18.75	37.5	37.5
DELAY2(ms)	125	93.75	125	181.25	100	125	106.25	125	156.25

Table 3 The ambiguous SOAs determined from the behavioral experiment, for all subjects. The SOAs are all multiple of 6.25ms, the interval between two frame refreshes on the 160hz monitor.

Overall, the behavior of subjects in the EEG experiment, for DELAY1 trials, was consistent with what we were aiming for – they reported about 50% “non simultaneous” percepts on average (Figure 53). However, the different conditions actually differ in the value of DELAY1, as can be seen from some systematic differences. One key factor is whether the two flashes are presented in the same hemisphere or not : at DELAY1, interhemispheric presentation (bottom and top) almost invariably leads to more “simultaneous” reports, suggesting that some aspects of the perception of temporal relationships (e.g., apparent motion perception) are more efficient within a visual hemifield (Figure 54).

Our use of multiple positions, without determining the ambiguous SOA for each position, may complicate the analysis of the EEG data. Based on behavior, it may be a good idea to look separately at interhemispheric and intrahemispheric presentations. Additionally, in the case of intrahemispheric presentations, it is possible that the brain rhythm governing perception of simultaneity will be found in the contralateral hemisphere. Hence, we will investigate the different flash locations mostly independently in the following.

c) Correlation between simultaneity judgement and phase of alpha at electrode O1?

In a first pass, we wanted to look precisely for the effect described by Varela. To do so, we filtered the EEG data in the alpha band (8-12hz), and extracted its phase at the time of presentation of the first flash via a Hilbert transform. We used 8 phase bins, and looked at simultaneity judgements in each of the phase bins. In a plot of percent “non simultaneous”

percept VS phase of alpha at the onset of the first flash (Figure 55), we expected to see a large sinusoidal modulation of report probabilities. At electrode O1, which is where Varela and colleagues found their effect, at the onset of the first flash,




Figure 53 Behavior in the EEG experiment at DELAY1. Each group of bars is for one subject; there are 8 conditions, each represented by a small bar (see legend for color assignment). The average percent “non simultaneous” responses over all conditions, for a given subject, is represented by a large, black edged bar encompassing the colored bars.





Figure 54 At DELAY1, subjects report on average over all conditions about 50% “non simultaneous” percept in the EEG experiment. However, there is a clear distinction between interhemispheric flash presentation and intrahemispheric flash presentation – at DELAY1, interhemispheric presentation leads to significantly more “simultaneous” reports (Student’s T-test, two-sided, $p < 0.01$).

we did not find a significant modulation in either of the 4 conditions – we show the best of the 4 conditions, when the flashes were presented above fixation, in Figure 55. Interestingly, the presentation of the flashes was interhemispheric in the original experiment performed in 1981, and it is with one of the interhemispheric conditions that we see the most promising modulation; it should be noted that their flashes were less far apart (2 degrees apart instead of 5 degrees) and they did not mention a vertical offset with respect to fixation. If our effect is a real effect (which is not warranted by the

data, on average over subjects, so far), the largest modulation we can see in any of the subjects is nowhere near the 50% amplitude that was originally reported (Figure 48). However, our effect could be diminished by the fact that we do not control for pre-stimulus alpha amplitude in this analysis. With their online approach, Varela and colleagues triggered the stimuli when the amplitude of alpha exceeded a threshold. Arguably, in trials for which the ongoing alpha rhythm has more power, we are more likely to see an effect of its phase than in trials for which it is very weak, both for a functional reason (a stronger alpha rhythm is more likely to influence perception) and for analytical reasons (the phase determination of a low power oscillation is less precise than for a high power oscillation).

We computed the event-related potentials (ERPs), separately for “simultaneous” and “non simultaneous”, in an approach similar to (Mathewson et al., 2009). If there is a phase effect dependent on the power of ongoing oscillations, it should appear using this method, which can be seen as automatically weighing the trials linearly by their power. Also, this approach allows to look at multiple pre-stimulus time points. We performed this trial averaging approach, after applying a 25Hz low pass filter to the data (to clean up the ERPs by removing high frequency noise), and the results are shown in Figure 56 (electrode O1). We expected to see pre-stimulus oscillations in the ERPs, at a frequency corresponding roughly to the alpha band, and with a phase shift of ~ 180 deg between the traces corresponding to the two percepts. We do see some oscillations in the prestimulus ERPs at each of the flash locations (and these cancel out when all trials are averaged together, which validates our studying different flash locations independently), but they are best in the TOP condition – 180 deg out of phase between the traces corresponding to the different judgements, as expected, and a period of about 100ms. Using multiple t-tests, we find some time points that differ “significantly” (at a very loose threshold, not corrected for multiple comparisons) between the prestimulus traces...



Phase of ongoing alpha rhythm at the onset of the first flash

Figure 55 Varela (F J Varela et al., 1981) reported a large effect of the phase of the alpha rhythm at electrode O1 on subjective simultaneity judgement. **(left)** When flashes are presented above fixation, we see a small, non-significant effect on average over 9 subjects. The average across subjects is the continuous red line; the dotted red lines represent the standard error **(right)** Experimental points (red line) and sine wave fits (pink dashed line) are reported for all subjects.

We had a precise hypothesis, from Varela and colleagues 1981 paper : that the phase of the alpha rhythm recorded at electrode O1 before the onset of the first flash would be predictive of the percept when two flashes are presented with an ambiguous delay. We tested it and there are some trends in the data, but we can not claim to have reproduced the effect from this dataset. I do see some encouraging results, which need to be confirmed with more subjects.

But why limit ourselves to the electrode that Varela recorded from? It is possible that, though the effect can be picked up on this O1 derivation, it would in fact be even more pronounced at another derivation. The method that we developed for PAPER 4 is well suited to answer this question of where the maximal effect lies: we can cast a wide net in terms of prestimulus times and frequencies of interest, as well as extend the analysis to a set of electrodes, and look for phase effects, backed up by randomization tests.

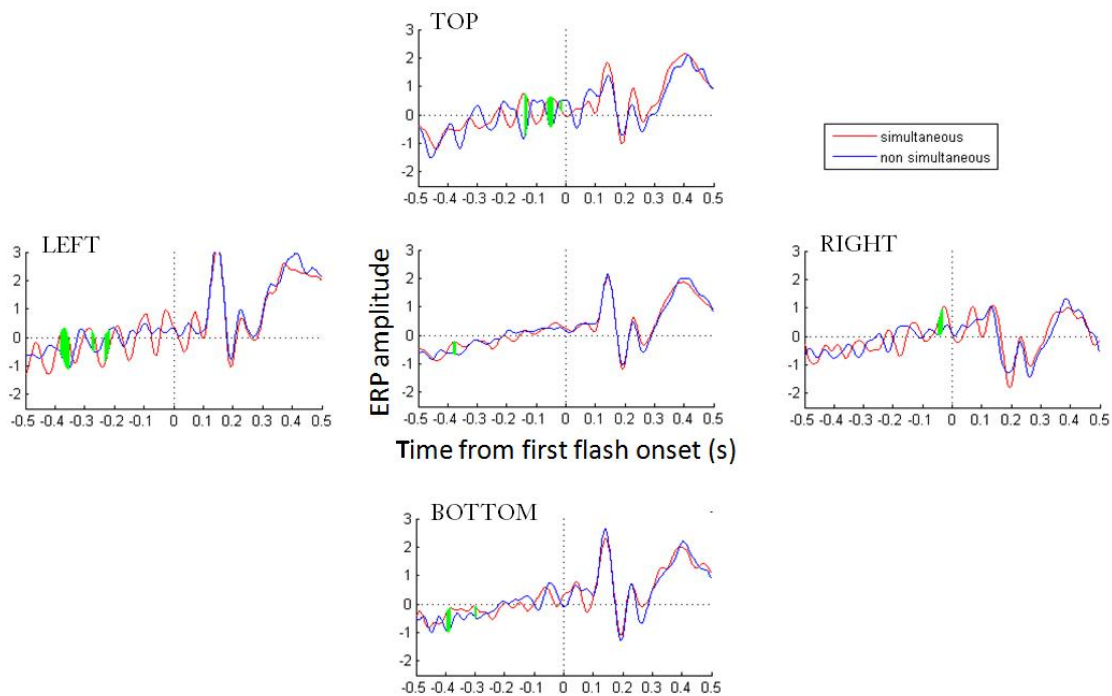


Figure 56 ERP analysis. The grand average Event Related Potential at electrode O1 is calculated, collapsing conditions together (plot in the center) or keeping them separate (the plots' locations corresponds to the locations of the first flash in the experiment, for easier visualization), after sorting trials by the reported perceptual outcome. In green are the time points that reach significance for a double sided Student's t-test, with $\alpha=0.05$.

d) Correlation between perception of simultaneity and phase of any rhythm at any electrode?

We computed a phase locking factor for trials when the percept was “simultaneous” and for trials when the percept was “non-simultaneous” separately, at each electrode, time point and frequency band. To extract the phase at different times and frequencies, we used a wavelet transform, whose characteristics (number of frequencies, number of cycles, corresponding time resolution) are plotted in Figure 57.




Figure 57 Structure of the wavelet transform used to extract phases. The number of cycles increases with the frequency of interest, in an arbitrary fashion, so as to have a “good” compromise between time resolution, frequency resolution and signal-to-noise.




Figure 58 Average of the phase locking factor maps computed separately for “simultaneous” and “non simultaneous” trials, and averaged across electrodes. The four plots correspond to the four possible flash locations : TOP, LEFT, RIGHT and BOTTOM. The effects of ongoing rhythms are seen to the at negative times (before the onset of the first flash). The large phase locking after stimulus presentation corresponds to the event related potential.




Figure 59 Topography of the increased phase locking in the 8-12Hz band from -270ms to -100ms, for the TOP flash location. The effect is maximum at electrode 26 (Oz).




Figure 60 (top) Average of the phase locking factor maps computed separately for “simultaneous” and “non simultaneous” trials, over electrode Oz. **(bottom)** corresponding p values, plotted as $-\log_{10}(p)$. No point survives FDR correction (FDR at 0.05 yields $\alpha=0.0039$, i.e. $-\log_{10}(\alpha)=2.41$).

We show the average of the two separately computed phase locking factor maps across electrodes in Figure 58. In these maps, we are not interested in the large phase locking which occurs after stimulus presentation (corresponding to the Event Related Potential), but in the small phase locking effects before the first flash appears – these reflect the ongoing rhythms. As you can see, the oscillations that we observed in the prestimulus ERP, for the TOP and LEFT positions, translate into increased phase locking in the 5-15Hz frequency band, which can be seen on average across electrodes. By zooming in on the region of increased phase locking, in the alpha band (8-12Hz) and from -270ms to -100ms, we can study its topography and we find that phase locking is maximum at electrode Oz (just between electrodes O1 and O2, the ones that

Varela was recording from). There is also a hot spot in the same location in the topography for the LEFT presentations. This prompts us to look at this electrode specifically, and perform randomization tests on its average phase locking maps. The results are represented on Figure 60: based on 100,000 surrogate grand means, the average phase locking factor was not significantly higher than predicted by the null hypothesis, after correction for multiple comparisons using FDR. Of course, this approach is again oblivious to the power of the ongoing alpha rhythm – however, to really address this problem and select trials accordingly, we need more data than we collected...

3. Summary and Discussion

We were initially expecting a huge effect of the phase of the alpha rhythm on the perception of simultaneity, when two visual flashes are presented sequentially at a delay leading to half “simultaneous”, half “non simultaneous” responses. Accordingly we presented flashes at different locations, thinking it would not matter much. However, we saw that the location of flashes, especially whether they are presented in one hemisphere or across hemispheres, has a significant effect on the perception of simultaneity. There are, in essence, three major flaws in the current dataset : i) there are no conditions when we chose the proper ambiguous delay ii) we have 4 times less trials than we could have collected, had we focused on a single location for the flashes iii) the percept was unclear when testing DELAY2, hence we have 2 times less trials than we could have collected, had we focused on DELAY1 only. Collecting a new dataset, with a single location, a single ambiguous delay, and as many total trials, may thus greatly increase our chances of finding an effect – however, we know for sure that the effect we will find, if we find one, will not be as large as reported by Varela.

We were not able to evidence any significant effect of the phase of the ongoing alpha rhythm (nor any other rhythms) on the perception of simultaneity in our paradigm, even when focusing the analysis on a single location. There were some promising trends in two out of the four conditions, but nothing reaching significance.

Another potential problem with our experiment is that the flashes were quite bright. It is possible that a strong transient resets the rhythm that normally underlies temporal framing. In our next attempt, we will use flashes that are close to detection threshold, so that the underlying rhythms are least perturbed by them. Also, we will do extensive psychophysics to try and find the conditions that maximize the perceptual difference between simultaneous and sequential, at an ambiguous delay – varying the distance between flashes, the size of the flashes, their position of the screen, ...

Linking the temporal framing phenomenon to ongoing brain rhythms has thus eluded us thus far. We do not rule out such a relationship for good yet – at this point it is clear that Varela’s 1981 result is not as robust as they suggested in the original paper, and that the relationship to be found may be a little more complex than previously envisioned – e.g., dependent on the location of the flashes. This in fact raises an important concern : perhaps there is a feature, computed from the continuous flow of information coming from the brain, which allows to judge simultaneity at a resolution finer than that of perceptual moments. There may for instance be a flicker detector, as there is a first order motion detector, and the output of this detector would be represented within each perceptual moment. The readout frequency by attention would, in that case, not affect perception of simultaneity...

D. Does the spotlight oscillate between concurrently attended locations?

A central part of the theory that we are pushing for throughout this thesis is that focal attention is inherently periodic, and takes samples at the attended locations. Some electrophysiological evidence for this was put forward in section III.1 (page 103), in a situation when only one location was monitored. What happens when more than one location is attended to, i.e. when attention is divided in a sustained fashion? We have addressed this question with a psychophysical paradigm in section II.2 (page 52). We mentioned then that some electrophysiological evidence had been published in favor of a sustained division of the attentional spotlight. We review this evidence here.

1. A sustained division of the attentional spotlight has not been convincingly shown yet

Müller, Malinowski, Gruber and Hillyard published a very nice study in *Nature* in 2003, entitled “Sustained Division of the Attentional Spotlight” (M M Müller et al., 2003). They make use of the frequency tagging technique in EEG to follow the deployment of attention in different conditions. We will explain the frequency tagging technique in more detail in the following section – for the purpose of understanding this study, it is enough to know that a flickering stimulus evokes an oscillatory neural response of the same temporal frequency as the driving stimulus (the steady-state visual evoked potential, or SSVEP) (D. Regan, 1977), which can be picked up on some scalp electrodes, and which is modulated by attention (Morgan, Hansen, & S. a Hillyard, 1996; M M Müller, Picton, et al., 1998). This is a very robust paradigm.

Müller’s idea was to concurrently present four RSVP (Rapid Serial Visual Presentation) streams, each updated at slightly different rates so as to evoke SSVEPs at different frequencies. Also, the locations at which the streams were presented, along the horizontal meridian, two on each side of fixation, ensured that the maximal response would be picked up on different electrodes (Figure 61). The same five symbols were used in all streams; one of the symbols was designated as the target symbol. In a given block of trials, subjects had to attend to two of the four streams; their task was to report when the target symbol appeared in both attended locations simultaneously, which could occur 0 to 3 times per trial (each trial being 3.06 seconds long).

There were four attentional conditions: attend to both streams in the left hemifield (1+2), attend to both streams in the right hemifield (3+4), or attend to one stream from each hemifield (1+3 or 2+4). In the last two conditions, there is one intervening irrelevant stream – the question asked in this study is : can attention select the two disjoint locations of interest without enhancing processing at the intervening location?

A summary of the results is given by the simple bar graph in Figure 62 (a good way to sell your results to *Nature* – they like when you don’t need a statistician...). In essence, when a location is attended, the SSVEP response is enhanced, compared to when it is unattended: this is the classical result. The key finding is that when an unattended location is between two attended locations, it does not receive any more attentional resources than an unattended location outside.




Figure 61 Schematic diagram of stimulus sequences, electrode positions and typical SSVPEPs for attended (dashed lines) and unattended (solid line) conditions. Subjects reported simultaneous occurrences of the target symbol “8” at the two attended positions (either 2+4, 1+3, 3+4, or 1+2 on different blocks of trials). Stimulus sequences show examples of targets at the attended positions under the four conditions. SSVPEP waveforms were obtained by a sliding average technique in the time domain. SSVPEPs shown were obtained under conditions of attention to positions 1+2 and to positions 3+4. Reproduced from (M M Müller et al., 2003)




Figure 62 Normalized amplitude values averaged across the 8.7 Hz and 20.3 Hz SSVPEPs (at positions 2 and 3, respectively) under different attention conditions. **Indicates the highly significant difference ($P < 0.0001$) between the conditions when the 8.7 Hz and 20.3 Hz stimulus sequences were attended (that is, attend 2+4 for 8.7 Hz and attend 1+3 for 20.3 Hz) as compared to conditions when these positions were intermediate and ignored (that is, attend 1+3 for 8.7 Hz and attend 2+4 for 20.3 Hz). *Indicates the significant difference ($P < 0.05$) between conditions when adjacent positions were attended (that is, attend 1+2 for 8.7 Hz and attend 3+4 for 20.3 Hz) as compared to conditions when separated positions were attended (that is, attend 2+4 for 8.7 Hz and attend 1+3 for 20.3 Hz). No significant differences were found when comparing unattended intermediate positions versus unattended positions in the hemifield opposite to the attended adjacent positions. Reproduced from (M M Müller et al., 2003)

This shows unambiguously that two disjoint locations can be attended to, without any enhancement of intervening locations – it rules out a single spotlight model which would be enlarged to encompass the attended locations. A secondary finding is that the attentional enhancement is larger when the two locations to be attended are disjoint – this could

however be due to a confounding intrahemispheric VS interhemispheric allocation of attention, given the postulated independence of attentional resources in the two hemispheres (Patrick Cavanagh & Alvarez, 2005).

While this is a very nicely conducted study, which has been taken by many as definite evidence for the possibility to divide attention, it does not demonstrate a true sustained split of attention. A closer look at the methods shows that the amplitude of the SSVEPs was estimated over a 2500ms window (from 526ms to 3026ms in each trial). This means that this method is unable to discriminate between a true sustained split and an oscillatory behavior of the spotlight, selecting a single location at any given time. Müller and colleagues were aware of this. They wrote :

The use of brief (181 ms) target durations rules out the possibility that subjects were rapidly switching attention between the separated locations to achieve accurate performance rather than dividing attention between the locations. Although different studies have obtained varying estimates of the minimal time for switching attention, there is general agreement that the minimum time needed to identify a target at one location and then switch attention to identify a target at a second location is in the range of 200–500 ms.

Their argument is insufficient to rule out the switching spotlight model, despite the reported 83% hit rate (and 5% false alarms). It is possible, until strong evidence is provided against it, for a switching spotlight to take a sample at each location in 181ms (a 7hz spotlight may for instance take two samples in ~140ms).

There was another study, a year later, using fMRI to make a very similar claim (McMains & Somers, 2004). The study was entitled “Multiple Spotlights of Attentional Selection in Human Visual Cortex” and published in *Neuron*. It also relied on multiple RSVP streams; this time, they were streams of letters, and the targets were digits. Subjects had to attend to two streams concurrently, and decide whether the digits that appeared simultaneously were the same or different (there was also a condition in which subjects attended to a single location – they just had to respond when they detected a given digit). The duration of presentation of each item was 173ms. McMains and Somers show evidence for attentional enhancement at disjoint spatial locations, with no enhancement at intervening locations; compared to the Müller study, they also show that there can be two disjoint spotlights of attention within a single hemifield (Figure 63a). Their study is usually also taken as strong evidence for a sustained division of the attentional spotlight – as opposed to a single zoom lens which would grow to encompass attended locations. Our criticism is of course valid in the case of this fMRI study : the BOLD imaging method is known to have extremely poor temporal resolution, and hence it provides no evidence whatsoever against the alternative account of a switching spotlight (in this experiment, the signal was averaged over 34s to get the pretty activation maps...). Again, McMains and Somers were aware of this limitation, and presented an argument similar to that of Müller and colleagues to discard the serial switching spotlight account. To be completely safe from criticism though, they conducted an additional psychophysics experiment to support their claim. They varied the rate of RSVP presentation and measured corresponding performance. They reasoned that a strictly serial model should predict that the performance in the ATTEND1 and the ATTEND2 conditions would differ by a factor of 2. Their results showed otherwise – attending to two streams entailed a very slight cost in terms of the performance vs RSVP presentation rate relationship (Figure 63b). They thus concluded that this invalidated the serial account in their paradigm. We do not think that this control is convincing – for one thing, the tasks in the attend 1 and attend 2 conditions are not comparable, which hinders any further speculations.




Figure 63 a (left) Right hemifield spatial layout of the five RSVP streams. In the ATTEND1 condition, subjects covertly attend to the uppermost stream (blue). In the ATTEND2 condition, subjects attend to the lowermost and central streams (red). **(right)** Subject C data in the right hemifield configuration **b** Psychophysical Results. Average ($n=5$) subject performance in ATTEND1 (filled/solid) and ATTEND2 (unfilled/dashed) tasks as a function of RSVP letter duration. Linear model fits (lines) reveal an 8 ms difference for threshold level performance ($d'=1$). Adapted from (McMains & Somers, 2004)

We therefore argue that the question of the mechanism of sustained divided attention is still an open question. Methods such as electroencephalography are usually thought of as having an extremely good temporal resolution – however, this resolution had to be sacrificed in the Müller et al study (M M Müller et al., 2003) in favor of signal-to-noise ratio. Is there a way to take full advantage of the temporal resolution of EEG?

2. Our methodological approach : from SSVEPs to broadband frequency tagging

a) How the brain responds to different stimulation frequencies

One of the first thorough descriptions of SSVEPs is found in (D. Regan, 1977)– much can be learnt from this “oldie”. Regan argues that the visual system is nonlinear, hence its response to periodic stimulation will yield information that is not redundant with its response to transient stimulation. Steady-state responses are evoked by “shaking the system gently”, rather than “kicking” it. Regan shows how reliable steady-state responses can be, provided the pass band for the Fourier analysis is narrow enough – meaning that the temporal window over which the analysis is conducted should be long enough (a 7s filter yields a very stable SSVEP). If the pass band is too wide, it quickly becomes difficult to tell the response from the background noise.

The SSVEP is dependent on many factors. The spatial frequency of the stimulus is a critical variable : for instance, an unpatterned flickering light will not evoke the same response as a flickering checkerboard (and the spatial frequency of the checkerboard will also influence the response). The intensity of stimulation, as well as its color properties, also affect the amplitude of the response. But most importantly, not all stimulation frequencies are equally potent in entraining neuronal populations. Regan showed a compilation of results from different studies, focusing on the amplitude of the steady state response VS the frequency of stimulation, for an unpatterned large patch of flickering light. He determined that there were

three distinct ranges of frequencies which seemed to elicit responses from different functional networks (Figure 64a) : a low-frequency range (around 10Hz), a medium-frequency range (13-25Hz, with a peak at 18Hz) and a high-frequency range (40-60Hz). Interestingly, the amplitude of steady-state evoked potentials does not seem to correlate with perceived flicker threshold or sensory magnitude. More than twenty years later, a systematic study of steady-state responses evoked by flicker at 1 to 100Hz (Christoph S. Herrmann, 2001) came to very similar conclusions, finding resonance peaks around 10, 20 and 40Hz (Figure 64b).

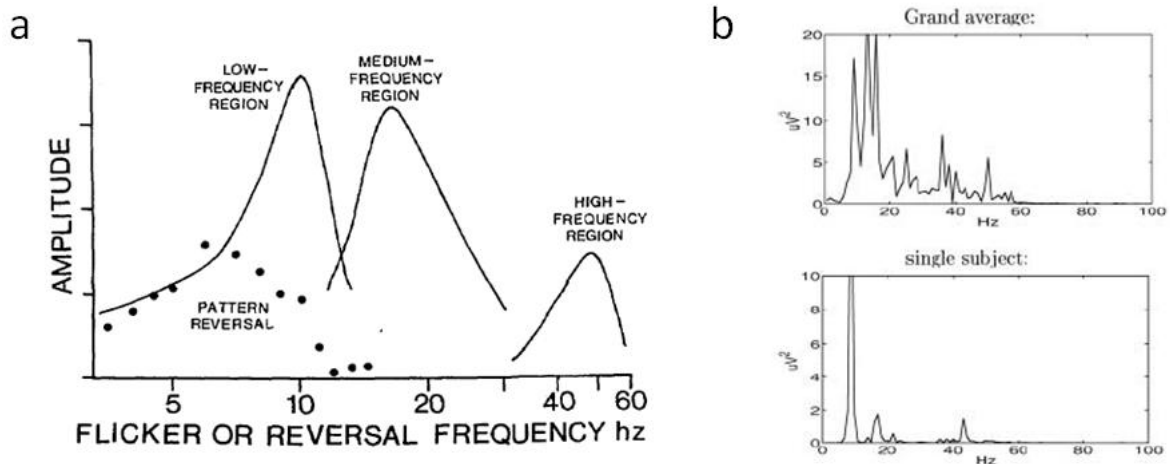


Figure 64 a Effects of stimulus frequency upon flicker SSVEPs and pattern SSVEPs. Continuous lines plot SSVEP amplitude versus flicker frequency, showing "low-frequency," "medium-frequency," and "high-frequency" SSVEPs. The SSVEPs were elicited by a spatially unpatterned (blank) patch of light. The dotted line shows the quite different plot of SSVEP amplitude versus pattern reversal frequency for pattern SSVEPs. Reproduced from (D. Regan, 1977) **b** Profile of the fundamental frequency for the average across all ten subjects (**top**) and for one single subject (**bottom**). Clear resonance peaks are visible around 10, 20 and 40 Hz. Reproduced from (Christoph S. Herrmann, 2001)

b) Attentional modulation of SSVEPs

Regan stated in 1977 (D. Regan, 1977):

steady-state EPs do not seem to be much affected by psychological variables such as attention. [...] if you wish to study, e.g., attention or the orienting response, then steady-state EPs are most likely completely useless, so that you must use transient EPs.

This statement has been seriously revised since. Morgan and colleagues (Morgan et al., 1996) published the first evidence of a significant modulation of the amplitude of SSVEPs by spatial focal attention. Subjects were cued to attend to one of two RSVP streams (one in each hemifield, on the horizontal meridian, 5.7deg from fixation), looking for a rare '5' among letters; the black letters were superimposed on a white flickering square, turning on and off at a frequency of 8.6Hz for one stream and 12Hz for the other. Each of the two corresponding SSVEPs was substantially enlarged when attention was directed towards it. The authors point out that the effect was found to be more pronounced over electrodes in the right hemisphere, in particular for the 12Hz SSVEP; which contrasts with attentional effects of classical VEPs after transient

stimulation, usually found contralateral to the hemifield of stimulation for some components and ipsilateral for some others. Following Morgan et al's study, Müller and colleagues (M M Müller, Picton, et al., 1998) extended the findings to medium-frequency range potentials. Since, SSVEPs have proved to be a powerful tool to characterize spatial attention (J. Ding, George Sperling, & Srinivasan, 2006; M M Müller et al., 2003; Matthias M Müller & Hübner, 2002).

There is also a whole corpus of studies which is largely ignored by visual scientists – the brain-computer interface (BCI) literature. Researchers in that field have quickly put the steady-state potential technique to good use – with perhaps the first report published in 2000 (Middendorf, McMillan, Calhoun, & K. S. Jones, 2000). These authors used the increase in SSVEP amplitude which results from directly looking at a flickering stimulus (compared to having it in the periphery) as signal – in other words, they created an overt-attention based BCI, capable of telling which of two locations (flickering at different rates) a subject was looking at; you may call it a two-position SSVEP-based eye-tracker. Much research has been done since, most interestingly focusing on covert attention shifts. For instance, Kelly and colleagues (Kelly, Lalor, Reilly, Member, & John J Foxe, 2005) achieved a 71% classification rate when attempting to discriminate which of two locations flickering at different frequencies (one in the left hemifield and one in the right hemifield, respectively 9.45Hz or 14.17Hz and 10.63Hz or 17.01Hz) was being attended to on single trials (using a roughly 3.5s long window for SSVEP amplitude estimation). This performance was obtained using solely information from SSVEP amplitudes; additional features can enhance performance, for instance the distribution of occipito-parietal alpha power (Bahramisharif, van Gerven, Heskes, & Ole Jensen, 2010; van Gerven, Bahramisharif, Heskes, & Ole Jensen, 2009; Kelly, Lalor, Reilly, & John J Foxe, 2006; Kelly et al., 2005). These results from the BCI literature are instrumental in quantifying the reliability of SSVEP amplitude modulations by attention. In brief : these modulations are reliable on single trials. Single-trial amplitude estimates of SSVEPs usually rely on temporal windows of about 3s, in the BCI literature as well as in the Müller et al study (M M Müller et al., 2003) that we described in detail. Is it possible to reduce the temporal window of analysis dramatically, to get to the question that we are interested in addressing: does the spotlight of attention oscillate between concurrently attended locations? From previous evidence, we may expect the spotlight of attention to take samples at about 7Hz, i.e. every 140ms. This means that we need to use a temporal window of about 70ms at most, preferably less, to accurately follow the position of the spotlight. Using such a narrow temporal window will mean using a fairly wide pass band, hence a lot of the background noise will be included in the amplitude estimates – and the signal-to-noise ratio will become too low for classification purposes. Our goal is to find a way around this.

c) Increasing the temporal resolution of the SSVEP technique

The original idea that we came up with is to integrate the signal over a range of frequencies instead of integrating it over time, in order to up the signal to noise ratio whilst keeping the best possible temporal resolution. As choosing specific frequencies would have been an arbitrary step (but it may be worth a try as well, informed by the individual resonance frequencies of each subject, cf Figure 64b), we decided to use broadband flicker – letting underlying neural networks pick whatever frequencies were present in the stimulus at any point in time. We tentatively coined the term “broadband frequency tagging” for this technique. Integrating information over frequencies seems like a good idea, as it has been shown that different stimulation frequencies drive different networks (J. Ding et al., 2006; D. Regan, 1977) – hence, it may be possible to drive a few different neuronal populations simultaneously, and combining the information that they yield will increase the signal to noise ratio. Will the brain be driven by broadband flicker – i.e., white noise – the same way that it is driven by a sustained, periodic stimulation? This is an open question. Clearly, the analysis of broadband frequency

tagging will differ from the classical analysis of SSVEPs, and will rely on a direct comparison of the instantaneous power and phase of oscillations in the stimulus and in the EEG.

3. Where is the spotlight of attention now?

We designed an experiment in which subjects either had to attend to one or two flickering locations, in order to detect a target (contrast decrement). We plan to use the ATTEND LEFT and ATTEND RIGHT conditions to find a reliable instantaneous EEG signature of attentional orientation, then to use this signature to follow attentional orientation in real time in the ATTEND BOTH condition. We expect to see an oscillation of the position of the attentional spotlight between the two attended locations in the latter case.

a) Stimuli, procedure and EEG acquisition

The experiments were created and run with the Psychophysics Toolbox version 3 (Brainard, 1997) and Matlab 2007a (The Mathworks, Inc., 2007) on a Windows PC attached to a 17" monitor at a resolution of 640 x 480 and a refresh rate of 160 Hz. A viewing distance of 57 cm was maintained throughout the experiment. Two luminance modulated discs of 3.5 degrees radius were presented simultaneously on either side of fixation at 7.5 degrees of visual angle, on a black background, for 6.25s (Figure 65). The luminances sequences were designed to be random and independent of each other on any given trial. Importantly, the sequences were constructed to have equal power at all frequencies between 0 and 80Hz – this is the range of frequencies that the monitor's 160Hz refresh rate permits. More specifically, to create a luminance sequence, we used the following procedure. We generated a sequence of 1000 random numbers (6.25s corresponds to 1000 frames at a refresh rate of 160Hz). In the Fourier domain, we then normalized the amplitudes of all components (between 0.16 and 80hz, the Nyquist frequency). Back in the time domain, we set any luminance values greater than 255 or less than 0 as white or black, respectively. Sequences thus ranged from black (0.1 cd/m²) to white (59 cd/m²).

Observers covertly monitored either or both of the stimuli (depending on the attentional condition, see next paragraph) to detect a 1s-long target square (size 3.75 degrees) appearing inside the disc on a random 25% of trials. The area within the square followed the same sequence of luminance changes as the disc stimulus, but scaled in amplitude (by multiplication by a number between 0 and 1, adjusted individually using a staircase procedure) so that detection performance was fixed at approximately 82%. Observers were instructed to press a button at the end of the sequence if they had detected the stimulus. Participants completed 480 trials, in 12 blocks of 40 trials – attentional conditions were blocked. There were three attentional conditions : attend to the left disk (3 blocks, 120 trials), attend to the right disk (3 blocks, 120 trials), and attend to both disks (6 blocks, 240 trials). Block order was randomly determined at the start of the experiment. In the first trial of each block, the locations of the to-be-attended disks were circled in red – this cue was 100% valid.

EEG was recorded with a sampling rate of 1024 Hz using a 64-channel ActiveTwo Biosemi system. Horizontal and vertical electro-oculograms were recorded by 3 additional electrodes around the subjects' eyes. Data were downsampled off-line to 500 Hz and epoched from 0.25 s before stimulus onset to 8.25 s after stimulus onset. Stimuli luminance sequences were upsampled to match the sampling rate of the EEG. EEG epochs were baseline-corrected. We applied a notch filter (49-

51Hz) to remove european 50Hz line noise. All epochs were inspected for artefacts and discarded where necessary. A wavelet transform was applied to the stimulus and the EEG timecourses, using built-in EEGLAB function `timefreq()`; its profile is shown on Figure 66.

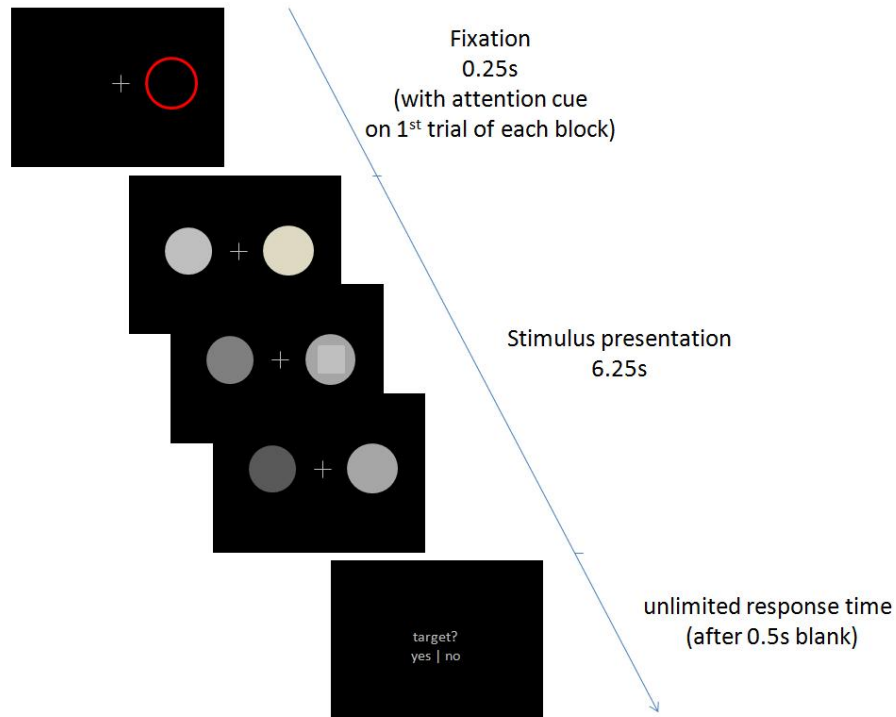


Figure 65 Stimulus display for the Broadband Frequency Tagging experiment




Figure 66 The wavelet transform used in the analysis.

b) *Occipito-parietal alpha indexes attentional allocation*

Our first concern is to make sure that our subjects were indeed performing the experiment correctly and deploying attention according to block-wise instructions. The problem is that our cue was 100% valid – hence, it is impossible to infer that the subjects correctly oriented their attention based on a comparison of performance between attended and non-attended locations, as is usually done in a Posner-type paradigm.

A now well established result in the literature (Bahramisharif et al., 2010; van Gerven et al., 2009; Kelly et al., 2006; Gregor Thut, Nietzel, S. a Brandt, & Alvaro Pascual-Leone, 2006; Worden, J J Foxe, Wang, & Simpson, 2000; Yamagishi, 2003) is that there is a lateralization of occipito-parietal alpha when attention is directed in a sustained manner to a location in one hemifield – contralateral alpha power decreases while ipsilateral alpha power increases, and the difference of alpha power between a subset of occipito-parietal electrodes in the left hemisphere and a corresponding subset of occipito-parietal electrodes in the right hemisphere is a reliable index of attentional allocation. In fact, the distribution of occipito-parietal alpha power has even been used to track the position of the attentional spotlight around a circle (Bahramisharif et al., 2010). This neural signature of attentional allocation was chosen as our way to control whether subjects behaved correctly. We thus computed an alpha lateralization index for our ten subjects, in each of the three attentional conditions (attend left, attend right and attend both). Note that on average, occipito-parietal alpha power tended to be higher in the left hemisphere than in the right hemisphere (Figure 67) in our subjects – in the ATTEND BOTH condition, power is not equal in the two hemispheres. However, compared to the baseline given by the ATTEND BOTH condition, we find the predicted modulation of alpha power lateralization (defined as average normalized amplitude in left occipital electrodes minus average normalized amplitude in right occipital electrodes) in all but two of our subjects (Figure 68). These two subjects were discarded from further analysis, since we could not be sure that they had followed the attentional instructions.

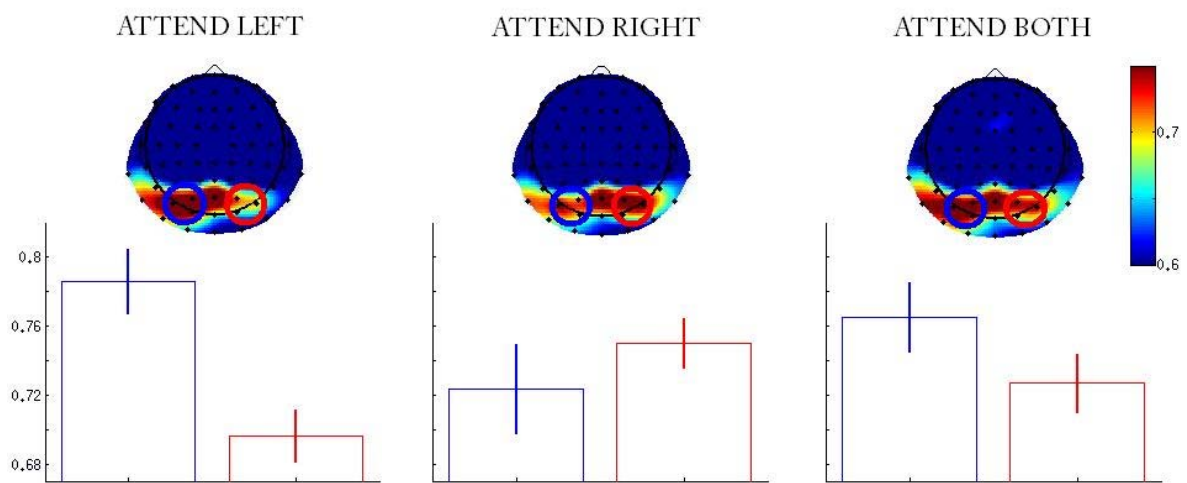


Figure 67 Average distribution of alpha power over 8 subjects, over the interval 1 to 5.5 seconds in each trial. (top) The scalp maps show alpha power at all electrodes, normalized for each trial to the maximum alpha power recorded at any of the electrodes. (bottom) The bar graphs show the average of three electrodes in each hemisphere. Notice that alpha power in the left occipito-parietal region is higher than in the right occipito-parietal region in the ATTEND BOTH condition (baseline). Attending to the left or right modulates the distribution of power.

One may wonder whether this lateralization index could serve as an “instantaneous” indicator of attentional allocation. It is more of a theoretical question than an empirical one. Can the alpha rhythm be modulated this quickly? Can an oscillation be modulated in amplitude at a frequency practically similar to that of its own rhythm? This does not seem plausible neurophysiologically. If attention indeed switches at 7hz or so, the distribution of occipito-parietal alpha power is unlikely to be the mechanism subserving the switching.

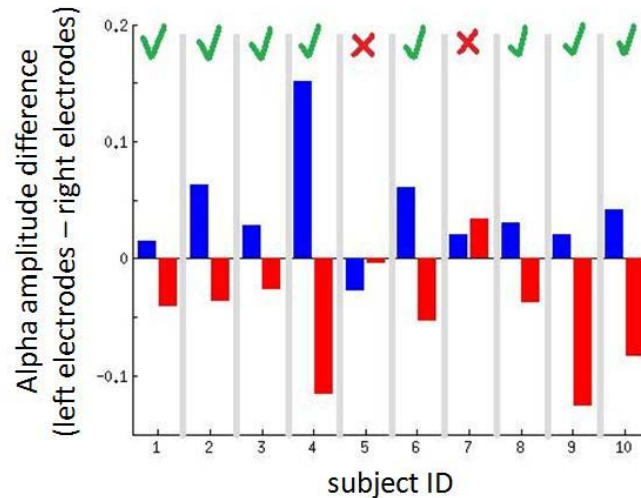


Figure 68 Baseline corrected alpha lateralization index for each of the ten subjects. The lateralization index is computed as the normalized alpha amplitude in the left occipital electrodes (cf Figure 67) minus the normalized alpha amplitude in the right occipital electrodes. The ATTEND BOTH condition serves as baseline. Two subjects do not show the expected modulation (red crosses) and were discarded from further analysis.

c) *The EEG phase locks to stimuli (mostly at lower frequencies)*

After establishing that our subjects were following attentional instructions, the next step was to verify that there was some sort of frequency-dependent response in the EEG. Indeed, our paradigm is not a classical SSVEP paradigm, and it was in fact unclear whether our stimuli would elicit a sustained response at all in each frequency band. We looked for power and phase relationships between the stimuli and the recorded EEG. More specifically, we tried to correlate the instantaneous power at different frequencies in the stimuli with the instantaneous power at the corresponding frequencies in the EEG, varying the delay between the stimulus and the EEG – we expected the brain to reflect the transient fluctuations of power in the EEG at the different frequencies under investigation, after a certain delay (processing time). We found no reliable effects using this approach. We then looked at phase synchrony between the stimulus and the EEG in each frequency band, i.e. the degree to which the phases of the two filtered time courses kept a constant relationship over time points and trials, again varying the delay between the stimulus and the EEG. We found that, on average over electrodes and attentional blocks, there was increased phase locking between the stimuli and the EEG, starting at a delay of about 80ms (Figure 69). The sustained EEG response, as indexed by phase locking, appeared most distinguishable at frequencies below about 30hz.

We looked at the topography of the main effect, in the 8-15Hz frequency band. We found that the early response (80-150ms) is mostly occipital, then it spreads to frontal electrodes (200-350ms) and the later response is mainly occipital again (400-550ms), as seen on Figure 70.

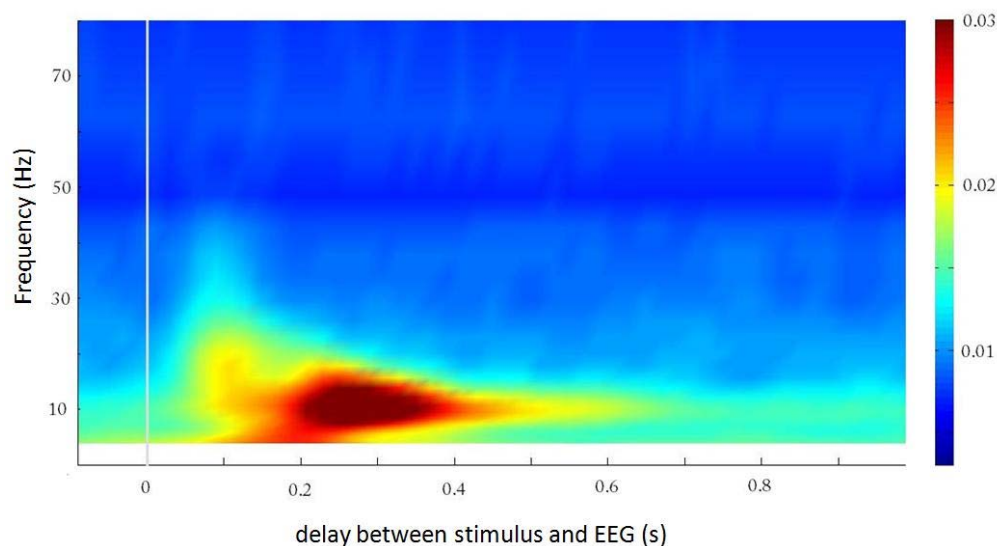


Figure 69 Grand mean phase locking of the EEG with the stimulus, for the 8 selected subjects (average over all electrodes and attentional conditions). There is a sustained response to broadband flicker, indexed by this phase locking measure, mostly for frequencies below about 30hz. Note the prolonged response in the alpha band.

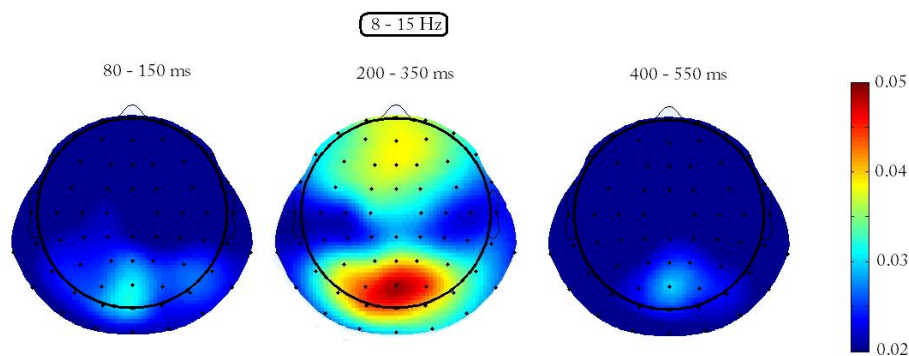


Figure 70 Topography of the average phase locking in the 8-15Hz band, at three different delay ranges. The early response of the EEG is mostly occipital then engages frontal areas as well.

d) *Attention increases phase locking*

Is the response that we evidenced affected by attentional allocation? Meaning, is there more phase locking with the attended stimulus than with the unattended stimulus on a given trial? This modulation would be crucial if we are to use this measure to decode the position of the spotlight. We computed the difference in phase locking between the attended and unattended stimuli, in the unequal allocation conditions (ATTEND L and ATTEND R). Figure 71 shows that there is indeed more phase locking with the luminance pattern of the attended stimulus, and Figure 72 shows the topography of the effect. Attention mainly affects phase locking at occipital electrodes in the two hemispheres, this time at delays of

about 200ms (indeed, it appears that phase locking to the stimulus is strongest at around 80-100ms, but that this 'automatic' response is not modulated by attention).




Figure 71 Grand mean phase locking difference between the attended stimulus and the unattended stimulus and the EEG. The effect of attention is mostly confined between 10 and 20Hz, at a delay of about 200-300ms.




Figure 72 Topography of the attentional phase locking difference, in the 12-20Hz frequency band, between 200 and 280ms

e) The phase locking information available on occipital electrodes can be used to decode attentional allocation (with modest performance)

We used a linear decoding framework to combine information optimally between stimuli, electrodes, delays and frequencies – we first picked some delays (50ms to 400ms), frequencies (12Hz to 40Hz) and electrodes (10 occipital electrodes) by hand, to reduce feature space (these choices were motivated by the findings described in the previous

pages). We had two classes to try and tell apart: time points belonging to trials in the ATTEND L condition, and time points belonging to trials in the ATTEND R condition.

A linear decoding framework is not readily applicable to features that depend on phase values. Indeed, phase is a circular variable, and it cannot be used directly in a thresholded weighted sum of the kind that linear classifiers rely on. We hence needed to come up with a way to linearize the phase locking information for each feature (Figure 73):

- i) the average phase difference between the stimulus and the EEG (over the entire training set, but not taking into account the test set) was computed in the most phase locked condition, when the stimulus was attended $\rightarrow |\phi|_{att}$
- ii) the difference between the phase difference for each example and this average phase difference was computed, and the cosine was taken, so as to yield a distance measure $\rightarrow \cos(\phi_i - |\phi|_{att})$. This feature takes value 1 when the phase difference in the example (ϕ_i) is exactly equal to the average phase difference in the attended condition ($|\phi|_{att}$), and -1 when it is diametrically opposite.

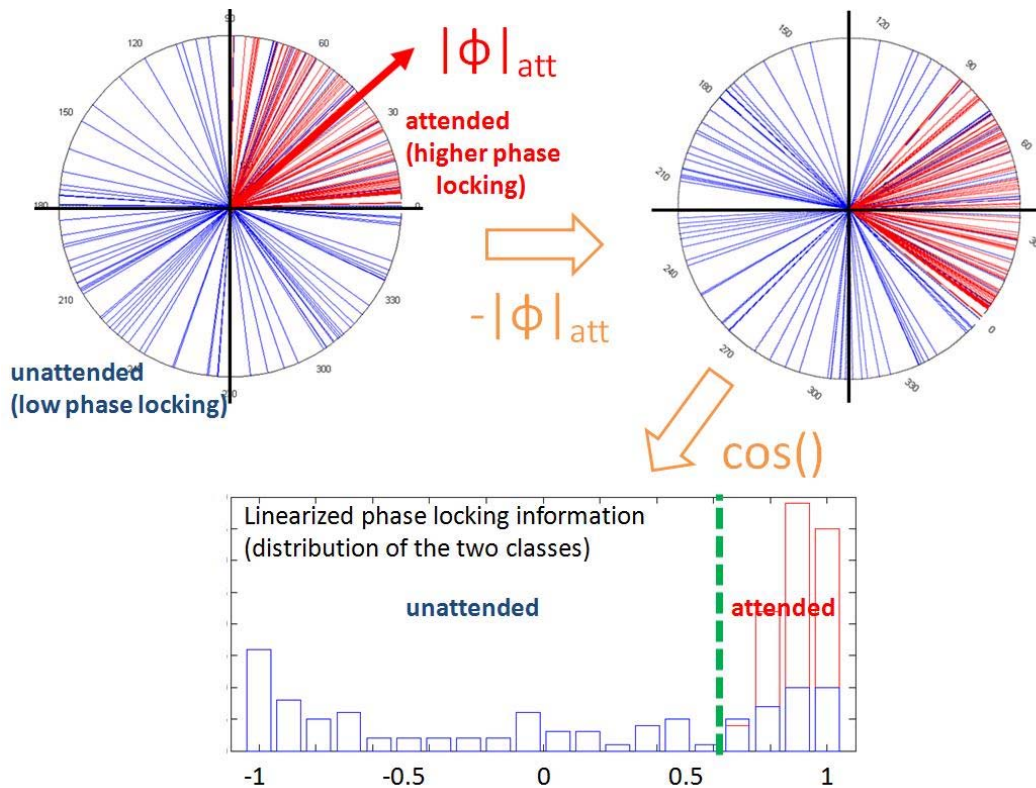


Figure 73 Linearizing the information available in the phase locking of the EEG and stimuli timecourses. This is SIMULATED data. The blue vectors correspond to the phase difference between the EEG and the left (right) stimulus, in time points when it is ignored (all time points in ATTEND R (L) blocks) at a given electrode, frequency and delay; the red vectors correspond to the timepoints when it is attended (i.e. all time points in ATTEND L (R) blocks). Refer to the text for explanations.

There were in total about 240 (trials) \times 225 (time points) = 54000 examples per subject. The time points were taken every 20ms from 1s to 5.5s of each trial. Since each subject had performed 3 blocks, we divided each dataset into training and

test sets accordingly to perform 9 cross-validations for each subject (Figure 74). We used a Support Vector Machine (SVM) classifier with a linear kernel (liblinear implementation by Chih-Jen Lin). We implemented feature selection with a filter approach using information theoretic measures (Maximum Relevance Minimum Redundance algorithm (Peng, Long, & C. Ding, 2005)) and enforced the use of 300 features; we performed grid search for the C parameter of the SVM classifier, choosing the best parameter based on four cross-validations within the training set.




Figure 74 Data for each subject was divided to perform 9 cross-validations. The testing set is never used for anything else but testing the final classifier, built from information in the training set alone.

We found that we could decode attentional allocation at each time point above chance for all subjects (Figure 75a), using the linearized phase locking information. To gauge how our classifier performed on classifying whole trials, we summed the decision values outputted by the classifier for each time point in a given trial and used the sum as a decision value for trial-wise classification. We got a performance well above chance for all included subjects (Figure 75b).

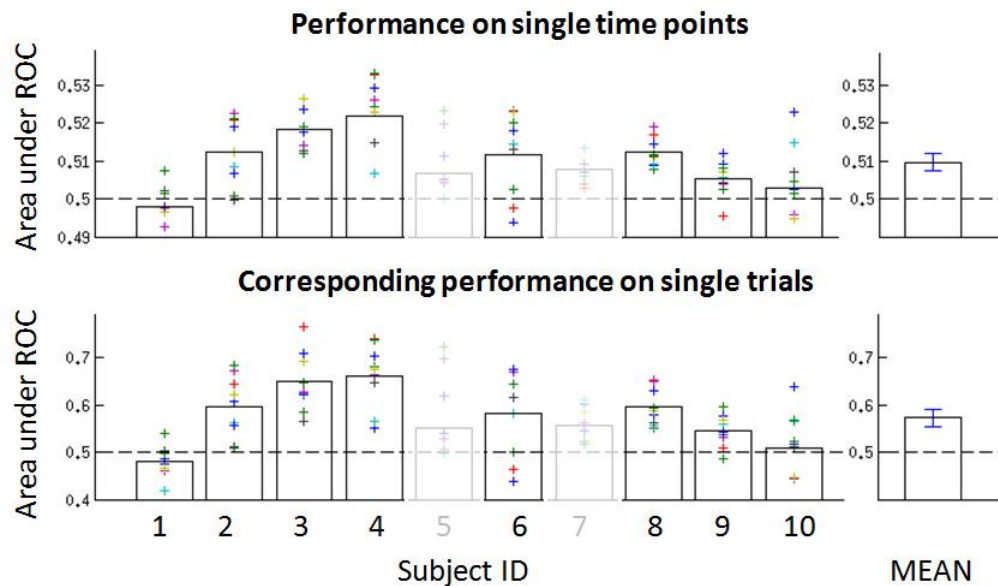


Figure 75 Performance of our linear classifier in distinguishing ATTEND L from ATTEND R time points (top). The performance for each subject is shown, and the colored crosses correspond to each cross-validation. The average performance is shown to the right. Summing the decision values in each trial, we get the corresponding single trial performance (bottom).

While the performance on single time points was low (about 51.5% on average) it was reliable since the accumulation of these predictions over the course of one trial (6.25s long, but only using 1s to 5.5s to avoid transients at stimulus onset and offset) allowed us to predict the location of attention with 62% accuracy on average (a reasonable level of performance with single-trial EEG data).

f) What is the temporal resolution of our technique?

While we used time points every 20ms as examples, it is obvious that we do not have 20ms time resolution with this technique – indeed, the time support of our wavelets (Figure 66) is likely to induce correlations between neighboring time points. To gauge the time resolution of our technique, we looked at the autocorrelogram of the decisions values outputted by the classifier in the test set (Figure 76). We collapsed both sides of the autocorrelogram, and determined at what time the autocorrelation exceeded the (mean + 5 standard deviations) of the segment between 0.2s and 0.5s, defined as the baseline. We found this time to be about 100ms – hence, our decision values were significantly correlated in a 100ms window. This limited the frequencies that we could safely look at in further power spectrum analyses.

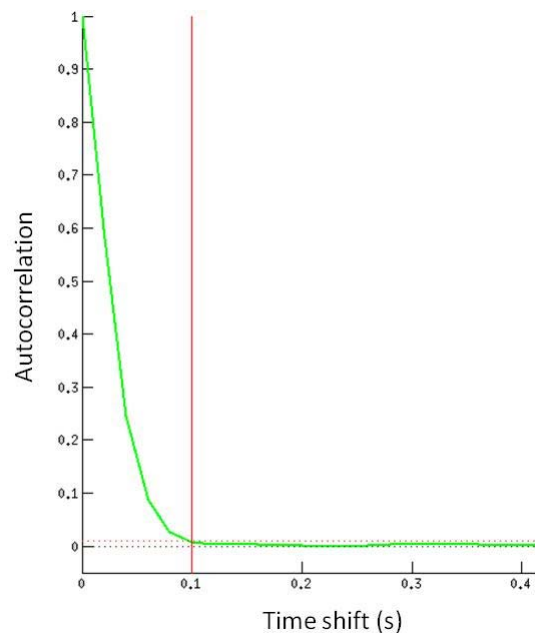


Figure 76 Autocorrelation, computed in each trial in the test set (ATTEND L and R), of the decision values outputted by the classifier. Our technique has a resolution better than 100ms.

g) The ~7Hz blinking spotlight of attention oscillates between simultaneously attended locations

Figure 77 shows (carefully chosen) representative 1-second-long segments of the decision values that we get in the three conditions. The classifier outputs negative values for ATTEND L and positive values for ATTEND R.

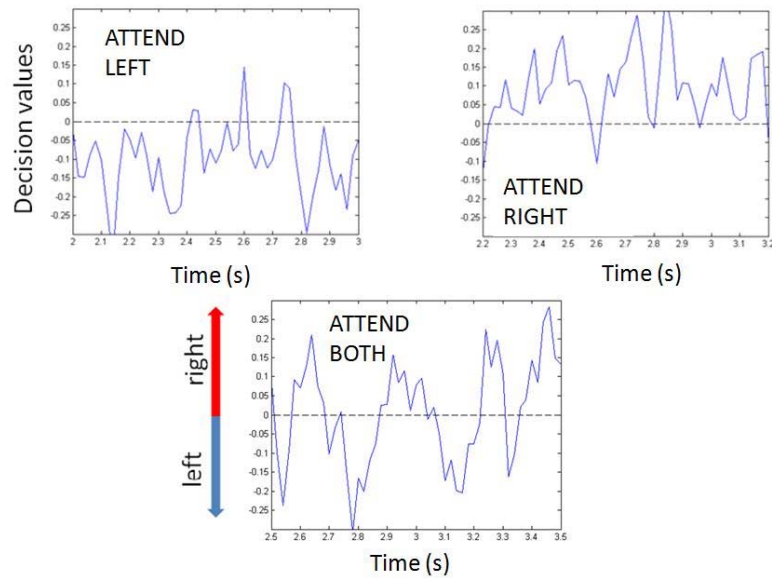


Figure 77 Representative time courses (1 second segments) of the decision values output by the classifier in the ATTEND L, ATTEND R and ATTEND B conditions. Negative decision values correspond to the prediction ATTEND L, while positive values correspond to the prediction ATTEND R.

Does the prediction oscillate in the ATTEND B condition? We can look at the autocorrelation of decision values in the ATTEND B condition : it gets negative (though not significantly) at around 150ms, indicating that the other location is being sampled at this time; importantly, at the same time shift, the autocorrelogram in the ATTEND 1 condition is positive.




Figure 78 T-values for the autocorrelation of decision values in the ATTEND B and ATTEND 1 conditions. Decision values are negatively correlated at a time shift of about 150ms in the ATTEND B condition, but not in the ATTEND 1 condition.

Encouraged by this result, we also conducted a power spectrum analysis of decision values in the ATTEND B condition, and compared it to the power spectrum of decision values in the ATTEND 1 (=ATTEND L or ATTEND R) condition. Figure 79

shows the difference between ATTEND B and ATTEND 1 in dB. There is a peak in the difference around 3.5-4.5Hz (not significant). This means that there is an oscillation in the decision values in the ATTEND B condition which is not present in the ATTEND 1 condition, of frequency ~ 4 Hz. We interpreted this as meaning that the spotlight of attention takes a sample at each location in turn (the decision values oscillate between positive and negative values from one sample to the next), coming back to the same location every 250ms or so.

Another note of interest is that there is a negative peak in the spectral difference, around 6.5 to 7Hz. Though it is not significant, it is possible that this is a signature of the blinking spotlight of attention, which takes samples at 7Hz when a single location is attended.

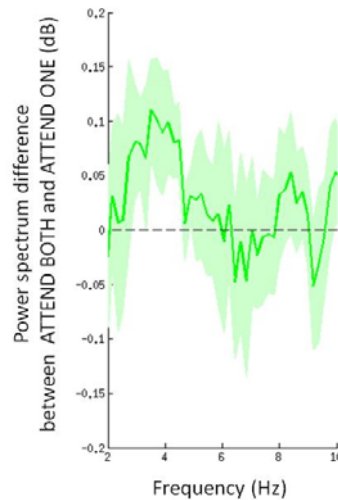


Figure 79 The power spectrum of decision values in the ATTEND B condition has more power in the 3.5-4.5Hz band than the power spectrum of decision values in the ATTEND 1 condition.

The effects that we evidenced are quite weak – the spectral difference is of the order of a few percent. This was expected from the very modest performance of our classifier on single time points. Another potential difficulty is that attention may not be switching in an all-or-none manner – the modulations may be much smaller : processes are seldom all-or-none in biology... The switching of the spotlight could in fact mean that we oscillated between a 40%-60% and a 60%-40% distribution of attention. Our result remains extremely encouraging. It may be possible to improve classification accuracy (while keeping a good temporal resolution), but collecting more subjects and sticking to the current method would be another way to improve the statistics and make sure that this is a true effect.

4. Summary and discussion

At the meeting of the Vision Sciences Society held in Naples from May 7th to May 12th 2011, Ayelet Landau, a postdoc in Pascal Fries's lab, gave a talk of high relevance to our endeavor. The study is unpublished, but I describe here what I remember of it. They ask the same question as we do, with a very similar display : when attention is directed to two locations simultaneously, does the spotlight of attention truly shine on both locations simultaneously or does it oscillate

between the two locations? Two drifting gratings, viewed through circular apertures, are presented on each side of a fixation point. The task of the subject is to detect a very brief ($\sim 10\text{ms}$) contrast decrement, which may occur in any of the two disks (or not at all). The contrast decrement is always in the same location, near the top of the disk, and it is presented at individual detection threshold as determined by a staircase procedure during a training period. Critically, at some point in each trial, 4 salient dots briefly flash around one of the two disks. The subjects are told to ignore this event, and indeed it does not help them with task performance – it is not predictive of target location. However, it does attract attention, involuntarily. This is the trick that Ayelet uses to reset the purported attentional rhythm: she knows where attention is at this exact time point. With the contrast decrement occurring between $\sim 200\text{ms}$ before and $\sim 600\text{ms}$ after the resetting event, in small steps ($\sim 10\text{ms}$), she can probe detection performance and plot its time course relative to the reset. Figure 80a shows a very schematic representation of the experiment.

The results show a detection performance which oscillates, with a period of about 300ms , at each of the two locations. The salient dots initially have a masking effect which hinders detection at the location where they appear, but the masking effect vanishes and detection performance oscillates thereafter. Critically, the oscillations of detection performance at the two locations have the same period (300ms) but are out of phase: it appears as if when performance is enhanced at one location, it is hindered at the other one, and vice versa. These results suggest that attention to both locations is in fact a serial process, sampling in turn the two locations...

This experiment appears to demonstrate exactly what we showed here: that attention to two spatial locations is in fact a serial process, enhancing processing at each location in turn, and that the switching occurs every 150ms or so. This corresponds to attentional sampling at about 7Hz (each location is focused on at every other sample in this situation, hence each location receives processing enhancement at about 3.5Hz).

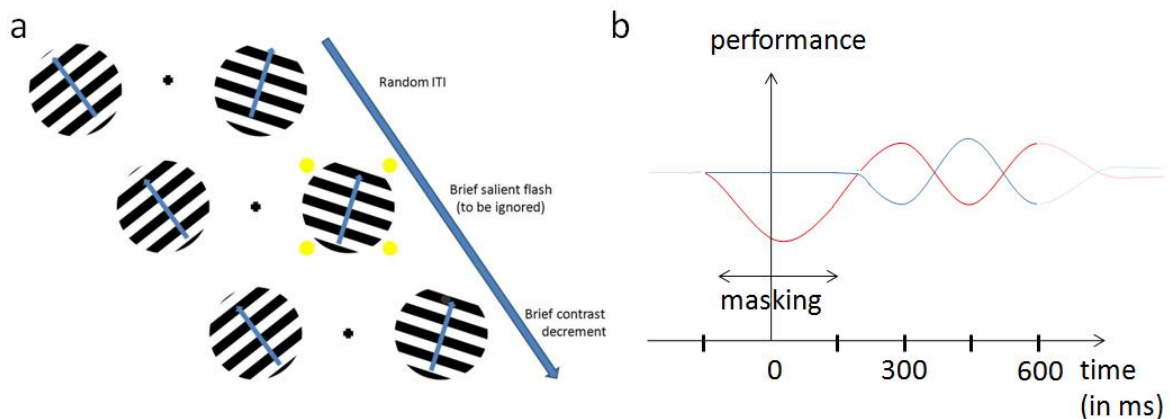


Figure 80 **a** Schematic representation of the display used by Landau and Fries. The orientation and spatial frequency of the gratings may not be the ones actually used in the experiment, but they are not of major interest (these stimuli were chosen as good gamma drivers, in preparation for a follow up in MEG). The brief contrast decrement appears near the top of the right disk in this case, and after the resetting event; it can also happen beforehand in some trials (baseline, negative times in b). **b** Schematic results obtained by Landau and Fries. After a dip in performance due to masking at the location where the salient flashes appeared (time 0), performance oscillates with a period of $\sim 300\text{ms}$. Interestingly, performance at the opposite location also oscillates, with the same period, but with an opposite phase. This is

indicative of an oscillation of the spotlight of attention, taking a sample every 150ms or so and alternating between the two attended locations.

There is one possible criticism though: we have been assuming that attention has an intrinsic rhythm, which is reset by the salient flashes, and can then be seen for a couple of cycles by averaging across trials. While this interpretation is the one that we favor, following the hypothesis that this thesis relies on, it could be that there is no intrinsic attentional rhythm – before the salient flashes, attention may not have been oscillating at all. The salient event was a necessary (and clever) trick to realign trials and evidence an oscillation by averaging. However, some may argue that the oscillation that we see after the salient event is transient and artificial, caused by the salient event, rather than reflecting an intrinsic attentional oscillation. While attention may have been truly simultaneously allocated to both locations prior to the salient event, it causes attention to focus on the side where the flashes appear, and it takes some time for the imbalance to be compensated for – during the return to a stable split attention configuration, attention can be seen to oscillate between the two locations. In a way, we could be seeing a ripple in the pond caused by the rock we just threw, rather than an intrinsic, orderly wave which we would have managed to reset... This criticism is valid, and the experiment presented by Landau and Fries can not address it.

Our approach does not suffer from the same confound, as we do not seek to interfere with attentional allocation. We really seek to observe an intrinsic rhythm. Our success was mitigated by a low signal-to-noise ratio, but we believe that our approach, if slightly improved on in terms of SNR, can potentially close the debate of serial VS parallel attentional allocation in the divided spatial attention literature. There is of course a possible criticism to our design : that our target (1 second long contrast decrement) lasted too long and thus encouraged a serial sampling strategy over a parallel, continuous monitoring strategy. This criticism can easily be fixed by having short targets, though. The reason why we chose a longer target duration was to be able to have a smaller contrast decrement (for the same performance level), hence less chances to reset attention by a salient transient. Another point which may be raised is that the attended locations were in separate hemifields; it is believed that the two hemispheres may have some independence regarding attentional allocation. If anything, this goes against us – having independent resources would favor a parallel strategy, hence a finding of serial behavior would be all the stronger and more convincing with this arrangement. This may mean that, to increase our chances for a significant result, we should consider a paradigm where the two locations are in the same hemifield.

E. PAPER 5 : VanRullen, Busch, Drewes and Dubois (2011) – Frontiers in Psychology



Ongoing EEG phase as a trial-by-trial predictor of perceptual and attentional variability

R. VanRullen^{1,2*}, N. A. Busch^{3,4}, J. Drewes^{1,2} and Julien Dubois^{1,2,5}

¹ Université de Toulouse, Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, Toulouse, France

² Centre National de la Recherche Scientifique, UMR 5549, Faculté de Médecine de Rangueil, Toulouse, France

³ Institute of Medical Psychology, Charité–Universitätsmedizin Berlin, Berlin, Germany

⁴ Berlin School of Mind and Brain, Humboldt Universität, Berlin, Germany

⁵ Division of Biology, California Institute of Technology, Pasadena, CA, USA

Edited by:

Guillaume A. Rousselet, University of Glasgow, UK

Reviewed by:

Peter Lakatos, Hungarian Academy of Sciences, Hungary

Gregor Thut, University of Glasgow, UK

*Correspondence:

R. VanRullen, Université de Toulouse, Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, 31062 Toulouse, France.
e-mail: rufin.vanrullen@cerco.ups-tlse.fr

Even in well-controlled laboratory environments, apparently identical repetitions of an experimental trial can give rise to highly variable perceptual outcomes and behavioral responses. This variability is generally discarded as a reflection of intrinsic noise in neuronal systems. However, part of this variability may be accounted for by trial-by-trial fluctuations of the phase of ongoing oscillations at the moment of stimulus presentation. For example, the phase of an electro-encephalogram (EEG) oscillation reflecting the rapid waxing and waning of sustained attention can predict the perception of a subsequent visual stimulus at threshold. Similar ongoing periodicities account for a portion of the trial-by-trial variability of visual reaction times. We review the available experimental evidence linking ongoing EEG phase to perceptual and attentional variability, and the corresponding methodology. We propose future tests of this relation, and discuss the theoretical implications for understanding the neuronal dynamics of sensory perception.

Keywords: EEG, oscillation, phase, pre-stimulus, spontaneous, ongoing, perception, attention

INTRODUCTION

Run a computer program twice with the same inputs: chances are, you should get the same output twice. As any experimenter knows, it is not so with the human brain. This unreliable device persistently fails to provide a consistent outcome: reaction times (RTs) vary by a factor of two or more, perception sometimes gets distorted and sometimes does not occur at all – even though the external world has been carefully controlled and equated, trial after trial. This variability gets in the way of any serious scientific measurement, and therefore scientists have dubbed it “noise” and found ways to discount it, generally by considering the mean response over several hundreds of trials as the true standard of brain function. Oftentimes, however, one comes across a signal in the brain that tells a lot about the subject’s perception on a given trial, or that can explain hitherto unexplained differences between individual trials. Such signals are the focus of the Special Topic to which this article belongs. More specifically, in this review we will consider situations in which the *phase* of ongoing brain oscillations (i.e., whether the oscillation is currently at its peak, its trough, or any particular point in between), even before any stimulus is actually presented to the subject, can inform us about their subsequent perception. Beyond the obvious implication that the brain has little to do with modern computers, these recent findings reveal much about its processing strategies.

OSCILLATORY PHASE INFLUENCES NEURAL RESPONSES

Neurons in the brain communicate by sending electrical pulses or “spikes,” which create electric potential differences at synapses and cell bodies. In turn, these voltage differences are responsible for the opening and closing of membrane channels and the subsequent

flow of electrically charged ions in and out of the cell bodies. Altogether, these processes induce large variations in local and long-distance electrical voltages at different temporal scales that can be considered as signatures of neuronal communication. These signatures are picked up by experimenters, for example using depth electrodes recording the extra-cellular potential and the local field potential (LFP), or using surface electrodes recording the electro-encephalogram (EEG).

The net effect of spikes and synaptic transmission at the level of neuronal populations often takes the form of an oscillation of the electric potential, in which the extra-cellular voltage increases and decreases at regular intervals. The responsiveness of single neurons to the same input intensity (i.e., the same number of spikes received) can vary greatly depending on the neurons’ present state (i.e., their membrane potential) as well as on whether the extra-cellular voltage oscillation is in its lower or higher stage. This influence of spontaneous oscillatory phase on neuronal processing has long been recognized *in vitro* (Calvin and Stevens, 1967; Levitan et al., 1968; Stern et al., 1997) but it is only recently that the potential effects of oscillatory phase on sensory processing have started to be investigated *in vivo*. For example, Fries et al. (2001) reported that the phase of pre-stimulus gamma (40–70 Hz) oscillations in cat visual cortex determined the latency of subsequent neuronal firing. Montemurro et al. (2008) found that the precise phase of an ongoing delta (1–4 Hz) oscillation at which neurons in primary visual cortex fired carried information about the visual stimulus that could not be extracted based on firing rate alone. In fact, the firing phase within each gamma oscillatory cycle is a reliable indicator of neuronal activation (Vinck et al., 2010a). These results support previously published theories proposing that the phase of spike

firing relative to an ongoing oscillatory signal could constitute a meaningful neural coding scheme (VanRullen et al., 2005a; Fries et al., 2007).

The same relation that exists between oscillatory phase at the moment a neuron receives its inputs, and this neuron's responsiveness to those inputs, can also be observed over larger-scale neuronal populations comprising entire brain areas. Indeed, past studies have also reported an influence of pre-stimulus EEG phase on the magnitude of various subsequent event-related potential (ERP) components – which represent a sensory system's response to its visual or auditory inputs (Jansen and Brandt, 1991; Brandt, 1997; Kruglikov and Schiff, 2003). Because neuronal firing ultimately generates subjective perception, and because ERPs are often regarded as external markers of this perception, the literature reviewed so far seems to point, albeit indirectly, to a possible relation between ongoing oscillatory phase and sensory perception. The direct measurement of this relation will be the topic of this review.

Our focusing on oscillatory phase does not imply, of course, that the *amplitude* of ongoing oscillations has no impact on perception. For one thing, the phase of an oscillatory signal can only be reliably computed when this signal has significant power. This is not only true in a mathematical sense, but also at the biophysical level: if membrane potential fluctuations were not synchronous over a reasonably large population of neurons, any influence of phase existing for individual neurons would average out at the population level. Furthermore, it is well accepted now that oscillatory power in various frequency bands bears significant relations to sensory perception and attention (Klimesch, 1999; Tallon-Baudry and Bertrand, 1999; Engel and Singer, 2001; Varela et al., 2001; Ergenoglu et al., 2004; Hanslmayr et al., 2005, 2007; Thut et al., 2006; van Dijk et al., 2008). Our motivation for concentrating on ongoing oscillatory phase is, simply, that this variable has been largely overlooked, at least until recent years. Similarly, we will restrict this review to cases of truly spontaneous oscillatory activity, even though numerous recent studies in human and non-human primates have reported an entrainment of the phase of brain oscillations to rhythmic stimulus presentation sequences, accompanied by periodic modulations of perception, attention, and RTs (Large and Jones, 1999; Lakatos et al., 2008; Schroeder and Lakatos, 2009; Mathewson et al., 2010).

Measuring the influence of ongoing oscillatory phase on perception cannot be performed using conventional methods, but requires single-trial analyses, which – one way or another – relate the variability of spontaneous brain signals to the changes in response variables across trials – rather than discarding this variability by averaging across trials or computing correlations across subjects. Before surveying the recent experimental advances in this area, we shall detail a few of the data analysis methods that have been used to uncover this relation.

METHODS FOR LINKING ONGOING EEG PHASE TO TRIAL-BY-TRIAL VARIABILITY

Our general problem can be defined in the following terms (see **Figure 1**). On each trial, an experimenter records both a temporal signal (e.g., an EEG waveform) and a behavioral response that can be graded (e.g., RT, perceived intensity on a continuous scale) or have only a few discrete labels (e.g., stimulus perceived vs. not

perceived, two- or multiple-alternative discrimination forced-choice; see **Figure 1A**). In certain situations, the graded behavioral response can be turned into a discrete variable by binning neighboring values (for example, labeling each RT with the corresponding quintile value, from 1 to 5). For our purposes, the temporal signal will correspond to brain activity *prior* to the presentation of the stimulus that must be perceived or classified by the observer (of course, similar methods can also be used for analysis of stimulus-evoked brain activity, but these will not be discussed here). In addition, we will assume that the experimental paradigm is designed using randomized inter-trial intervals, such that the moment of stimulus onset is unpredictable, and therefore the distribution of oscillatory phase values at or before stimulus onset is uniform across all trials. In technical terms, our main question is whether this phase distribution will significantly depart from uniformity, once the behavioral outcome is taken into account.

In practice, for a discrete behavioral variable, trials are grouped according to the behavioral response, and the uniformity of the distribution of phases is evaluated for each trial group. This last step can be done explicitly, by averaging across trials in the complex domain – with phase being represented by the angle of the complex vector (see **Figure 1C**); this was the approach used to compute results in **Figures 2A,C**. It can also be performed implicitly, in the temporal domain (see **Figure 1B**); for example, Fries et al. (2001) compared pre-stimulus LFP averages for groups of trials separated by short vs. long firing latencies; similarly, Mathewson et al. (2009) compared band-passed pre-stimulus EEG averages for perceived vs. unperceived visual stimuli. Finally, for a continuous behavioral variable, specific methods exist (Berens, 2009) that estimate the correlation between the pre-stimulus EEG phase (a circular variable) and the behavioral response (generally given on a linear scale); this was the approach used for the analysis illustrated in **Figure 2B**. There are, of course, other alternatives to measure phase dependency, but the methods listed here already cover most of those used in the existing literature.

Obviously, each of the approaches listed above also needs to be accompanied by appropriate statistics. Simple parametric tests are sometimes sufficient: for example, time-domain signals can be directly compared between two groups of trials corresponding to two distinct behavioral outcomes, using a Student's *t*-test – the null hypothesis being that the pre-stimulus means for the two conditions are equivalent at each point in time. Care must be taken, however, to correct for the increased likelihood of false positives due to the number of multiple comparisons (in the above example, each time point yields a distinct, but not necessarily independent, statistical comparison). Such correction methods (Bonferroni's, among others) are beyond the scope of this article. In addition, circular variables (e.g., phase angles expressed in the complex domain) are highly non-linear (for example, the mean of two angles of 10° and 350° does not correspond to the arithmetic mean of 180; instead, the circular mean of these two vectors is 0° – or 360°). Therefore, it is often preferable to devise non-parametric statistical tests based on permutation or bootstrapping methods (Vinck et al., 2010a,b). For example, phase-locking values measured across trials grouped by behavioral outcome cannot be directly compared with a null hypothesis of zero phase locking, because the null hypothesis actually depends on the exact number of trials in each




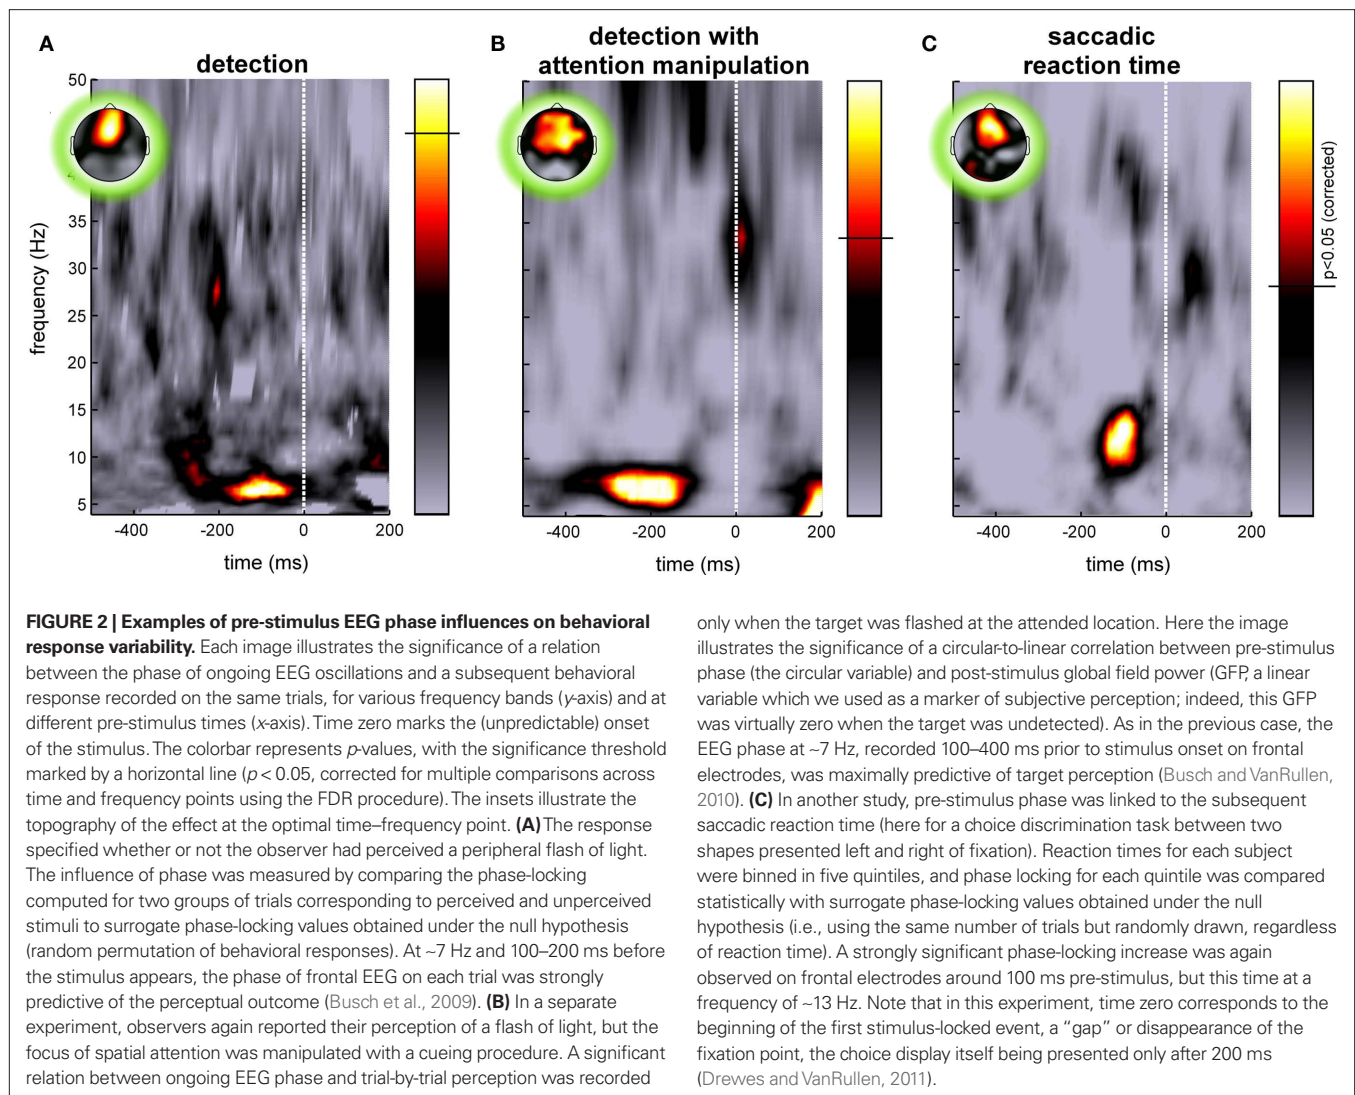
FIGURE 1 | Methods for linking pre-stimulus EEG phase to trial-by-trial variations of a discrete behavioral response. (A) On each trial, a time-varying signal (such as an EEG waveform) is recorded before stimulus onset. Here the signal is represented as oscillatory, but it could also correspond to a band-pass filtered version of a noisy, non-oscillatory signal. For each trial, the observer provides a behavioral response with discrete outcomes [here two possible outcomes (A,B)] corresponding, e.g., to distinct percepts (seen vs. unseen, category (A vs. B)), or to quantized reaction times (e.g., slow vs. fast), etc. On the right, the time-domain signal recorded on each trial is expressed as a vector in the complex plane (which can be calculated using Fourier or wavelet time–frequency decomposition methods). The direction of the vector and its length represent the phase and the amplitude (respectively) of the oscillatory signal recorded in a given time window of interest. **(B)** Since the moment of stimulus onset is unpredictable, the pre-stimulus mean over all trials is near-zero. However, if certain oscillatory phases systematically induce one or the other behavioral outcomes, then selectively averaging the signals for trials grouped according to the behavioral response should reveal oscillations with distinct phase angles. Here the

magnitude of the resulting oscillation appears constant at all pre-stimulus times but in practice, due to external and measurement noise, the phase difference is more likely to be visible just before stimulus onset. This method essentially corresponds to an ERP computed before, rather than after the stimulus onset. **(C)** Another method consists in selectively averaging the vectors for each trial group in the complex domain; prior to averaging, each vector is normalized to a unit length, implying that its phase will always equally contribute to the average, regardless of its amplitude [without this normalizing step, the method would actually give equivalent results to the one described in **(B)**]. This step is important because amplitude modulations that would occur independently of phase effects would tend to obscure them. The length of the resulting vector after averaging is called “phase-locking value” (PLV) or “inter-trial coherence” (ITC; Lachaux et al., 1999). The phase locking is weak when the distribution of phase angles across trials is uniform (as should occur when all trials are pooled together regardless of behavioral outcome); if certain phase angles systematically induce one specific behavioral outcome, on the other hand, this phase locking should be significantly increased by considering only the trials with this outcome.

group (even for a uniform phase distribution, the expected phase locking in any *finite* group of trials is significantly above zero; see **Figure 1C**). A solution is to randomly reassign each trial to one of the behavior-defined groups, keeping the respective number of trials constant, and then re-calculate phase locking for this surrogate dataset; repeating this operation several times provides a distribution of phase-locking values under the null hypothesis, with which the real phase-locking value can be compared to estimate its statistical significance. The same approach can be applied to a situation with a continuous behavioral variable, by shuffling the assignment of behavioral values (e.g., RTs) to the corresponding EEG signals, each time re-calculating the circular test statistic under the null hypothesis (e.g., circular-to-linear correlation). In our experiments (**Figure 2**), we have favored such permutation methods because of their robustness and relative lack of assumptions about the data structure.

ONGOING EEG PHASE PREDICTS PERCEPTUAL VARIABILITY

Recent studies by our group and others have started exploring the impact of the phase of ongoing pre-stimulus EEG oscillations on the subsequent perception of a visual stimulus. In a first study (Busch et al., 2009), we presented brief (6 ms) and dim peripheral flashes of light to our observers ($n = 12$), with the luminance of the flash adjusted individually so that the exact same stimulus would be perceived on approximately half of the trials, but go completely unnoticed on the other half. We computed pre-stimulus phase-locking separately for the two trial groups corresponding to perceived and unperceived flashes, and found for each group a significant increase (compared to phase-locking computed on the same number of trials but drawn randomly, irrespective of perceptual outcome). This increase occurred just before stimulus onset, at a frequency of ~ 7 Hz (**Figure 2A**), and the effect was maximal over fronto-central electrodes. In fact, by considering the phase of the 7-Hz band-pass



filtered EEG recorded at those electrodes just before stimulus onset in each trial, we could predict the subsequent response of the subject well above chance. Up to 16% of the trial-by-trial differences in perception were accounted for by comparing trials having the optimal phase angle with those at the opposite angle (Busch et al., 2009).

These findings were globally consistent with the conclusions of an independent study by Mathewson et al. (2009), who also reported that the phase of low-frequency oscillations (around 10 Hz) just before stimulus onset predicted trial-by-trial perception, in a situation where only half of the targets were consciously detected. However, important differences also exist between the two studies. The paradigm differed from ours, first, in that conscious visibility was regulated not by using dim stimuli, but by displaying a strong-contrast stimulus (a “mask”) shortly after the target. In addition, the inter-trial interval duration was fixed in that experiment; the possibility that certain oscillatory rhythms could have been reset by the stimulus onset in the previous trial thus makes it difficult to draw strong conclusions in terms of ongoing or spontaneous oscillations. Finally, the analysis method relied on time-domain averaging (as

illustrated in **Figure 1A**) which renders phase effects dependent on potential oscillatory amplitude differences between the perceptual conditions; such pre-stimulus amplitude differences between correctly and incorrectly perceived trials are known to exist, particularly over occipital regions at alpha-band frequencies around 10 Hz (Ergenoglu et al., 2004; Hanslmayr et al., 2005, 2007; Thut et al., 2006; van Dijk et al., 2008). This may also contribute to explain why the principal phase effect was observed by Mathewson et al. (2009) at 10 Hz on occipital electrodes, instead of 7 Hz on frontal electrodes in our study. Nonetheless, the fact that both studies point to a similar conclusion reinforces the general idea that pre-stimulus oscillatory phase at 7 and/or 10 Hz can determine to some extent the trial-by-trial changes in our conscious perception of a repetitive event.

ONGOING EEG PHASE REFLECTS PERIODIC ATTENTIONAL SAMPLING

In our next study, we asked whether the influence of ongoing phase on perception was mediated by top-down attentional factors (Busch and VanRullen, 2010). The previous results had been obtained

under conditions in which target location was always known in advance, and therefore subjects may have paid covert attention to that location in order to improve their detection performance. Would the same ongoing phase influence still occur for a target appearing at an unattended location?

Before each trial began, a central cue indicated to the observers ($n = 13$) the location on the screen where they should expect the target (left or right). When the flash of light did occur on this side, everything happened in fact exactly as in the previous experiment – and indeed we confirmed our previous results in this condition, with a strong impact of ~ 7 Hz pre-stimulus EEG phase recorded at fronto-central electrodes on the probability of target perception (**Figure 2B**). However, the target also sometimes occurred on the opposite side of the screen. In this case the subjects had more difficulty in perceiving the light (as indicated by the higher light intensity that proved necessary to achieve a 50% detection rate in this condition); this confirms that the observers were focusing on the cued side at the detriment of the rest of the screen. Critically, for those trials where the target appeared outside the focus of attention, no significant phase-locking effect was recorded for the perceived or unperceived trials. In other words, perception was related to ongoing EEG phase only via the action of attention. We thus hypothesized that attention samples visual information periodically, and that each ~ 7 Hz ongoing EEG cycle is the reflection of a new attentional sample (Busch and VanRullen, 2010). Stimuli occurring at around the optimal phase enjoy all the benefits of attention, while others are processed merely as if they were out of the attention focus. For some reason, likely related to its architecture and its neuronal substrates, the attention system could not apply the optimal strategy (optimal for such a detection task with unpredictable target onset) of steadily monitoring the expected location. These findings concur with conclusions from a previous psychophysical study in which we reported that attention samples information periodically at ~ 7 Hz, even when only a single item needs to be attended (VanRullen et al., 2007). In this context, the topographic localization of the phase effects over fronto-central electrodes may reveal the contribution of the frontal eye field (FEF), an area known, among other things, for its involvement in visual attention (Crowne, 1983; Kodaka et al., 1997; Corbetta and Shulman, 1998; Schall, 2004; Wardak et al., 2006). However, our EEG data would require independent corroboration using more accurate anatomical localization methods, before the implication of FEF can be definitely established.

ONGOING EEG PHASE PREDICTS REACTION TIME VARIABILITY

Attention and saccade programming are heavily intertwined brain functions (Rizzolatti et al., 1987; Deubel and Schneider, 1996; Smith et al., 2004). In particular, the FEF is involved in both visual attention orienting and saccadic motor outputs (Moore and Fallah, 2001; Murthy et al., 2001; Juan et al., 2004; Schall, 2004; Wardak et al., 2006). Our next experiment thus tested whether saccadic responses would also be affected by ongoing pre-stimulus phase (Drewes and VanRullen, 2011).

We used three different versions of a saccadic response task, performed by the same 13 observers. All three tasks required speeded choice responses using leftward or rightward eye movements, but

the difficulty of the choice varied. In the easiest case, subjects simply made alternating left and right saccades toward a target whose position was fully predictable, trial after trial. In the second task the position was unpredictable, but since only one target appeared on each trial the task could be performed using mostly reflexive responses. Finally, the third task required a discrimination between two shapes displayed simultaneously; the saccade was made toward the shape that presented an opening at the top. In all three tasks, the fixation point disappeared 200 ms before the appearance of the saccade target display; this so-called “gap” procedure is normally used to maximize the occurrence of rapid “express” saccades (Fischer and Boch, 1983; Fischer and Ramsperger, 1984). The disappearance of the fixation point is itself a transient event that can modify or even reset ongoing oscillations; therefore, we considered the beginning of the gap as time zero, and concentrated only on spontaneous oscillations occurring before this time.

Behavioral results revealed that, as expected, the mean discrimination performance decreased while the average RT increased with task difficulty. This time, we focused on RT variability across trials, and how this variability would relate to ongoing EEG phase differences. In order to apply the analysis methods described previously (see **Figure 1C**), we binned the RTs of each observer into five groups, corresponding to the five quintiles of the RT distribution (this was done separately for each of the three tasks). For each trial group, we calculated the pre-stimulus phase locking (at each frequency and each pre-stimulus time point) and compared it to surrogate phase-locking values calculated using the same number of trials which were randomly drawn, irrespective of RT. Again, a significant increase in pre-stimulus phase locking was found over frontal electrodes for all five quintiles; this time, however, the effect was maximal at a frequency around 13 Hz. This phase effect was stronger in the easy and in the medium difficulty tasks; in fact, during the difficult task, the effect was only observed for subjects who provided very rapid (but often inaccurate) responses (**Figure 2C**). Thus the influence of pre-stimulus phase on RT variability seems to depend on cognitive load and subject strategy; for tasks involving a considerable cognitive effort, many other factors (discrimination and decision processes, motivation) may come into play and contribute their own variability, which would act to conceal the effects of pre-stimulus phase.

The eventuality of a trial-by-trial relationship between pre-stimulus oscillatory phase and subsequent RTs had already been investigated in the past, but these early studies provided only mixed results (Walsh, 1952; Lansing, 1957; Callaway and Yeager, 1960; Dustman and Beck, 1965). One specificity of our experiment that could explain its comparative success, is that our analysis did not assume a one-to-one relationship between phases and RTs – contrary to previous studies that all searched for specific phase angles systematically inducing the fastest, or the slowest RTs. As we found out, the relation between ongoing EEG rhythms and subsequent RTs can actually span more than a single oscillatory cycle. Each range of RT values will be associated with a specific and unique phase angle (as our analysis revealed), but the reverse is not true, that is, a given range of phases might be linked to two, or even more distinct RT values. This sort of relation would be missed by an analysis that would first group the trials by phase, and then compare the RTs in each group (a strategy employed in most of the studies cited above).

It would also be missed by directly calculating the circular-to-linear correlation between phase angles and RTs. A proper analysis for such a situation is the one described in **Figure 1**, in which trials are first grouped according to RTs and phase values are then considered in each group – in other words, the very analysis that we used (**Figure 2C**). The existence of a phase–RT relationship spanning multiple oscillatory cycles suggests that the underlying ~13 Hz ongoing oscillation creates multiple successive and regularly spaced “windows of opportunity” for saccade production.

WHAT ELSE DOES ONGOING EEG PHASE PREDICT?

The evidence reviewed so far unambiguously indicates that ongoing ~7 and ~13 Hz EEG oscillations dynamically modulate information processing in the visual system, and in particular the sampling of visual information by attention. Therefore, the single-trial phase recorded just before stimulus onset can be used as a predictor for subsequent behavioral and perceptual variables. Aside from conscious visual detection and RTs, ongoing EEG phase could also contribute to the trial-by-trial variability of several other cognitive functions. Some of the possible associations are listed below, together with a preview of the significance that their discovery could have for our understanding of brain function. Needless to say, our group is currently exploring several of these issues.

- *Does ongoing EEG phase predict transcranial magnetic stimulation (TMS) phosphene perception at threshold?* The perception of a flash of light is the result of a complex sequence of neuronal processing events, from the retina to the cortex via the thalamus. It is unclear (and heavily debated) which cortical region, or which network of cortical areas is critical for conscious perception to occur. Within this context, our results of a rhythmic influence of ongoing oscillations onto the conscious detection of a flash are difficult to interpret. However, the conscious experience of light can also be induced by direct stimulation of the occipital cortex (the seat of the visual system), for example using TMS pulses. Just as in our experiments, the intensity of the TMS pulse can be individually adjusted so that the perception of the induced “phosphene” only takes place in half of the trials, and it is possible to record EEG while applying TMS pulses (Thut et al., 2005; Taylor et al., 2008; Thut and Miniussi, 2009). Furthermore, focusing on pre-stimulation oscillations means that the study would be immune to the numerous artifacts generally evoked by the pulse in concurrent TMS/EEG studies (Ilmoniemi and Kicic, 2010; Thut and Pascual-Leone, 2010). Would the perceptual outcome in this experiment also depend on ongoing 7 Hz frontal EEG phase? If these oscillations are the reflection of attentional sampling, and assuming that visual attention enhances phosphene perception (Bestmann et al., 2007), the answer is likely to be positive. In addition, would the perceptual outcome on each trial also be affected by the phase of locally generated oscillations within occipital cortex itself – and if so, at what frequency? The fact that TMS-induced perception bypasses many of the early visual processing stages should maximize the chances of directly observing the local interplay between ongoing activity and visual responses, which so far has eluded our previous experimental efforts.
- *Does ongoing EEG phase predict threshold perception in other sensory modalities (e.g., audition)?* It is easy to adjust auditory, or even somato-sensory stimuli so that they are consciously registered only half of the time. Would a pre-stimulus phase difference be observed between perceived and unperceived stimuli? If yes, are the same or different mechanisms involved as in the visual modality, in terms of oscillation frequency, and cortical origin? The results of this experiment may indicate whether the source of the periodic modulation is central, or more sensory-specific. Existing evidence in animals already hints at a supramodal coordination of theta-band oscillatory activities (Lakatos et al., 2009).
- *Does ongoing EEG phase predict the latency of attentional shifts?* Just like saccadic RTs are variable when you move your eyes, it also takes variable amounts of time to shift your attention covertly from one location to another. This variability can be measured using specific paradigms, e.g., using a running analog “clock” at the target location and asking the observer to report the first “time” they can read on the clock after an attention orienting event. This time is generally delayed with respect to the onset of the orienting cue, and the delay is taken to reflect the latency of shifting attention to the target location (Carlson et al., 2006). Even for identical repetitions of an experimental trial, the latency is not fixed but varies by tens of milliseconds. Can the method employed to study the latency of saccadic responses (often taken to represent shifts of “overt” attention) also be used to reveal a relation between ongoing EEG phase and covert attentional shifts? In this case, would the relevant oscillation frequency be found at ~7 Hz (like in our study of attentional sampling; Busch and VanRullen, 2010) or at ~13 Hz (like in our study of saccadic latency; Drewes and VanRullen, 2011)? In any case, a positive outcome would lend credence to our proposal that ongoing ~7 and/or ~13 Hz oscillations mirror the rhythm at which attention samples visual information (VanRullen et al., 2005b, 2006, 2007; Busch and VanRullen, 2010).
- *Does ongoing EEG phase predict the capacity and/or ordering of items in visual short-term memory?* Visual attention and visual working memory share several traits, such as their limited capacity (Luck and Vogel, 1997; Cowan, 2001; Alvarez and Cavanagh, 2004), and it has even been proposed that the two cognitive functions overlap in part (LaBar et al., 1999; Awh et al., 2000, 2006; Downing, 2000; Awh and Jonides, 2001; de Fockert et al., 2001), although this conclusion is debated (Woodman et al., 2001, 2007). An influential model of working memory organization posits that remembered items are maintained in memory as an ordered sequence (Sternberg, 1966), each element being represented by one cycle of a high-frequency oscillation (e.g., 30–80 Hz gamma activity) nested within a lower-frequency cycle (e.g., 4–8 Hz theta activity) supposed to encompass the entire list (Lisman and Idiart, 1995). The limited capacity of working memory (about four to seven items) is explained in this model by the number of high-frequency cycles that can be slotted in one period of the low-frequency oscillation. A recent study of neural responses in monkey prefrontal cortex reported that spikes fired at distinct phases of an ongoing ~32 Hz oscillatory rhythm carried

information about distinct objects from a to-be-remembered list (Siegel et al., 2009). Although the study was presented as evidence for Lisman's model, its findings depart significantly from Lisman's predictions (indeed, the phase dependence should occur at the lower rhythm frequency, not at the higher one). Nonetheless, the findings clearly suggest that pre-stimulus oscillations could relate to short-term memory performance. Remaining questions include (i) whether the phase of ongoing lower-frequency oscillations (e.g., in the theta band) at which a test item is presented would predict the trial-by-trial variability in response time for deciding whether or not the item belongs to the remembered list – this prediction follows naturally from Lisman's model since the phase encodes the item's rank in the list and the rank determines the time needed for retrieval; (ii) whether inter-individual or inter-trial differences in peak oscillation frequency would correlate with capacity and performance measures; (iii) whether interfering with ongoing oscillations at specific phases (for example, using TMS applied on frontal areas) would disrupt working memory maintenance only for specific items within the list.

- *Does ongoing EEG phase predict long-term memory encoding and/or recall?* Memory encoding over longer time scales (minutes, days, or even more) depends both on frontal structures and on the medial temporal lobe system, including the hippocampus (Poldrack and Gabrieli, 1997; Desgranges et al., 1998; Kramer et al., 2005; Ramus et al., 2007). This latter area displays very large amplitude oscillations in the theta band (4–8 Hz) which are known to underlie spatial memory formation in rodents (O'Keefe, 1993; Buzsaki, 2002, 2006). More precisely, the firing of certain hippocampal neurons signals a remembered location in the rat's environment, and the specific phase of the ongoing theta rhythm at which this firing occurs reflects the relative position of the rat with respect to this location – a mechanism coined “theta phase precession” (O'Keefe and Recce, 1993; Skaggs et al., 1996). In fact, theta phase precession also coordinates the firing of prefrontal neurons to the hippocampal theta rhythm (Jones and Wilson, 2005; Siapas et al., 2005). A recent study using single-neuron recordings in humans revealed that trials in which long-term memory formation was successful were characterized by stronger phase locking of hippocampal neurons to the ongoing theta rhythm, even before the onset of the stimulus to be recalled (Rutishauser et al., 2010). This neuronal result does not directly imply the existence of a relationship between the phase of ongoing theta oscillations at the moment of presentation of a visual stimulus and the subsequent recall of this stimulus, but it makes such a relationship worth testing in future experiments.

DO ONGOING OSCILLATIONS PRODUCE PERCEPTUAL SNAPSHOTS?

One critical test of the relation between ongoing rhythms and conscious perception has eluded the community for so long that it deserves a dedicated section in this review. In 1981, inspired by earlier theories claiming that our perceptual experience was built upon discrete processing events, similar to the discrete frames of a movie sequence (Pitts and McCulloch, 1947; Stroud, 1956; Harter, 1967; Allport, 1968), Varela et al. (1981) reported apparently direct evidence for such perceptual “frames” or “snapshots.” By presenting identical stimuli (two successive flashes separated by a short delay) at different phases of the ongoing alpha (10 Hz) EEG, they induced drastic changes in the observers' experience of temporal simultaneity: while at one alpha phase they judged the two flashes to have occurred at the same time, at the opposite phase they perceived the *same* two flashes as clearly separate events. The implication is that each alpha cycle slices the continuous temporal sequence of visual inputs into a new discrete chunk or snapshot – when the two flashes straddle the critical phase of the cycle, they are sliced into separate snapshots. Unfortunately, these promising first results did not prove as clear-cut in a follow-up study by the same author (Gho and Varela, 1988), and could simply not be replicated afterward, despite repeated efforts by our group and others (D. Eagleman, personal communication).

This failure is fateful: whereas the evidence reviewed so far of a relation between ongoing oscillatory phase and trial-by-trial variations in conscious detection, attention, or motor outputs implies the existence of periodic components in perception, it only indirectly alludes to the issue of discrete vs. continuous perception. A relation between ongoing phase and temporal framing, on the other hand, would directly, and unambiguously demonstrate the discrete nature of perception (VanRullen and Koch, 2003). Until such a demonstration is provided, the ongoing debate must continue to rely on indirect experimental signatures of the postulated discreteness, such as the motion reversals occurring in continuous light during the “wagon wheel illusion” (Purves et al., 1996; Kline et al., 2004; Andrews and Purves, 2005; Andrews et al., 2005; Holcombe et al., 2005; VanRullen et al., 2005b, 2006; Kline and Eagleman, 2008). In the end, even if it turns out that discrete temporal framing does not occur after all, or that it is restricted to specific sensory domains or experimental situations, the studies reviewed in this article should hopefully convince the reader that the outcome of many important brain functions depends in a periodic manner on the ongoing state of the brain, as reflected by the phase of certain pre-stimulus oscillations; and further, that it is possible to reveal this dependence using careful analysis of single-trial EEG activity.

REFERENCES

- Allport, D. A. (1968). Phenomenal simultaneity and the perceptual moment hypothesis. *Br. J. Psychol.* 59, 395–406.
- Alvarez, G. A., and Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychol. Sci.* 15, 106–111.
- Andrews, T., and Purves, D. (2005). The wagon-wheel illusion in continuous light. *Trends Cogn. Sci. (Regul. Ed.)* 9, 261–263.
- Andrews, T., Purves, D., Simpson, W. A., and VanRullen, R. (2005). The wheel keeps turning: reply to Holcombe et al. *Trends Cogn. Sci. (Regul. Ed.)* 9, 561.
- Awh, E., Anllo-Vento, L., and Hillyard, S. A. (2000). The role of spatial selective attention in working memory for locations: evidence from event-related potentials. *J. Cogn. Neurosci.* 12, 840–847.
- Awh, E., and Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends Cogn. Sci. (Regul. Ed.)* 5, 119–126.
- Awh, E., Vogel, E. K., and Oh, S. H. (2006). Interactions between attention and working memory. *Neuroscience* 139, 201–208.
- Berens, P. (2009). CircStat: a MATLAB Toolbox for circular statistics. *J. Stat. Softw.* 31, 1–21.

- Bestmann, S., Ruff, C. C., Blakemore, C., Driver, J., and Thilo, K. V. (2007). Spatial attention changes excitability of human visual cortex to direct stimulation. *Curr. Biol.* 17, 134–139.
- Brandt, M. E. (1997). Visual and auditory evoked phase resetting of the alpha EEG. *Int. J. Psychophysiol.* 26, 285–298.
- Busch, N. A., Dubois, J., and VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *J. Neurosci.* 29, 7869–7876.
- Busch, N. A., and VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proc. Natl. Acad. Sci. U.S.A.* 107, 16048–16053.
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron* 33, 325–340.
- Buzsáki, G. (2006). *Rhythms of the Brain*. New York: Oxford University Press.
- Callaway, E. I., and Yeager, C. L. (1960). Relationship between reaction time and electroencephalographic alpha phase. *Science* 132, 1765–1766.
- Calvin, W. H., and Stevens, C. F. (1967). Synaptic noise as a source of variability in the interval between action potentials. *Science* 155, 842–844.
- Carlson, T. A., Hogendoorn, H., and Verstraten, F. A. (2006). The speed of visual attention: what time is it? *J. Vis.* 6, 1406–1411.
- Corbetta, M., and Shulman, G. L. (1998). Human cortical mechanisms of visual attention during orienting and search. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 353, 1353–1362.
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav. Brain Sci.* 24, 87–114.
- Crowne, D. P. (1983). The frontal eye field and attention. *Psychol. Bull.* 93, 232–260.
- de Fockert, J. W., Rees, G., Frith, C. D., and Lavie, N. (2001). The role of working memory in visual selective attention. *Science* 291, 1803–1806.
- Desgranges, B., Baron, J. C., and Eustache, F. (1998). The functional neuroanatomy of episodic memory: the role of the frontal lobes, the hippocampal formation, and other areas. *Neuroimage* 8, 198–213.
- Deubel, H., and Schneider, W. X. (1996). Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res.* 36, 1827–1837.
- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychol. Sci.* 11, 467–473.
- Drewes, J., and VanRullen, R. (2011). This is the rhythm of your eyes: the phase of ongoing EEG oscillations modulates saccadic reaction time. *J. Neurosci.* 31, 4698–4708.
- Dustman, R. E., and Beck, E. C. (1965). Phase of alpha brain waves, reaction time and visually evoked potentials. *Electroencephalogr. Clin. Neurophysiol.* 18, 433–440.
- Engel, A. K., and Singer, W. (2001). Temporal binding and the neural correlates of sensory awareness. *Trends Cogn. Sci. (Regul. Ed.)* 5, 16–25.
- Ergenoglu, T., Demiralp, T., Bayraktaroglu, Z., Ergen, M., Beydagi, H., and Uresin, Y. (2004). Alpha rhythm of the EEG modulates visual detection performance in humans. *Brain Res. Cogn. Brain Res.* 20, 376–383.
- Fischer, B., and Boch, R. (1983). Saccadic eye movements after extremely short reaction times in the monkey. *Brain Res.* 260, 21–26.
- Fischer, B., and Ramsperger, E. (1984). Human express saccades: extremely short reaction times of goal directed eye movements. *Exp. Brain Res.* 57, 191–195.
- Fries, P., Neuenschwander, S., Engel, A. K., Goebel, R., and Singer, W. (2001). Rapid feature selective neuronal synchronization through correlated latency shifting. *Nat. Neurosci.* 4, 194–200.
- Fries, P., Nikolic, D., and Singer, W. (2007). The gamma cycle. *Trends Neurosci.* 30, 309–316.
- Gho, M., and Varela, F. J. (1988). A quantitative assessment of the dependency of the visual temporal frame upon the cortical rhythm. *J. Physiol.* 83, 95–101.
- Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., and Bauml, K. H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *Neuroimage* 37, 1465–1473.
- Hanslmayr, S., Klimesch, W., Sauseng, P., Gruber, W., Doppelmayr, M., Freunberger, R., and Pecherstorfer, T. (2005). Visual discrimination performance is related to decreased alpha amplitude but increased phase locking. *Neurosci. Lett.* 375, 64–68.
- Harter, M. R. (1967). Excitability cycles and cortical scanning: a review of two hypotheses of central intermittency in perception. *Psychol. Bull.* 68, 47–58.
- Holcombe, A. O., Clifford, C. W., Eagleman, D. M., and Pakarian, P. (2005). Illusory motion reversal in tune with motion detectors. *Trends Cogn. Sci. (Regul. Ed.)* 9, 559–560.
- Ilmoniemi, R. J., and Kicic, D. (2010). Methodology for combined TMS and EEG. *Brain Topogr.* 22, 233–248.
- Jansen, B. H., and Brandt, M. E. (1991). The effect of the phase of prestimulus alpha activity on the averaged visual evoked response. *Electroencephalogr. Clin. Neurophysiol.* 80, 241–250.
- Jones, M. W., and Wilson, M. A. (2005). Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. *Hippocampus* 15, 867–873.
- Juan, C. H., Shorter-Jacobi, S. M., and Schall, J. D. (2004). Dissociation of spatial attention and saccade preparation. *Proc. Natl. Acad. Sci. U.S.A.* 101, 15541–15544.
- Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Res. Brain Res. Rev.* 29, 169–195.
- Kline, K., and Eagleman, D. M. (2008). Evidence against the temporal subsampling account of illusory motion reversal. *J. Vis.* 8, 13.1–13.5.
- Kline, K., Holcombe, A. O., and Eagleman, D. M. (2004). Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field. *Vision Res.* 44, 2653–2658.
- Kodaka, Y., Mikami, A., and Kubota, K. (1997). Neuronal activity in the frontal eye field of the monkey is modulated while attention is focused on to a stimulus in the peripheral visual field, irrespective of eye movement. *Neurosci. Res.* 28, 291–298.
- Kramer, J. H., Rosen, H. J., Du, A. T., Schuff, N., Hollnagel, C., Weiner, M. W., Miller, B. L., and Delis, D. C. (2005). Dissociations in hippocampal and frontal contributions to episodic memory performance. *Neuropsychology* 19, 799–805.
- Kruglikov, S. Y., and Schiff, S. J. (2003). Interplay of electroencephalogram phase and auditory-evoked neural activity. *J. Neurosci.* 23, 10122–10127.
- LaBar, K. S., Gitelman, D. R., Parrish, T. B., and Mesulam, M. (1999). Neuroanatomic overlap of working memory and spatial attention networks: a functional MRI comparison within subjects. *Neuroimage* 10, 695–704.
- Lachaux, J. P., Rodriguez, E., Martinerie, J., and Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Hum. Brain Mapp.* 8, 194–208.
- Lakatos, P., Karmos, G., Mehta, A. D., Ulbert, I., and Schroeder, C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320, 110–113.
- Lakatos, P., O'Connell, M. N., Barczak, A., Mills, A., Javitt, D. C., and Schroeder, C. E. (2009). The leading sense: supramodal control of neurophysiological context by attention. *Neuron* 64, 419–430.
- Lansing, R. W. (1957). Relation of brain and tremor rhythms to visual reaction time. *Electroencephalogr. Clin. Neurophysiol.* 9, 497–504.
- Large, E. W., and Jones, M. R. (1999). The dynamics of attending: how people track time-varying events. *Psychol. Rev.* 106, 119–159.
- Levitan, H., Segundo, J. P., Moore, G. P., and Perkel, D. H. (1968). Statistical analysis of membrane potential fluctuations. Relation with presynaptic spike train. *Biophys. J.* 8, 1256–1274.
- Lisman, J. E., and Idiart, M. A. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267, 1512–1515.
- Luck, S. J., and Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature* 390, 279–281.
- Mathewson, K. E., Fabiani, M., Gratton, G., Beck, D. M., and Lleras, A. (2010). Rescuing stimuli from invisibility: inducing a momentary release from visual masking with pre-target entrainment. *Cognition* 115, 186–191.
- Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., and Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *J. Neurosci.* 29, 2725–2732.
- Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr. Biol.* 18, 375–380.
- Moore, T., and Fallah, M. (2001). Control of eye movements and spatial attention. *Proc. Natl. Acad. Sci. U.S.A.* 98, 1273–1276.
- Murthy, A., Thompson, K. G., and Schall, J. D. (2001). Dynamic dissociation of visual selection from saccade programming in frontal eye field. *J. Neurophysiol.* 86, 2634–2637.
- O'Keefe, J. (1993). Hippocampus, theta, and spatial memory. *Curr. Opin. Neurobiol.* 3, 917–924.
- O'Keefe, J., and Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330.
- Pitts, W., and McCulloch, W. S. (1947). How we know universals: the perception of auditory and visual forms. *Bull. Math. Biophys.* 9, 127–147.
- Poldrack, R. A., and Gabrieli, J. D. (1997). Functional anatomy of long-term memory. *J. Clin. Neurophysiol.* 14, 294–310.
- Purves, D., Paydarfar, J. A., and Andrews, T. J. (1996). The wagon wheel illusion in movies and reality. *Proc. Natl. Acad. Sci. U.S.A.* 93, 3693–3697.
- Ramus, S. J., Davis, J. B., Donahue, R. J., Disenza, C. B., and Waite, A. A. (2007). Interactions between the orbitofrontal cortex and the hippocampal memory

- system during the storage of long-term memory. *Ann. N. Y. Acad. Sci.* 1121, 216–231.
- Rizzolatti, G., Riggio, L., Dascola, I., and Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* 25, 31–40.
- Rutishauser, U., Ross, I. B., Mamelak, A. N., and Schuman, E. M. (2010). Human memory strength is predicted by theta-frequency phase-locking of single neurons. *Nature* 464, 903–907.
- Schall, J. D. (2004). On the role of frontal eye field in guiding attention and saccades. *Vision Res.* 44, 1453–1467.
- Schroeder, C. E., and Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci.* 32, 9–18.
- Siapas, A. G., Lubenov, E. V., and Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron* 46, 141–151.
- Siegel, M., Warden, M. R., and Miller, E. K. (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proc. Natl. Acad. Sci. U.S.A.* 106, 21341–21346.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., and Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172.
- Smith, D. T., Rorden, C., and Jackson, S. R. (2004). Exogenous orienting of attention depends upon the ability to execute eye movements. *Curr. Biol.* 14, 792–795.
- Stern, E. A., Kincaid, A. E., and Wilson, C. J. (1997). Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *J. Neurophysiol.* 77, 1697–1715.
- Sternberg, S. (1966). High-speed scanning in human memory. *Science* 153, 652–654.
- Stroud, J. M. (1956). “The fine structure of psychological time,” in *Information Theory in Psychology*, ed. H. Quastler (Chicago, IL: Free Press) 174–205.
- Tallon-Baudry, C., and Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends Cogn. Sci. (Regul. Ed.)* 3, 151–162.
- Taylor, P. C., Walsh, V., and Eimer, M. (2008). Combining TMS and EEG to study cognitive function and cortico-cortico interactions. *Behav. Brain Res.* 191, 141–147.
- Thut, G., Ives, J. R., Kampmann, F., Pastor, M. A., and Pascual-Leone, A. (2005). A new device and protocol for combining TMS and online recordings of EEG and evoked potentials. *J. Neurosci. Methods* 141, 207–217.
- Thut, G., and Miniussi, C. (2009). New insights into rhythmic brain activity from TMS-EEG studies. *Trends Cogn. Sci. (Regul. Ed.)* 13, 182–189.
- Thut, G., Nietzel, A., Brandt, S. A., and Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *J. Neurosci.* 26, 9494–9502.
- Thut, G., and Pascual-Leone, A. (2010). Integrating TMS with EEG: how and what for? *Brain Topogr.* 22, 215–218.
- van Dijk, H., Schoffelen, J. M., Oostenveld, R., and Jensen, O. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *J. Neurosci.* 28, 1816–1823.
- VanRullen, R., Carlson, T., and Cavanagh, P. (2007). The blinking spotlight of attention. *Proc. Natl. Acad. Sci. U.S.A.* 104, 19204–19209.
- VanRullen, R., Guyonneau, R., and Thorpe, S. J. (2005a). Spike times make sense. *Trends Neurosci.* 28, 1–4.
- VanRullen, R., Reddy, L., and Koch, C. (2005b). Attention-driven discrete sampling of motion perception. *Proc. Natl. Acad. Sci. U.S.A.* 102, 5291–5296.
- VanRullen, R., and Koch, C. (2003). Is perception discrete or continuous? *Trends Cogn. Sci. (Regul. Ed.)* 7, 207–213.
- VanRullen, R., Reddy, L., and Koch, C. (2006). The continuous wagon wheel illusion is associated with changes in electroencephalogram power at approximately 13 Hz. *J. Neurosci.* 26, 502–507.
- Varela, F., Lachaux, J. P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nat. Rev. Neurosci.* 2, 229–239.
- Varela, F. J., Toro, A., John, E. R., and Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia* 19, 675–686.
- Vinck, M., Lima, B., Womelsdorf, T., Oostenveld, R., Singer, W., Neuenschwander, S., and Fries, P. (2010a). Gamma-phase shifting in awake monkey visual cortex. *J. Neurosci.* 30, 1250–1257.
- Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C. M. (2010b). The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *Neuroimage* 51, 112–122.
- Walsh, E. G. (1952). Visual reaction time and the alpha-rhythm, an investigation of a scanning hypothesis. *J. Physiol.* 118, 500–508.
- Wardak, C., Ibos, G., Duhamel, J. R., and Olivier, E. (2006). Contribution of the monkey frontal eye field to covert visual attention. *J. Neurosci.* 26, 4228–4235.
- Woodman, G. F., Luck, S. J., and Schall, J. D. (2007). The role of working memory representations in the control of attention. *Cereb. Cortex* 17(Suppl. 1), i118–i124.
- Woodman, G. F., Vogel, E. K., and Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychol. Sci.* 12, 219–224.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 12 January 2011; paper pending published: 26 January 2011; accepted: 24 March 2011; published online: 09 April 2011.

Citation: VanRullen R, Busch NA, Drewes J and Dubois J (2011) Ongoing EEG phase as a trial-by-trial predictor of perceptual and attentional variability. *Front. Psychology* 2:60. doi: 10.3389/fpsyg.2011.00060
This article was submitted to *Frontiers in Perception Science*, a specialty of *Frontiers in Psychology*.

Copyright © 2011 VanRullen, Busch, Drewes and Dubois. This is an open-access article subject to a non-exclusive license between the authors and Frontiers Media SA, which permits use, distribution and reproduction in other forums, provided the original authors and source are credited and other Frontiers conditions are complied with.

IV. Discussion

Four years ago, I started this thesis work guided by the intriguing possibility that visual perception may rely on discrete samples taken quasi-periodically from the incoming visual input. This is a view that differs from the way the community classically thinks about the visual brain. The prominent stance is that, as visual information comes in, it is treated in parallel and continuously by various feature and shape detectors whose output is fed, millisecond by millisecond, to conscious experience – as put by Alex Holcombe (Alex Holcombe, 2011). This “continuous perception” view is compatible with our subjective, seamless visual experience.

From the very start of this thesis manuscript, I warned the reader against introspection as a guide to the underpinnings of perception : visual perception is a constructive process, with many intermediate stages which we have no awareness of. Perceptual decisions are made silently in the face of ambiguous inputs (e.g. in the lab, binocular rivalry and ambiguous figures) and missing information (e.g. the blind spot, occlusions); processing artefacts (motion blur) are compensated for to give a crisper interpretation of the visual input.

The hypothesis that visual perception may be updated periodically has emerged many times in the psychological literature, with perhaps the most interest in the 50s-70s with Stroud’s formulation of the “perceptual moment theory” (Stroud, 1967). Arguably, most (but not all) of the evidence that was put forward in those years could be explained away with a simple integration window, by which visual perception would be updated millisecond by millisecond based on what happened in the previous 80-100ms. Recently, a paradigm has revived the largely forgotten discrete/continuous perception debate : the continuous Wagon Wheel Illusion. A continuously rotating spoked wheel watched under continuous lighting can be seen to reverse direction, intermittently. This observation has been interpreted by some as a temporal aliasing artefact, thus demonstrating that there is a sampling mechanism at work in the visual pathways for perception – and, critically, that attention is at the origin of sampling. We discussed this illusion at length in section II.B (page 34). It has so far failed to convince the community. It is evident that the brain is not “running a clock”. The visual brain is not a movie camera, taking periodic snapshots of the scene at a fixed frame rate. This caricatural view is often the one chosen by opponents of the discrete perception hypothesis, to discredit it. However, there is increasing evidence that perception does indeed rely on periodic processes. Our proposal is as follows: that attention may be a periodic process, relying on oscillatory brain activity, and that perception is shaped by this periodicity.

My thesis work was meant to add to this evidence, thereby urging the community to consider brain oscillations and their functional relevance to the computations underlying perception. In this discussion, I will first ponder on the difficulty of finding evidence of underlying periodicities behaviorally, which is the main reason why the community has generally been reluctant to the discrete perception idea, and is largely unaware of the relevance of brain rhythms to perception. Second, I consider the periodicities that we have found in the various experiments that I presented in this thesis, and other periodicities evidenced in our group and others. Is there truly a characteristic rhythm for attention (and the ensuing perception), or does the rhythm depend on specific task requirements? Is it possible to entrain or interfere with the rhythm of attention, and consequently alter task performance, for the better or for the worse? What is the relationship of covert attentional sampling with overt rhythmic sampling (eye movements, whisking, sniffing, ...)? Finally, we allow

ourselves to go slightly beyond the data and get more speculative: how does the 7Hz rhythm of covert attention fit with existing results, e.g. the recent progress on the role of the occipito-parietal alpha rhythm? We formulate a novel hypothesis to account for empirical observations, which will need to be put to the test of further experimental investigations...

A. Why are oscillatory mechanisms not obvious perceptually?

If oscillations are a key feature of neuronal communication, and have a central role in the computations underlying perception, why are they not obvious perceptually? It has proved really hard to find any strong evidence of underlying periodicities in behavioral experiments – temporal aliasing in the continuous Wagon Wheel Illusion paradigm remains, to date, the most telling evidence of an underlying rhythm (rather than an integration period), but this interpretation remains a subject of criticism. The difficulties that we experienced in extending the temporal aliasing paradigm (section II.1, page 49) seem to further drive the point that if there are oscillatory components to perceptual processes, the brain is extremely secretive about them. How does the brain protect itself from our scrutinizing attempts?

1. An absence of perception cannot be perceived

The interpretation of the cWWI is that periodic samples are taken of the moving stimulus, while little or no information is acquired in the time interval between two successive samples. If this interpretation is an accurate description of reality, it means that there are times at which we do not perceive as readily as at other times – this is consistent with the findings that we exposed in section III.1 (page 103) of the dependence of visual detection at threshold on the phase of an ongoing brain oscillation, i.e. on the precise timing of incoming information relative to endogenous brain activity. How are we not aware of these recurring periods in which we gather less information?

You may remember from the introduction how our eyes move around all the time, making fast saccades to foveate on different locations or objects, at a mean rate of about three per second. Perception is turned off during saccades. An easy way to convince yourself of this is to face a mirror and look at your right eye, then at your left eye, and go back and forth like this a few times. Do you think you'll see your eyes moving? You won't. Are your eyes moving so fast that you just can't see the motion? If you ask someone to stand in front of you and look at your right eye, then your left, and go back and forth a few times, you can clearly see their eyes moving, hence the eyes are not moving too fast. The reality is that visual perception is partially suppressed during eye movements. This so-called saccadic suppression has been extensively studied psychophysically (e.g. (D C Burr, M. C. Morrone, & Ross, 1994) and electrophysiologically (e.g. (Thiele, Henning, Kubischik, & Hoffmann, 2002)). Subjectively, we are totally oblivious to this phenomenon – there doesn't appear to be any gap in our perception as we make eye movements! This demonstrates that we do not perceive perceptual blanks at times when the brain is not accepting any incoming information. Either these intervals are simply cut from our subjective experience of time, or they are actively filled in by interpolation – some evidence of time compression associated with saccades may point to the first interpretation (M. Morrone, Ross, & David C Burr, 2005).

The point is that, if there are any time points when little or no information is acquired for perception, they are unlikely to be evident – the brain knows how to conceal them from our awareness.

2. Early areas buffer incoming visual information until the next sample

It is one thing not to perceive the perceptual blanks subjectively. However, if information is only acquired periodically for perception, a direct consequence is that we should miss short lived stimuli that occur during the perceptual blanks. Again, this is what our results for the visual detection task at threshold point to in section III.1 (page 103). But shouldn't it be more obvious, psychophysically? Do we need to be at threshold to see any effects? Can we really talk of "discrete perception" when our experiment shows a mere 16% detection probability modulation? Discrete perception does not necessarily mean that there are times at which we do not perceive at all – the duty cycle of perceptual cycles may well be 100%, with a mere fluctuating gain for perception, hitting a low at the border between perceptual epochs and a high in the middle of each perceptual epoch. Hence, there is no good reason to predict that we will entirely miss some events that fall at times when the gain is lowest – only events that elicit very little activity may be missed, as in our visual detection experiment at threshold. If the duty cycle was not 100%, a mechanism that can certainly cover the tracks of periodic sampling for perception is the buffering of incoming visual information by early visual cortices. It is likely that information flows in continuously from the retina to the LGN and to V1 and some higher level areas. Neuronal activity in early visual cortex in response to an impulse is not as short lived as the stimulus; for example, in the visual cortex of a cat, the temporal impulse response of a simple cell may last for more than 250 milliseconds (DeAngelis, Ohzawa, & Freeman, 1995)– clearly long enough to be detected in the next perceptual epoch in the event that the brief stimulus happened during a state of low perceptual gain. Therefore, unless the response elicited by a brief stimulus is very short lived (which may happen if the stimulus is very faint), it will not be missed by discrete perceptual sampling. This is possibly the main reason why it is so difficult to find evidence of discrete sampling – it is necessary to use very brief threshold stimuli in order for them to have a chance to be missed.

3. Motion perception mechanisms are the glue that links discrete epochs

If information is acquired for perception periodically, and if the world is stationary, there is no problem, as nothing will have changed dramatically from one perceptual epoch to the next. But if there are substantial changes from one acquisition to the next, there is a need to smoothly interpolate the successive frames to provide us with a continuous, seamless interpretation. In the absence of the mechanisms which link successive frames, our perception of the world would likely be jittery. In section II.3 (page 70), we argued that motion perception may be involved in glueing successive snapshots. Indeed, a complete disruption of motion perception leads to cinematographic vision or akinetopsia : the patient only sees one frame at a time, and her subjective visual experience is a jittery succession of static frames. We also described another, less dramatic but also more frequent disturbance of motion perception, which we referred to as visual trailing : patients perceive a series of slowly decaying positive afterimages in the wake of moving objects. We argued that this disturbance may be a case in which underlying perceptual epochs become consciously available, though other interpretations are also likely.

4. Resetting the rhythm may be the only way to find a behavioral effect

If you want to show behaviorally that there is a rhythmic modulation of the probability of detection of a target stimulus (the key prediction of the perceptual sampling hypothesis), you will run into the problem that you do not have control over the internal state of the subject on a given trial. We circumvented this problem by recording from the brain and identifying the rhythms that correlated with detection performance in section III.1 (page 103). But if you do not have access to brain activity, you need to be able to reset the ongoing sampling rhythm at each trial at will, so that you can realign different trials and see the effects of the internally generated rhythm after reset. This approach remains slightly problematic : are you really resetting an ongoing sampling rhythm, or are you generating a rhythm with your “resetting” event, as when you throw a stone in a pond? (see discussion in II.3)

We have a better understanding of why it is difficult to find behavioral evidence of perceptual rhythms. Many mechanisms are in place to hide their effects – and this is all for the best, as we wouldn’t want to have a flickery perception of our environment. Recently, our lab came across a new illusion: the center of a wagon wheel with a high radial spatial frequency (e.g., 32 spokes) is seen to flicker when you move your eyes (typically, while reading some text) as shown in Figure 81. This illusion may be linked to the alpha rhythm of the subject (Sokoliuk and VanRullen, in preparation). Sokoliuk and VanRullen suggested that this illusion is a case in which we perceive the oscillatory activity underlying perception, however the specific way in which the physical stimulus may cause this is not yet understood.




Figure 81 You should see the center of the wheel flicker as you read this.

Despite the inherent difficulties that we just described, some experiments have claimed to isolate a perceptual sampling rhythm (which we have equated with the rhythm of attention throughout this thesis). So, what is it?

B. The rhythm of covert attention

1. 7hz? 10hz?13hz?

We previously discussed a recent psychophysical model fitting study (Rufin VanRullen et al., 2007), in the context of a difficult contrast decrement detection task performed with divided attention to up to 4 locations, which concluded that the spotlight of attention was blinking at 7hz (even when a single location is attended to). Our EEG experiment in a visual detection at threshold paradigm (N.A. Busch et al., 2009) presented in this thesis (PAPER 4, page 104), and its follow up experiment (Niko A. Busch & Rufin VanRullen, 2010), both indicate that attentional signals recorded over fronto-central electrodes modulate visual detection threshold at about 7hz, consistent with the rhythm of the blinking spotlight. Our attempts to measure the switch rate of spatial attention between two locations in section III.3 (page 131) also suggested that the spotlight may take samples at a rate of about 7hz, which was independently confirmed by a psychophysical experiment by Landau and Fries presented at VSS 2011 and described in the discussion of section II.3 (page 147).

A Transcranial Magnetic Stimulation (TMS) study was just completed in our group (Dugue, Marque, & Rufin VanRullen, 2011), which addressed the same question that we asked in PAPER 4 (page 104), using the detection of a magnetically induced ‘visual phosphene’ instead of a visual flash. It also found an influence of the ongoing phase at fronto-central electrodes, albeit it was maximal at around 10hz – attention was not explicitly manipulated in that experiment, hence it is unclear that this is an attentional rhythm, but the most parcimonious interpretation would be that it corresponds to the same rhythm found in (N.A. Busch et al., 2009; Niko A. Busch & Rufin VanRullen, 2010). In this context, why is the rhythm found to be slightly different? does it have to do with intersubject variability, or with the task?

Finally, the interpretation of the continuous Wagon Wheel Illusion that we have favored until now, and which largely inspired this thesis work, suggests that attentional motion computation relies on samples taken at around 13hz (Rufin VanRullen, Reddy, et al., 2005) – an EEG correlate of the illusion, at 13hz, was evidenced over right parietal lobe electrodes. Interestingly, Drewes and VanRullen (Drewes & R. VanRullen, 2011) found an effect of the phase of a 13hz rhythm over fronto-central electrodes on saccadic latency across three different tasks, which they tentatively associated with an attentional rhythm. I’ll come back to this when discussing the relationship of overt rhythms with the covert rhythm.

Do all these findings reflect the same process? For simplicity we sometimes summarize our findings by talking about effects in the 7-13hz range. Such statements could be obscuring important details and giving other researchers reason to criticize our work. Dissecting the tasks that are performed in the various experiments may yield some insight on how to interpret the discrepant findings. My view is that evidence is accumulating in favor of about 7hz as the intrinsic rhythm of spatial attention – whether it is selecting a single location or multiple locations “concurrently”. The frontocentral rhythm (whose effect is maximal at about 10hz) evidenced in the recent TMS experiment in our group (Dugue et al., 2011) could be the same rhythm. The task is quite different though: the “stimulus” to be detected is internally generated rather than being an external flash appearing on the display. It is unclear that the neural activity elicited by a pulse of TMS has much in common with the orderly activity evoked by a flashed external stimulus (the temporal impulse response to a pulse of TMS is likely to be very different from the temporal impulse response to a brief flash, given that many stages of processing are bypassed and activity is elicited in the early visual cortex in a very unspecific manner), and in that regard it is unclear that spatial attention would function as it does in natural conditions. Furthermore, the spatial extent of the percept in the TMS

experiment is less well defined, and much larger, than in the case of the externally flashed stimulus. For all these reasons, I do not wish to overemphasize this experiment in our discussion of the rules underlying normal vision. The remaining issue is that of the 13hz rhythm suggested by the cWWI paradigm. It is clearly distinct from the 7hz rhythm that we have evidenced elsewhere – a 7hz rhythm would not explain the findings of the cWWI. Recently, our group tried to test a simple prediction in the context of multiple simultaneously attended wagon wheels: the prediction of attentional sampling is that when e.g. 3 wagon wheels are attended to, each wagon wheel will be sampled at roughly 1/3 of the attentional sampling frequency for a single wagon wheel; we thus expected a linear decrease in the tuning frequency of the cWWI as a function of the number of concurrently presented stimuli. The experiments were performed by Macdonald, Cavanagh and VanRullen (Macdonald, Patrick Cavanagh, & Rufin VanRullen, 2011). While they did find a decrease in the optimal frequency for perceptual reversals as a function of set size, the decrease was not linear: there was in fact only a trend for a decrease in the tuning frequency of the illusion when going from set sizes 2 to 3 and 3 to 4 (Figure 82). This complicates the initial hypothesis that attention would sample each of the attended wheels in turn, at its previously found 13hz rate. An irregular sampling strategy, in which samples from the concurrent wheels are not taken orderly but irregularly, may still account for the results. Another possibility is that the 13hz rhythm evidenced by the cWWI is in fact specific to (third-order) motion perception, and altogether distinct from the attentional rhythm that we evidence in other experiments. This may be the best explanation for the failure of our prediction. An interaction between the intrinsic rhythm of attentional allocation and the rhythm of attentional sampling for motion perception somehow leads to the effects that Macdonald and colleagues reported.




Figure 82 A nonlinear decrease of the optimal temporal frequency for illusory reversals when the number of simultaneously attended Wagon Wheels increases from 1 to 4. Reproduced from (Macdonald et al., 2011)

In conclusion, the evidence points to a rhythm of about 7hz for spatial attention. Is this consistent with other threads in the literature, which have estimated the speed and the processing time of attention in different contexts?

2. Other estimates of the attentional rhythm in the literature

a) The speed of attentional shifts

There exists a large literature which we have not mentioned yet, but which is highly relevant in the context of a discussion on the rhythm of attention : studies of the speed of attentional shifts. Several studies, using different methodologies, have attempted to determine the time required to shift attention from one location to another: these include for example attentional orienting studies – varying the cue-target asynchrony and measuring the effect on performance, to time attentional shifts (endogenous: (Charles W Eriksen & Collins, 1969; H. J. Müller & Rabbitt, 1989; G L Shulman et al., 1979); exogenous: (Egly & Homa, 1991; H. J. Müller & Rabbitt, 1989; Nakayama & Mackeben, 1989; Tsai, 1983)), and attentional gating studies – shifting attention from one RSVP stream to another upon detection of a predefined target and reporting the first letter in the new RSVP stream (endogenous: (Reeves & George Sperling, 1986; George Sperling & Weichselgartner, 1995)) or reporting the first item after a luminance transient in a single RSVP stream (exogenous: (Weichselgartner & George Sperling, 1987)). More recently, a clever paradigm was developed by Carlson and colleagues (Carlson, Hinze Hogendoorn, & F. a J. Verstraten, 2006) to get rid of some conceptual problems plaguing previous studies, such as the need to define a performance threshold in attention orienting studies, and attentional blink or time required for higher level object recognition in attention gating studies. The paradigm was recently revisited in our group (Chakravarthi & Rufin VanRullen, 2011) with the introduction of a subtraction method to avoid the confound of processing time and thus yield a pure measure of the time required for an attentional shift. In conclusion : endogenous attention takes a minimum of about 250ms to relocate, and the time depends critically on how far the new location is from the current one; whereas exogenous attention takes roughly 100ms and this shift is just as fast if the new location is far than if it is close by. The speed of exogenous attention is more likely to speak to an intrinsic rhythm of attention, independent of voluntary control. According to this line of evidence, the intrinsic attentional rhythm should be close to 10hz – an estimate which is not too discrepant with the 7hz that we described previously.

b) Serial search

Another paradigm that may come to mind in the debate about the rhythm of attention is serial visual search. Though the classical way to think about serial search was in terms of search slopes (time that attention spends on each item), there has been much debate about the validity of strictly serial models of visual search, hence I will not tread through these waters (Wolfe, 1998)... More to the point, direct neural evidence of serial shifts during visual search was reported in macaque monkeys in a study by Buschman and Miller (Buschman & Miller, 2009). The shifts were paced by an 18-34Hz oscillation (variable across trials), a clocking signal in FEF. Of course, one can not directly compare rhythms between humans and monkeys – in general connections are shorter and neural communication is faster in the monkey brain. Also, the switching of attention between items in an overtrained search task may not reflect the intrinsic rhythm of attention implemented in a sustained divided attention paradigm. Sadly, there is to date no study of FEF activity in a paradigm involving sustained attention to multiple locations – but such a study may come out very soon (in fact, Maxime Cauchoix, another PhD student at the CerCo, ran such an experiment while recording with an implanted electrode matrix; the analysis is pending...). Then

it will become clearer whether the rhythm of attention is clocked by the same process in the trained serial search and divided attention paradigms.

3. Task dependence

Our understanding of oscillatory attentional mechanisms is still in its infancy. It is foreseeable that the rhythm of attention could be modulated as a function of task requirements : some tasks do not necessitate as fine sampling as others, and the level of alertness may significantly modulate the attentional rhythm.

a) When time slows down

A common observation is often cited as an argument in favor of the task demand adaptability of the attentional rate : that, during a frightening event (typically, a car accident), time appears to slow down and we notice many more things within a short time span. An increase of the rhythm of attention, taking in more samples in a given time window, seems a very likely candidate to account for this observation. This possibility was directly addressed by a clever and ludic study in David Eagleman's group (Stetson, Fiesta, & Eagleman, 2007). They used a wrist mounted LED display to show two digits in alternation with their negative images – when the alternation rate increases, they report that there was a sharp threshold, beyond which subjects could not correctly identify the presented digits any more. A determination of the threshold by a simplified staircase procedure (starting at a slow alternation rate, and increasing it at each correct response, until three successive mistakes are made) was performed on the ground for each subject (the sequence is presented for 2s). Then, subjects were taken to the platform of a Suspended Catch Air Device tower – and performed a 46m free fall, which takes 2.49s. While free falling, they attempted to read the digits on their wrist mounted device, which was set to an alternation rate slightly faster than their individual threshold. The authors report that the subjects were not able to perform the task while free falling any better than on the ground – which they take as evidence that there is no increase in the perceptual/attentional sampling rate. This study is very interesting, however it could be argued that it does not offer bulletproof evidence against an increase of attentional sampling in some situations. There are many confounding factors plaguing this study, such as the length of time that the participants really looked at or attended to the display during the free fall (they may have closed their eyes partway, or just have been too distracted by bodily responses to fully concentrate on the task). It is my belief that the question should receive further investigation – perhaps once a definite oscillatory EEG signature of attention is identified, variations in its frequency will be evidenced as a function of task requirements...

b) A rhythmic and a continuous mode

An interesting theory has been pushed forward by Charles Schroeder and Peter Lakatos, which deserves to be mentioned in the context of a discussion on task dependence. They have proposed that the brain may dynamically shift between variably rhythmic operating modes as a function of task demands; at the extremes, they describe a “continuous/vigilance” mode and a “rhythmic” mode. A recurring example of the “continuous” mode in their writings (C Schroeder & P Lakatos, 2009; Charles Schroeder, Peter Lakatos, C.-ming Chen, Radman, & Barczak, 2009) is that of a cat watching a mouse hole,

waiting for a mouse to appear – the time at which the mouse will come out of the hole is unpredictable and thus they argue that it would be detrimental for the brain to instantiate a rhythmic mode, as it is extremely important to react extremely fast at the first sign of the mouse exiting the hole. They thus propose that the brain can take in information continuously, millisecond per millisecond, in the right settings. However, it is metabolically taxing to maintain neuronal population in such an up state for a long period of time, and whenever possible the brain will adopt a rhythmic mode. The precise rhythm at which information is taken can entrain to the structure of the stimulus stream (they argue that natural stimulation is often explicitly rhythmic and predictable). This has been proposed before : it is one of the premises of Mari Riess Jones’s “dynamic attending theory” (M. R. Jones & Boltz, 1989). Lakatos and Schroeder have presented experimental evidence of such entrainment using neural recordings : attention could selectively enhance one of two rhythmic streams by phase locking a neuronal oscillation to the attended stream – this was done with a rhythmic stimulus at 1.5hz. The framework thus posits that the brain can entrain to a host of externally imposed rhythms. This does seem like a efficient scheme, but can the brain entrain to higher frequencies? According to Schroeder and his colleagues, the brain can go all the way to a continuous processing mode. However, I am not aware of any direct evidence of a “continuous” processing mode. Our finding of an average intrinsic attentional rhythm of about 7hz may eventually speak to an upper limit on the external rhythms that attention can entrain to.

4. Entraining the rhythm of covert attention

A natural oscillator, perturbed by a periodic force, can in the right conditions become synchronized to the externally imposed rhythm. This means that the oscillating element may start cycling with the same period as the external force, becoming entrained, or locked, to the external events. The best way to achieve entrainment is to use a periodic force of the same frequency as the natural rhythm. Such 1:1 frequency locking entrainment of neuronal oscillations has been reliably induced *in vitro* (Gouwens et al., 2010) and *in vivo* in rats (Ozen et al., 2010).

There are two threads in the literature which deal with the entrainment of brain rhythms in humans: entrainment through rhythmic stimulus presentation (typically, visual flashes or auditory tones), and entrainment through externally applied stimulation (e.g., repetitive transcranial magnetic stimulation or transcranial alternative current stimulation). As argued in a recent review by Thut, Schyns and Gross (Gregor Thut, Schyns, & Gross, 2011), one key requirement for classical entrainment of an oscillation is that external force influences the oscillating element by direct interaction. This is a thorny issue in human entrainment paradigms – unless the entraining rhythmic stimulation is imperceptible, its periodicity will engage attentional orienting mechanisms, e.g. anticipation of the next predictable event in the sequence. This will likely cause entrainment of oscillatory activity as well, but it is an indirect, endogenous interaction rather than a direct one. This is a recognized confound for people interested in direct entrainment, but this is also what we are most interested in, in the context of this thesis : in what range of frequencies can attention synchronize itself to an external stimulus? Thut and colleagues point out that the upper limit of internally controlled entrainment should lie in the alpha- to theta- range, using 100-300ms as a lower limit to the periodicity that can be achieved, based on the purported dynamics for the build up of expectation (M M Müller, Teder-Sälejärvi, & S. A. Hillyard, 1998). Lakatos and Schroeder have undoubtedly demonstrated attentional entrainment at 1.5hz. The upper limit of attentional entrainment remains to be tested empirically. A recent study performed by Kyle Mathewson and colleagues (Mathewson, Fabiani, Gratton, D. M. Beck, & Lleras, 2010) demonstrated that pulsed sensory stimulation at alpha-frequency through visual entrainers prior to presentation of a target leads to enhanced target detection, but only when the target appeared at the time when the next entraining

stimulus would have occurred (metacontrast masking paradigm). Whether this truly reflects attentional entrainment is not clear though, and should be addressed. The rhythm that was imposed was higher than the intrinsic rhythm of attention which we have evidenced in this thesis, and if attention was truly entrained, this may result in other enhancements that should be investigated (in terms of reaction time, simultaneity windows, etc).

5. Covert rhythmic sampling and overt rhythmic sampling: the egg and the chicken?

From the onset of this thesis, we chose to study covert attention – meaning that the focus of attention is dissociated from the center of gaze. This is, arguably, a rather unnatural situation : spatial attention is usually tightly linked to eye movements, in humans, simply because we have much better vision at the fovea and we tend to move our eyes to the focus of attention (unless it is not socially appropriate to do so). Because attentional processes and eye movements remain clearly dissociable, we have focused on attention isolated from eye movements as our main interest. But, can we completely ignore a feature as prominent as eye movements? How does the rhythm of the eyes relate to the purported 7hz covert rhythm?

I mentioned in the introduction how our eyes are constantly moving – on average, in free-viewing conditions, we perform about 3-4 saccades per second (Otero-millan, Troncoso, Macknik, Serrano-pedraza, & Martinez-conde, 2008). Each eye movement induces a major change in the retinal input, hence a major transient. In most of our experiments, subjects fixate, meaning that they voluntarily keep their eyes still, usually anchored at a central fixation point. Even in such a situation, it has been shown that the eyes are not entirely still. Three classes of so-called fixational eye movements have been described: drift (slow eye movements), tremor (high frequency, low amplitude movements) and microsaccades. The last class, microsaccades, consists of small amplitude, high velocity shifts of eye position. They have been found to share many properties with regular saccades (e.g., they follow the “main sequence”, a linear relationship between peak velocity and magnitude, when both are plotted on a logarithmic scale (Otero-millan et al., 2008)), and in fact they share the same neural generators (Hafed, Goffart, & Krauzlis, 2009). While trained subjects can to some extent suppress their microsaccades, it is difficult to do so for a long period. Microsaccadic rate, during prolonged fixation, can vary, but it is usually found to be 1Hz or so (Otero-millan et al., 2008) (maybe faster in monkeys, 3.3Hz in (Bosman, Womelsdorf, Robert Desimone, & Fries, 2009)).

Attention and eye movements are tightly linked, as we argued previously – some researchers have proposed that spatial attention and saccade planning are in fact two descriptions of the same process (Rizzolatti et al., 1987). Why does the rhythm of eye movements not correspond to the intrinsic rhythm of attention (7Hz)? I can think of a simple argument : operating muscles is metabolically costly – the organism can not afford to move the eyes at more than 3-4hz on average. My speculative claim is that eye movements should be phase locked to the attentional rhythm – but not every attentional sample necessarily gives rise to an eye movement. A recent experiment in our group (Drewes & R. VanRullen, 2011) may speak partly to this claim: Drewes and VanRullen looked for a brain oscillation whose phase correlated with the time it took to generate an eye movement, in three different saccadic reaction time tasks. The main phase effect was found in the 11-17hz band, at fronto-central electrodes (a close look at the plots may also show an effect at ~7hz, though less significant). Though this is an interesting result, an experiment which would correlate the timing of naturally performed exploratory saccades with an ongoing brain rhythm may speak more directly to my claim – I may set it up shortly. There is in fact

already some evidence in the literature for a 7Hz rhythm governing (micro)saccadic eye movements. Otero-Millan and colleagues (Otero-millan et al., 2008) looked at InterSaccadic Intervals (ISI) in a variety of tasks, including free-viewing – (1) free exploration of a blank screen (2) free exploration of a natural scene (3) visual search using a visual puzzle display, where subjects have to compare two versions of the same image presented side by side and (4) visual search using a ‘where is Waldo?’ display, where subjects have to, well, look for Waldo – and fixation – same displays, the task being to maintain fixation. They show that the normalized distributions of ISIs, pooled across tasks, are very similar for saccade-saccade intervals, microsaccade-microsaccade intervals and even saccade-microsaccade (and vice-versa) intervals (Figure 83a). This is a very strong argument that microsaccades and saccades are governed by the same rhythmic activity. There are some differences in the distributions of ISIs with respect to the task that the subjects are engaged in : when subjects fixate, or free-view a blank screen, long ISIs become more frequent. However, the peak of the ISI distribution (μ parameter in Figure 83b) is almost invariant – and it lies between 120 and 150ms. This period corresponds very well to an underlying 7Hz generator...

The distribution of intermicrosaccadic intervals reported by Bosman and colleagues in monkeys (Bosman et al., 2009) peaks at about 0.25s, which is about twice as long. Does it mean the rhythm underlying monkeys’ sampling of their visual environment is slower (about 4Hz)? What does it mean in terms of their conscious experience?

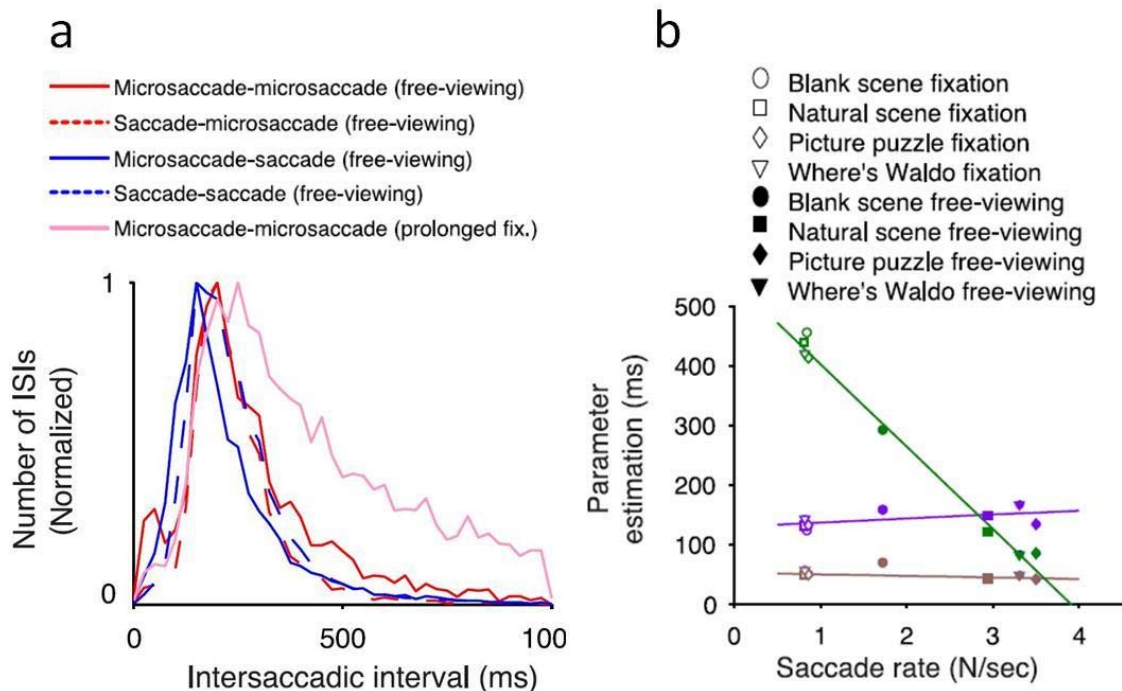


Figure 83 a InterSaccadic Interval (ISI) distributions are very similar between saccades and microsaccades, pointing to the same rhythm underlying their generation. **b** Distributions of ISIs can be fitted with ex-gaussian density functions, as is classically done with reaction time distributions (Van Zandt, 2000). Schematically, it is as if there were two stages : earlier ISIs are normally distributed (described by a gaussian density function, with mean μ and standard deviation σ), while later ISIs are exponentially distributed (described by the time constant of the decay τ). Here, you see that the τ parameter (green) is the only one affected by the task subjects are engaged into – the μ parameter (purple), which corresponds roughly to the peak of the distribution, is constant, and corresponds well to an underlying 7Hz generator.

Overt sampling behaviors, like the saccadic eye movements which we just focused on, are the rule rather than the exception. When rodents want to process a smell, they often automatically switch from regular breathing to a sniffing behavior, intaking air in small, rapid puffs at 4-10Hz (Wesson, Verhagen, & Wachowiak, 2009). For active haptic exploration, rats use their whiskers in a rhythmic pattern, with a frequency of 4-10Hz (Fanselow & Nicolelis, 1999; Kleinfeld, Berg, & O'Connor, 1999). Investigating the relationships between such overt sampling rhythms and central rhythms is an avenue of research which may bring very important insights into the way the brain processes information for perception.

C. Putting it all together : could perception rely on a 7Hz attentional sampling of information nested in occipito- parietal alpha oscillations?

1. The pulsed inhibition account of attentional selection

The function of the alpha rhythm, the prominent rhythm in the EEG, has long been a matter of controversy. Initially, it was shown to be modulated by the level of attention paid to the visual environment – losing focus, performing internal tasks such as mental calculations or working memory, closing the eyes all lead to increased alpha rhythm amplitude, hence it was interpreted as a neural correlate of idling in the visual system. But recently, a host of experimental findings has led to this view being reconsidered, with much evidence that the alpha rhythm and its precise dynamics play an active role in attention and conscious perception. Reviews of the function of the alpha rhythm have flourished in recent years which all propose slightly different theoretical accounts on how we should understand existing empirical evidence (Simon Hanslmayr, Gross, Wolfgang Klimesch, & K. L. Shapiro, 2011; Ole Jensen & Mazaheri, 2010; Wolfgang Klimesch, Paul Sauseng, & Simon Hanslmayr, 2007; Mathewson et al., 2011; Mazaheri & Ole Jensen, 2010; S. Palva & J. M. Palva, 2007).

Spontaneous variations of the *amplitude* of the ongoing occipito-parietal alpha rhythm have been correlated with visual performance at threshold, for the detection (Ergenoglu et al., 2004) and discrimination (van Dijk, Schoffelen, Oostenveld, & Ole Jensen, 2008) of brief unmasked visual stimuli, the detection of a metacontrast masked visual stimulus (Mathewson et al., 2009) as well as the detection of phosphenes elicited by transcranial magnetic stimulation (V. Romei, V. Brodbeck, et al., 2008). Occipito-parietal alpha amplitude is also predictive of between-subject performance variability, as shown in a backward masked letter discrimination task (Simon Hanslmayr et al., 2007, 2005) as well as in the detection of phosphenes elicited by transcranial magnetic stimulation (Vincenzo Romei, Rihs, Verena Brodbeck, & Gregor Thut, 2008). There is thus little doubt left that the amplitude of occipito-parietal alpha oscillations indexes visual cortex excitability (but see (S. Palva & J. M. Palva, 2007)). The *phase* of occipito-parietal alpha oscillations has also been linked to fluctuations of visual cortex excitability on a shorter time scale, influencing the detection of a metacontrast masked target stimulus (Mathewson et al., 2009), and the detection of a phosphene elicited by transcranial magnetic stimulation (Dugue et al., 2011).

Though an influence of alpha phase on excitability of the human visual cortex was not yet demonstrated, it had logically been strongly suspected for a long time. Klimesch theorized that alpha represents an “inhibitory timing” mechanism (Wolfgang Klimesch et al., 2007): a reduction in alpha amplitude reflects a state of higher excitability, whereas high amplitudes reflect a state of inhibition – instantiating cycles of excitability, whereby the timing of neuronal firing is restricted to short time windows. Mathewson’s finding of alpha phase effects on the detection of a metacontrast mask was restricted to trials when alpha amplitude was high (Mathewson et al., 2009), which corroborates Klimesch’s original proposal. Mathewson coins the term “pulsed inhibition” for his theoretical account, but the premises are basically the same. An additional, critical piece of evidence which adds support to this account is the revelation of an amplitude asymmetry, a.k.a. baseline shift, in the alpha rhythm (Nikulin et al., 2007): the peaks (troughs) are modulated stronger than the troughs (peaks), irrespective of the DC offset of the signal (Figure 84a). This non-sinusoidal feature of alpha band

could perception rely on a 7Hz attentional sampling of information nested in occipito-parietal alpha oscillations?

activity was further investigated by Jensen and colleagues, in the context of the debate over the genesis of classical evoked potentials (van Dijk, van der Werf, Mazaheri, Medendorp, & Ole Jensen, 2010; Mazaheri & Ole Jensen, 2008). However, Mazaheri and Jensen recently proposed that amplitude asymmetry may also play a key function in the dynamic active role of the alpha rhythm, coining the term “rhythmic pulsing” (Mazaheri & Ole Jensen, 2010). This twist provides a nice mechanism, with a plausible physiological implementation (unidirectional current flow coming from GABAergic interneurons), through which alpha instantiates “bouts of inhibition”, which affect processing as indexed by gamma activity (Osipova, Hermes, & Ole Jensen, 2008; Voytek et al., 2010)(Figure 84a). Mathewson incorporated the asymmetry idea in a recent reformulation of his theory, remaining agnostic regarding the signature of visual processing and referring loosely to cortical excitability (Mathewson et al., 2011)(Figure 84b).

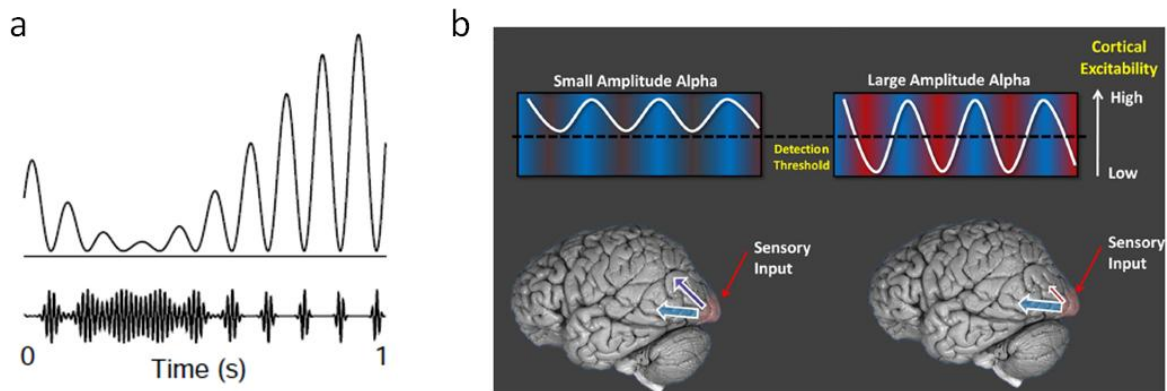


Figure 84 a. Rhythmic pulsing and its implications for processing – a hypothesis. The gamma burst can only occur when the alpha signal is low enough, e.g. at the troughs. Thus, the period of gamma activity becomes briefer with stronger alpha. When the alpha activity is sufficiently weak, gamma can occur during the whole cycle. Reproduced from (Mazaheri & Ole Jensen, 2010) **b.** A schematic of the pulsed-inhibition account of alpha oscillations. **(left)** when alpha amplitude is low, cortical excitability is sufficiently high to produce consistently high levels of processing, and constant detection regardless of phase. **(right)** when alpha power is high, certain periods of its phase are inhibitory for visual processing and target detection differs as a function of the phase at which the target was presented. Reproduced from (Mathewson et al., 2011)

Shifting or maintaining attention voluntarily to one side of the visual field leads to fluctuations in the power and topography of alpha oscillations : occipito-parietal alpha power increases ipsilaterally and decreases contralaterally, expectedly affecting task performance (Capotosto, Babiloni, Romani, & Maurizio Corbetta, 2009; Kelly et al., 2006; Rihs, C. M. Michel, & Gregor Thut, 2007, 2009; P Sauseng et al., 2005; Gregor Thut et al., 2006; Worden et al., 2000; Wyart & Catherine Tallon-Baudry, 2009). Interestingly, this effect may not be restricted to spatial attention : attending either to the color or the motion of visual stimuli induces shifts in alpha power and scalp distribution as well (Min & Christoph S Herrmann, 2007; A. C. Snyder & John J Foxe, 2010). Alpha oscillations can thus be controlled by top-down signals – they emerge as a potential mechanism for attentional selection. Jensen and Mazaheri go one step further in the interpretation (Ole Jensen & Mazaheri, 2010): they propose that pulsed inhibition by the alpha rhythm functionally blocks off task-irrelevant pathways, a mechanism for dynamic routing of information. This is presented as a viable alternative to other gating mechanisms, such as phase synchronization between areas (Figure 85).




Figure 85 Consider a situation in which information is supposed to be routed from node a to node b but not from node a to node c. Two possible gating mechanisms. **a** Gating by phase synchronization – the information flow from a to c is blocked by adjusting the phase difference. **b** Gating by inhibition – node c is actively suppressed by functional inhibition (reflected by high alpha band activity), gating the information flow from a to b.

2. Attentional selection through pulsed inhibition works if downstream areas integrate information over time.

Imagine that the brain is trying to reduce processing at an ignored location : it does it by increasing alpha oscillations at the visual neurons involved in processing this location, as argued previously. The duty cycle of the superimposed gamma rhythm (the main oscillatory signature of active neural processing (Ole Jensen et al., 2007; Jerbi et al., 2009); see the discussion of multiplexed representations in the visual system, page 22) decreases and less information is passed on to downstream areas *on average*. In a hypothetical situation where downstream areas would read the amount of information millisecond by millisecond, you would find that you would still get quite a bit of information about this location, every 100ms in the case of a 10Hz alpha rhythm. There are two options to reduce the amount of information downstream areas get about the ignored stimulus at any point in time:

- 1) either the brain makes sure that downstream areas oscillate at 10Hz as well, in counterphase, so that no information can go through
- 2) or, downstream areas integrate information over some time.

The first mechanism may be slightly difficult to implement if there are many locations that need to be ignored, but it is an option. I performed a toy model simulation of the second mechanism, to get a better idea of how it would work. I modeled two input signals, both modulated at 10Hz, with a random phase relationship. The alpha modulation is high for the ignored signal, low for the attended one (Figure 86a). I modeled the information received by a downstream area as an integration of the signals over a temporal window of a certain length, using an arbitrarily chosen weighing of information within this window (a sigmoid function). As shown on Figure 86b, information needs to be integrated over at least a few tens of milliseconds for the signal from attended location to dominate all the time. The worst case scenario corresponds to the two incoming signals being out of phase. Even then, it is easy to see that an integration period larger than one half-cycle (if all points in the integration window are weighed equally; longer otherwise) of the alpha oscillation should protect the system from spurious bouts of attending to the wrong side.

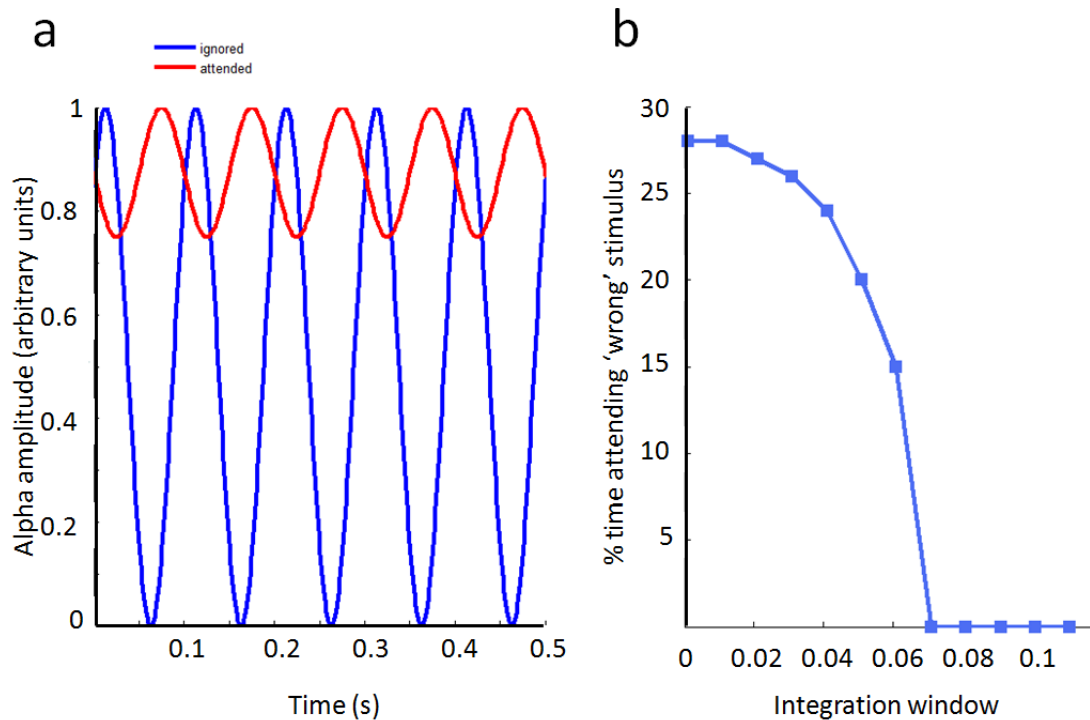


Figure 86 A toy model of attentional selection through pulsed inhibition. **a** The output of two locations, attended (red) and ignored (blue), over 0.5 seconds. The two signals are modulated at 10Hz, and have a random phase relationship. Higher (asymmetric, modulation of troughs) alpha amplitude implements pulsed inhibition, whereby information output is restricted to shorter time windows. **b** The percentage of time that the ignored location receives more processing than the attended location decreases as the length of the integration window increases. In this implementation, it is necessary to integrate over at least 70ms for attention to be effective at all times.

3. The 7Hz attentional rhythm as a readout rhythm

In the toy simulation, we used a sliding window, which can easily be implemented in neural hardware. However, integration could also rely on an oscillation in the readout area – information would be collected most when the downstream area is in a receptive state, and integrated within each cycle. This could be the role of the 7Hz attentional rhythm which we have evidenced in this work. 7Hz corresponds to a 144ms period – if information is collected mostly in half of a full cycle, it would be enough for efficient attentional selection in the pulsed inhibition framework.

Why would evolution have chosen to rely on rhythmic sampling rather than continuous readout with some temporal integration? A metabolic argument may be advanced : by closing the doors of perception intermittently, we spare some neural activity... One could also argue, as psychologists of the 1950s did, that it is necessary to have discrete perceptual moments to detect change in the dynamic environment that we live in, and to sense time.

We simulated a 7Hz-modulated attentional readout, for a situation similar to the one depicted in Figure 86. Using a duty cycle corresponding to half the period, we expectedly got similar attentional effects to the ones observed with a sliding integration window of similar length (~70ms). The interesting result came from simulating a situation when both locations

are simultaneously attended. This is done by having an intermediate amplitude alpha rhythm for both locations (Figure 87a) – there is likely to be a cost to attending to two locations simultaneously, and thus the amplitude of alpha should be higher than when a single location is attended (I am not aware of any empirical data pertaining to this in the literature, though). When implementing this, with a random phase shift between the two input signals, we found that the two locations seem to receive preferential processing in turn (Figure 87b) – a feature that we did not enforce, but that comes from the interaction between the 10Hz and 7Hz frequencies. This can be understood using plain logic: suppose that, at some point, a sample taken by our 7Hz attentional rhythm is biased towards stimulus A; since the next sample will be taken about 150ms later (i.e. 1.5 periods of the alpha oscillation) the bias will necessarily be reversed (towards B) at the next sample. The oscillation is of weak amplitude, and can be somewhat irregular (two samples in a row may sometimes be taken at the same location). Our toy model thus predicts a behavior of attention which we have evidenced in this thesis (broadband tagging paradigm, page 131). Note that we had previously argued that a serial behavior of attention implies a sampling rhythm, whereas the contrary is not necessarily true (attention could in principle be periodic, but take samples including all attended stimuli when it is focused on several of them): yet, in this model, the serial behavior of attention comes as a natural consequence of attentional sampling.




Figure 87 **a** Signals from two concurrently attended locations, over 1.5s, with a randomly chosen phase relationship **b** The amount of information received from the red signal compared to the information received from the blue signal. The locations seem to receive preferential processing in turn, in successive samples (though two samples in a row may be taken at the same location on occasion).

This is nothing more than a toy model, which offers a way to think of the 7Hz attentional rhythm which we have evidenced in our work. It aims at emphasizing the rhythm of covert attention is distinct from the occipito-parietal alpha rhythm, and proposes a link (very speculative at this point) between these rhythms. Much work remains to be done to fully understand the interplay of the various brain oscillations that shape perception...

Epilogue

Throughout this work, we were guided by a radical hypothesis : that (visual) perception may rely on a rhythmic sampling of incoming (visual) information by attention. Such a proposal affects the way we should think about conscious perception in many respects. It is thus only natural that this work made use of numerous converging paradigms in trying to find empirical support for this hypothesis. We explored such avenues as temporal aliasing in perception, divided attention, drug-induced disturbances of motion perception, fluctuations of visual detection thresholds and temporal framing. We used various techniques as well, from surveying to psychophysical experimentation and electrical brain recordings (with state-of-the-art analysis techniques). Each approach has yielded some insights and lent some support to our central hypothesis. As this work comes to a close, I do feel that it is a significant contribution to the field, and it definitely allowed us to refine our hypothesis – my own vision of how it may all work has greatly evolved over the past four years.

Since I started this work, the field has gotten quite crowded... Investigations of the phase of ongoing oscillations are now taken on by an increasing number of laboratories around the world. This is a very welcome evolution. Joining hands in the quest for oscillatory mechanisms in perception is likely to yield tremendous insights in the years to come. While we have focused, due to the techniques we used, on lower frequency rhythms, a complete account of how oscillations shape perception will very likely involve higher frequencies – we are for example very keen on the multiplexing idea presented briefly at the end of the introduction (page 22). More experimental work, guided by theoretical proposals, is critically needed to constrain our imaginative minds.

It's a pretty exciting time.

References

- [1] Adamo, M., Pun, C., Pratt, J., & Ferber, S. (2008). Your divided attention, please! The maintenance of multiple attentional control sets over distinct regions in space. *Cognition*, *107*(1), 295-303. doi:10.1016/j.cognition.2007.07.003
- [2] Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, *2*(2), 284-99.
- [3] Ahissar, E., Sosnik, R., & Haidarliu, S. (2000). Transformation from temporal to rate coding in a somatosensory thalamocortical pathway. *Nature*, *406*(6793), 302-6. doi:10.1038/35018568
- [4] Ahissar, Ehud, & Arieli, A. (2001). Figuring Space by Time. *Neuron*, *32*(2), 185-201. doi:10.1016/S0896-6273(01)00466-4
- [5] Alais, D., & Blake, R. (1999). Neural strength of visual attention gauged by motion adaptation. *Nature neuroscience*, *2*(11), 1015-8. doi:10.1038/14814
- [6] Allport, D. A. (1968). Phenomenal simultaneity and the perceptual moment hypothesis. *British Journal of Psychology*, *59*(4), 395-406. doi:10.1111/j.2044-8295.1968.tb01154.x
- [7] Amano, K., Arnold, D., Takeda, T., & Johnston, A. (2008). Alpha band amplification during illusory jitter perception. *Journal of Vision*, *8*(10), 3. doi:10.1167/8.10.3.Introduction
- [8] Anderson, C. H., Van Essen, David C, & Olshausen, B. A. (2005). Directed visual attention and the dynamic control of information flow. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of attention* (p. 744). New York, NY, USA: Academic Press.
- [9] Anderson, S. J., & Burr, D C. (1985). Spatial and temporal selectivity of the human motion detection system. *Vision research*, *25*(8), 1147-54.
- [10] Anderson, W. H., & O'Malley, J. E. (1972). Trifluoperazine for the "trailing" phenomenon. *JAMA: the journal of the American Medical Association*, *220*(9), 1244-5.
- [11] Andrews, T. J., & Purves, D. (1997). Similarities in normal and binocularly rivalrous viewing. *Proceedings of the National Academy of Sciences of the United States of America*, *94*(18), 9905-8.
- [12] Armstrong, K. M., Fitzgerald, J. K., & Moore, Tirin. (2006). Changes in visual receptive fields with microstimulation of frontal cortex. *Neuron*, *50*(5), 791-8. doi:10.1016/j.neuron.2006.05.010
- [13] Arnold, D. H., & Johnston, A. (2003). Motion-induced spatial conflict. *Nature*, *425*(September), 181-184. doi:10.1038/nature01866.1.
- [14] Awh, E., & Pashler, H. (2000). Evidence for split attentional foci. *Journal of Experimental Psychology*, *26*(2), 834-846. APA AMERICAN PSYCHOLOGICAL ASSOCIATION. doi:10.10371/0096-1523.26.2.834
- [15] Bahramisharif, A., van Gerven, M., Heskes, T., & Jensen, Ole. (2010). Covert attention allows for continuous control of brain-computer interfaces. *European Journal of Neuroscience*, *31*(8), 1501-1508. doi:10.1111/j.1460-9568.2010.07174.x
- [16] Baker, C. L., Hess, Robert F, & Zihl, J. (1991). Residual motion perception in a "motion-blind" patient, assessed with limited-lifetime random dot stimuli. *Journal of Neuroscience*, *11*(2), 454-61.

- [17] Balan, P. F., & Gottlieb, J. (2009). Functional significance of nonspatial information in monkey lateral intraparietal area. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 29(25), 8166-76. doi:10.1523/JNEUROSCI.0243-09.2009
- [18] Barry, R. J., Rushby, J. A., Johnstone, S. J., Clarke, A. R., Croft, R. J., & Lawrence, C. A. (2004). Event-related potentials in the auditory oddball as a function of EEG alpha phase at stimulus onset. *Clinical neurophysiology: official journal of the International Federation of Clinical Neurophysiology*, 115(11), 2593-601. doi:10.1016/j.clinph.2004.06.004
- [19] Bartley, S. H. (1951). The psychophysiology of vision. In S. S. Stevens (Ed.), *Handbook of experimental psychology* (pp. 921-984). Oxford, England: Wiley.
- [20] Bauer, M., Oostenveld, R., Peeters, M., & Fries, P. (2006). Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 26(2), 490-501. doi:10.1523/JNEUROSCI.5228-04.2006
- [21] Beckers, G., & Homberg, V. (1992). Cerebral visual motion blindness: transitory akinetopsia induced by transcranial magnetic stimulation of human area V5. *Proceedings of the Royal Society B: Biological Sciences*, 249, 173-178.
- [22] Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289-300.
- [23] Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87(1), 527-570. doi:10.1007/BF01797193
- [24] Bichot, N. P., Cave, K R, & Pashler, H. (1999). Visual selection mediated by location: feature-based selection of noncontiguous locations. *Perception & psychophysics*, 61(3), 403-23.
- [25] Bishop, G. H. (1932). Cyclic changes in excitability of the optic pathway of the rabbit. *American Journal of Physiology*, 103(1), 213-224.
- [26] Bisley, J. W. (2010). The Neural Basis of Visual Attention. *The Journal of physiology*, 1, 49-57. doi:10.1113/jphysiol.2010.192666
- [27] Blanke, O., Landis, T., Mermoud, C., Spinelli, L., & Safran, A. B. (2003). Direction-selective motion blindness after unilateral posterior brain damage. *European Journal of Neuroscience*, 18(3), 709-722. doi:10.1046/j.1460-9568.2003.02771.x
- [28] Bosman, C. a, Womelsdorf, T., Desimone, Robert, & Fries, P. (2009). A microsaccadic rhythm modulates gamma-band synchronization and behavior. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 29(30), 9471-80. doi:10.1523/JNEUROSCI.1193-09.2009
- [29] Bowen, R. W. (1989). Two pulses seen as three flashes: a superposition analysis. *Vision research*, 29(4), 409-17.
- [30] Braddick, O. (1974). A short-range process in apparent motion. *Vision research*, 14(7), 519-27.
- [31] Bragin, A., Jandó, G., Nádasdy, Z., Hetke, J., Wise, K., & Buzsáki, G. (1995). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 15(1 Pt 1), 47-60.
- [32] Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial vision*, 10(4), 433-6.
- [33] Broadbent, D. (1958). *Perception and Communication* (p. 338). London, England: Pergamon Press.
- [34] Bruce, V., Green, P. R., & Georgeson, M. A. (1996). *Visual Perception: Physiology, Psychology, Ecology*. Hove, UK: Laurence Erlbaum Associates.
- [35] Burr, D C, Morrone, M. C., & Ross, J. (1994). Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature*, 371(6497), 511-3. doi:10.1038/371511a0

- [36] Busch, Niko A., & VanRullen, Rufin. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, 1-6. doi:10.1073/pnas.1004801107
- [37] Busch, N.A., Dubois, J., & VanRullen, Rufin. (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience*, 29(24), 7869. Soc Neuroscience. doi:10.1523/JNEUROSCI.0113-09.2009
- [38] Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science (New York, N.Y.)*, 315(5820), 1860-2. doi:10.1126/science.1138071
- [39] Buschman, T. J., & Miller, E. K. (2009). Serial, Covert Shifts of Attention during Visual Search Are Reflected by the Frontal Eye Fields and Correlated with Population Oscillations. *Neuron*, 63, 386-396. doi:10.1016/j.neuron.2009.06.020
- [40] Buzsáki, György. (2006). *Rhythms of the Brain* (p. 464). Oxford University Press, USA.
- [41] Buzsáki, György, & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science (New York, N.Y.)*, 304(5679), 1926-9. doi:10.1126/science.1099745
- [42] Callaway, E., & Yeager, C. L. (1960). Relationship between reaction time and electroencephalographic alpha phase. *Science*, 132(3441), 1765-1766.
- [43] Cang, J., & Isaacson, J. S. (2003). In vivo whole-cell recording of odor-evoked synaptic transmission in the rat olfactory bulb. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(10), 4108-16.
- [44] Capotosto, P., Babiloni, C., Romani, G. L., & Corbetta, Maurizio. (2009). Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(18), 5863-72. doi:10.1523/JNEUROSCI.0539-09.2009
- [45] Carlson, T. a, Hogendoorn, Hinze, & Verstraten, F. a J. (2006). The speed of visual attention: what time is it? *Journal of vision*, 6(12), 1406-11. doi:10.1167/6.12.6
- [46] Carmel, D., Walsh, V., Lavie, N., & Rees, G. (2010). Right parietal TMS shortens dominance durations in binocular rivalry. *Current biology*, 20(18), R799-800. doi:10.1016/j.cub.2010.07.036
- [47] Carrasco, M, & Yeshurun, Y. (1998). The contribution of covert attention to the set-size and eccentricity effects in visual search. *Journal of experimental psychology. Human perception and performance*, 24(2), 673-92.
- [48] Cavanagh, P. (1991). Short-range vs long-range motion: not a valid distinction. *Spatial vision*, 5(4), 303-9.
- [49] Cavanagh, P, & Mather, G. (1989). Motion: the long and short of it. *Spatial vision*, 4(2-3), 103-29.
- [50] Cavanagh, Patrick, & Alvarez, G. A. (2005). Tracking multiple targets with multifocal attention. *Trends in cognitive sciences*, 9(7), 349-54. doi:10.1016/j.tics.2005.05.009
- [51] Cavanaugh, James, & Wurtz, Robert H. (2004). Subcortical modulation of attention counters change blindness. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(50), 11236-43. doi:10.1523/JNEUROSCI.3724-04.2004
- [52] Cave, Kyle R, Bush, W. S., & Taylor, T. G. G. (2010). Split attention as part of a flexible attentional system for complex scenes: comment on Jans, Peters, and De Weerd (2010). *Psychological review*, 117(2), 685-96. doi:10.1037/a0019083
- [53] Chakravarthi, R., & VanRullen, Rufin. (2011). Bullet trains and steam engines: Exogenous attention zips but endogenous attention chugs along. *Journal of Vision*, in press.
- [54] Chaudhuri, A., & Chaudhuri, A. (1990). Modulation of the motion aftereffect by selective attention. *Nature*, 344(6261), 60-2. doi:10.1038/344060a0

- [55] Cheatham, P G, & White, C T. (1952). Temporal numerosity: I. Perceived number as a function of flash number and rate. *Journal of experimental psychology*, 44(6), 447-51.
- [56] Clifford, C. W. G., & Ibbotson, M. R. (2002). Fundamental mechanisms of visual motion detection: models, cells and functions. *Progress in neurobiology*, 68(6), 409-37. doi:10.1016/S0301-0082(02)00154-5
- [57] Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature neuroscience*, 12(12), 1594-600. doi:10.1038/nn.2439
- [58] Corbetta, M., & Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3), 215–229. doi:10.1038/nrn755
- [59] Corbetta, Maurizio, Patel, G., & Shulman, Gordon L. (2008). The reorienting system of the human brain: from environment to theory of mind. *Neuron*, 58(3), 306-24. doi:10.1016/j.neuron.2008.04.017
- [60] Cowey, A., Campana, G., Walsh, V., & Vaina, L. M. (2006). The role of human extra-striate visual areas V5/MT and V2/V3 in the perception of the direction of global motion: a transcranial magnetic stimulation study. *Experimental Brain Research*, 171, 558-562. doi:10.1007/s00221-006-0479-6
- [61] Crick, F. (1984). Function of the thalamic reticular complex: The searchlight hypothesis. *Proceedings of the National Academy of Sciences*, 81(14), 4586. National Acad Sciences.
- [62] Culham, J. C., Verstraten, F. A. J., Ashida, Hiroshi, & Cavanagh, Patrick. (2000). Independent Aftereffects of Attention and Motion. *Neuron*, 28, 607-615.
- [63] Cutrell, E. B., & Marrocco, R. T. (2002). Electrical microstimulation of primate posterior parietal cortex initiates orienting and alerting components of covert attention. *Experimental brain research*, 144, 103-113. doi:10.1007/s00221-002-1032-x
- [64] DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in neurosciences*, 18(10), 451-8.
- [65] Dehaene, S. (1993). Temporal Oscillations in Human Perception. *Psychological Science*, 4(4), 264-270. doi:10.1111/j.1467-9280.1993.tb00273.x
- [66] Desimone, R, & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual review of neuroscience*, 18, 193-222. doi:10.1146/annurev.ne.18.030195.001205
- [67] Desimone, Robert, Moran, J., & Spitzer, H. (1988). Neural mechanisms of attention in extrastriate cortex of monkeys. *Dynamic interactions in neural networks*. New York, NY, USA: Springer-Verlag.
- [68] van Dijk, H., Schoffelen, J.-M., Oostenveld, R., & Jensen, Ole. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(8), 1816-23. doi:10.1523/JNEUROSCI.1853-07.2008
- [69] van Dijk, H., van der Werf, J., Mazaheri, A., Medendorp, W. P., & Jensen, Ole. (2010). Modulations in oscillatory activity with amplitude asymmetry can produce cognitively relevant event-related responses. *Proceedings of the National Academy of Sciences of the United States of America*, 107(2), 900-5. doi:10.1073/pnas.0908821107
- [70] Ding, J., Sperling, George, & Srinivasan, R. (2006). Attentional modulation of SSVEP power depends on the network tagged by the flicker frequency. *Cerebral Cortex*, 16(7), 1016-29. doi:10.1093/cercor/bhj044
- [71] Drewes, J., & VanRullen, R. (2011). This Is the Rhythm of Your Eyes: The Phase of Ongoing Electroencephalogram Oscillations Modulates Saccadic Reaction Time. *Journal of Neuroscience*, 31(12), 4698-4708. doi:10.1523/JNEUROSCI.4795-10.2011
- [72] Dugue, L., Marque, P., & VanRullen, Rufin. (2011). The phase of ongoing oscillations mediates the causal relation between brain excitation and visual perception. *Journal of Neuroscience, under revi.*

- [73] Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology*, 113(4), 501–517.
- [74] Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature*, 369(6478), 313–5. doi:10.1038/369313a0
- [75] Dustman, R., & Beck, E. (1965). Phase of alpha brain waves, reaction time and visually evoked potentials. *Electroencephalography and Clinical Neurophysiology*, 18(5), 433–440. doi:10.1016/0013-4694(65)90123-9
- [76] Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. J. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biological Cybernetics*, 60(2), 121–30.
- [77] Eckstein, M. P. (1998). The Lower Visual Search Efficiency for Conjunctions Is Due to Noise and not Serial Attentional Processing. *Psychological Science*, 9(2), 111–118. doi:10.1111/1467-9280.00020
- [78] Eckstein, Miguel P., Thomas, J. P., Palmer, J., & Shimozaki, S. S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Perception & Psychophysics*, 62(3), 425–451. doi:10.3758/BF03212096
- [79] Efron, R. (1973). Conservation of temporal information by perceptual systems. *Perception and Psychophysics*, 14(3).
- [80] Egly, R., & Homa, D. (1991). Reallocation of visual attention. *Journal of experimental psychology. Human perception and performance*, 17(1), 142–59.
- [81] Ergenoglu, T., Demiralp, T., Bayraktaroglu, Z., Ergen, M., Beydagi, H., & Uresin, Y. (2004). Alpha rhythm of the EEG modulates visual detection performance in humans. *Cognitive Brain Research*, 20(3), 376–383. Elsevier. doi:10.1016/j.cogbrainres.2004.03.009
- [82] Eriksen, B. A., & Eriksen, C.W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception and Psychophysics*, 16(1), 143–149.
- [83] Eriksen, C.W., & St James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception and Psychophysics*, 40(4), 225–240.
- [84] Eriksen, C.W., & Yeh, Y. Y. (1985). Allocation of attention in the visual field. *Journal of Experimental Psychology*, 11(5), 583–597.
- [85] Eriksen, Charles W., & Collins, J. F. (1969). Temporal course of selective attention. *Journal of Experimental Psychology*, 80(2, Pt.1), 254–261. doi:10.1037/h0027268
- [86] Eriksen, Charles W., & Spencer, T. (1969). Rate of information processing in visual perception: Some results and methodological considerations. *Journal of Experimental Psychology*, 79(2, Pt.2), 1–16. doi:10.1037/h0026873
- [87] Exner, S. (1875). Experimentelle Untersuchung der einfachsten psychischen Prozesse. *Archiv für die Gesamte Physiologie des Menschen und der Thiere*, 11(1), 403–432. doi:10.1007/BF01659311
- [88] Fanselow, E. E., & Nicolelis, M. a. (1999). Behavioral modulation of tactile responses in the rat somatosensory system. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 19(17), 7603–16.
- [89] Felleman, D. J., & Van Essen, D C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, N.Y. : 1991)*, 1(1), 1–47.
- [90] Fennema, C., & Thompson, W. (1979). Velocity determination in scenes containing several moving objects. *Computer Graphics and Image Processing*, 9(4), 301–315. doi:10.1016/0146-664X(79)90097-2

- [91] Foucher, J. R., Lacambre, M., Pham, B.-T., Giersch, a, & Elliott, M a. (2007). Low time resolution in schizophrenia Lengthened windows of simultaneity for visual, auditory and bimodal stimuli. *Schizophrenia research*, 97(1-3), 118-27. doi:10.1016/j.schres.2007.08.013
- [92] Fries, P., Nikolić, D., & Singer, W. (2007). The gamma cycle. *Trends in neurosciences*, 30(7), 309-16. doi:10.1016/j.tins.2007.05.005
- [93] Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, Robert. (2001). Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science*, 291(5508), 1560-1563. doi:10.1126/science.1055465
- [94] Fries, P., Womelsdorf, T., Oostenveld, R., & Desimone, Robert. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(18), 4823-35. doi:10.1523/JNEUROSCI.4499-07.2008
- [95] Friston, K J, & Frith, C. D. (1995). Schizophrenia: a disconnection syndrome? *Clinical neuroscience (New York, N.Y.)*, 3(2), 89-97.
- [96] Fuchs, T. (2007). The temporal structure of intentionality and its disturbance in schizophrenia. *Psychopathology*, 40(4), 229-35. doi:10.1159/000101365
- [97] Féron, F.-X., Frissen, I., Boissinot, J., & Guastavino, C. (2010). Upper limits of auditory rotational motion perception. *The Journal of the Acoustical Society of America*, 128(6), 3703-14. doi:10.1121/1.3502456
- [98] van Gerven, M., Bahramisharif, A., Heskes, T., & Jensen, Ole. (2009). Selecting features for BCI control based on a covert spatial attention paradigm. *Neural networks : the official journal of the International Neural Network Society*, 22(9), 1271-7. Elsevier Ltd. doi:10.1016/j.neunet.2009.06.004
- [99] Gho, M., & Varela, Francisco J. (1988). A quantitative assessment of the dependency of the visual temporal frame upon the cortical rhythm. *Journal of physiology, Paris*, 83, 95-101.
- [100] Giersch, A., Lalanne, L., Corves, C., Seubert, J., Shi, Z., Foucher, Jack, & Elliott, Mark a. (2009). Extended visual simultaneity thresholds in patients with schizophrenia. *Schizophrenia bulletin*, 35(4), 816-25. doi:10.1093/schbul/sbn016
- [101] Gobell, J. L., Tseng, C.-huei, & Sperling, George. (2004). The spatial distribution of visual attention. *Vision research*, 44(12), 1273-96. doi:10.1016/j.visres.2004.01.012
- [102] Gottlieb, J. P., Kusunoki, M., & Goldberg, M E. (1998). The representation of visual salience in monkey parietal cortex. *Nature*, 391(6666), 481-4. doi:10.1038/35135
- [103] Gouwens, N. W., Zeberg, H., Tsumoto, K., Tateno, T., Aihara, K., & Robinson, H. P. C. (2010). Synchronization of Firing in Cortical Fast-Spiking Interneurons at Gamma Frequencies: A Phase-Resetting Analysis. (L. J. Graham, Ed.) *PLoS Computational Biology*, 6(9), e1000951. doi:10.1371/journal.pcbi.1000951
- [104] Le Grand, Y. (1937). Recherches sur la diffusion de la lumière dans l'œil humain. *Revue d'Optique*, 16, 201-214.
- [105] Granziera, C., DaSilva, A. F. M., Snyder, J., Tuch, D. S., & Hadjikhani, N. (2006). Anatomical alterations of the visual motion processing network in migraine with and without aura. *PLoS medicine*, 3(10), e402. doi:10.1371/journal.pmed.0030402
- [106] Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, Robert. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science (New York, N.Y.)*, 324(5931), 1207-10. doi:10.1126/science.1171402
- [107] Gruber, T., Müller, M M, Keil, a, & Elbert, T. (1999). Selective visual-spatial attention alters induced gamma band responses in the human EEG. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, 110(12), 2074-85.

- [108] Haber, R. N., & Hershenson, M. (1973). *The psychology of visual perception*. (Holt, Rinehart, & Winston, Eds.) (p. 398). Oxford, England. doi:0-03-081413-8
- [109] Hafed, Z. M., Goffart, L., & Krauzlis, R. J. (2009). A neural mechanism for microsaccade generation in the primate superior colliculus. *Science (New York, N.Y.)*, 323(5916), 940-3. doi:10.1126/science.1166112
- [110] Haig, A. R., & Gordon, E. (1998). Prestimulus EEG alpha phase synchronicity influences N100 amplitude and reaction time. *Psychophysiology*, 35(5), 591-595. doi:10.1017/S0048577298970512
- [111] Hanslmayr, Simon, Aslan, A., Staudigl, T., Klimesch, Wolfgang, Herrmann, C.S., & Bäuml, K. H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *Neuroimage*, 37(4), 1465–1473. Elsevier. doi:10.1016/j.neuroimage.2007.07.011
- [112] Hanslmayr, Simon, Gross, J., Klimesch, Wolfgang, & Shapiro, K. L. (2011). The role of alpha oscillations in temporal attention. *Brain Research Reviews*, 1-13. Elsevier B.V. doi:10.1016/j.brainresrev.2011.04.002
- [113] Hanslmayr, Simon, Klimesch, Wolfgang, Sauseng, Paul, Gruber, W., Doppelmayr, Michael, Freunberger, R., & Pecherstorfer, T. (2005). Visual discrimination performance is related to decreased alpha amplitude but increased phase locking. *Neuroscience Letters*, 375, 64-68. doi:10.1016/j.neulet.2004.10.092
- [114] Harel, J., Koch, Christof, & Perona, P. (2007). Graph-Based Visual Saliency. *Neural Information Processing Systems*, 19.
- [115] Herrmann, Christoph S. (2001). Human EEG responses to 1-100Hz flicker: resonance phenomena in visual cortex and their potential correlation to cognitive phenomena. *Experimental Brain Research*, 137(3-4), 346-353. doi:10.1007/s002210100682
- [116] Hess, R., Baker, C. L., & Zihl, J. (1989). The “ motion-blind ” patient: low-level spatial and temporal filters. *Journal of Neuroscience*, 9(5), 1628. Soc Neuroscience.
- [117] Hesselmann, G., Kell, C. a, & Kleinschmidt, A. (2008). Ongoing activity fluctuations in hMT+ bias the perception of coherent visual motion. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(53), 14481-5. doi:10.1523/JNEUROSCI.4398-08.2008
- [118] Hesselmann, G., Kell, C. a, Eger, E., & Kleinschmidt, A. (2008). Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proceedings of the National Academy of Sciences of the United States of America*, 105(31), 10984-9. doi:10.1073/pnas.0712043105
- [119] Hirsh, I. J., & Sherrick, C. E. (1961). Perceived order in different sense modalities. *Journal of experimental psychology*, 62(5), 423-32.
- [120] Hogendoorn, H., Carlson, T., & Verstraten, F. A. J. (2007). The time course of attentive tracking. *Journal of Vision*, 7(14), 2. doi:10.1167/7.14.2.Introduction
- [121] Holcombe, Alex. (2011). Cracks in the edifice of visual time? Retrieved May 25, 2011, from <http://alexholcombe.wordpress.com/2010/07/12/cracks-in-the-edifice-of-visual-time/>
- [122] Holcombe, AO. (2009). Seeing slow and seeing fast: two limits on perception. *Trends in Cognitive Sciences*, 13(5), 216–221. Elsevier.
- [123] Holcombe, AO, & Cavanagh, Patrick. (2001). Early binding of feature pairs for visual perception. *nature neuroscience*, 4, 127–128. NATURE AMERICA.
- [124] Holcombe, AO, & Seizova-cajic, T. (2008). Illusory motion reversals from unambiguous motion with visual, proprioceptive, and tactile stimuli. *Vision Research*, 48(17), 1743–1757. Elsevier. doi:10.1016/j.visres.2008.05.019
- [125] Holcombe, AO, Kanwisher, Nancy, & Treisman, A. (2001). The midstream order deficit. *Perception and Psychophysics*, 63(2), 322–329. [Austin, Tex., etc., Psychonomic Society, etc.].

- [126] Hoogenboom, N., & Romei, Vincenzo. (2010). Probing different time-scales of oscillatory fluctuations in visual awareness: from behavior to phase. *Frontiers in human neuroscience*, 4(November), 178. doi:10.3389/fnhum.2010.00178
- [127] Hotson, J. R., & Anand, S. (1999). The selectivity and timing of motion processing in human temporo-parieto-occipital and occipital cortex: a transcranial magnetic stimulation study. *Neuropsychologia*, 37(2), 169-79.
- [128] Howard, C. J., & Holcombe, A. O. (2008). Tracking the changing features of multiple objects: progressively poorer perceptual precision and progressively greater perceptual lag. *Vision research*, 48(9), 1164-80. doi:10.1016/j.visres.2008.01.023
- [129] Huddleston, W. E. W. E., Lewis, J. W. J. W., Phinney, R. E. R. E., & DeYoe, E. A. E. (2008). Auditory and visual attention-based apparent motion share functional parallels. *Perception and Psychophysics*, 70(7), 1207-1216. Psychonomic Society Publications. doi:10.3758/PP
- [130] Ipata, A. E., Gee, A. L., Bisley, J. W., & Goldberg, Michael E. (2009). Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. *Experimental brain research*, 192(3), 479-88. doi:10.1007/s00221-008-1557-8
- [131] Itti, L., & Koch, Christof. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision research*, 40(10-12), 1489-506.
- [132] James, W. (1890). *The Principles of Psychology. Search* (p. 711). New York, NY, USA: Henry Holt and Company.
- [133] Jans, B., Peters, J. C., & De Weerd, P. (2010). Visual spatial attention to multiple locations at once: The jury is still out. *Psychological review*, 117(2), 637-82. doi:10.1037/a0019082
- [134] Jansen, B., & Brandt, M. (1991). The effect of the phase of prestimulus alpha activity on the averaged visual evoked response☆. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 80(4), 241-250. doi:10.1016/0168-5597(91)90107-9
- [135] Jensen, O. (2001). Information transfer between rhythmically coupled networks: reading the hippocampal phase code. *Neural computation*, 13(12), 2743-61. doi:10.1162/089976601317098510
- [136] Jensen, O, & Lisman, J E. (1996). Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learning & memory (Cold Spring Harbor, N.Y.)*, 3(2-3), 279-87.
- [137] Jensen, Ole, Kaiser, J., & Lachaux, J.-P. (2007). Human gamma-frequency oscillations associated with attention and memory. *Trends in neurosciences*, 30(7), 317-24. doi:10.1016/j.tins.2007.05.001
- [138] Jensen, Ole, & Lisman, John E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends in neurosciences*, 28(2), 67-72. doi:10.1016/j.tins.2004.12.001
- [139] Jensen, Ole, & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. *Frontiers in Human Neuroscience*, 4(November), 1-8. doi:10.3389/fnhum.2010.00186
- [140] Jerbi, K., Ossandón, T., Hamamé, C. M., Senova, S., Dalal, S. S., Jung, J., Minotti, L., et al. (2009). Task-related gamma-band dynamics from an intracerebral perspective: review and implications for surface EEG and MEG. *Human brain mapping*, 30(6), 1758-71. doi:10.1002/hbm.20750
- [141] Jones, M. R., & Boltz, M. (1989). Dynamic attending and responses to time. *Psychological review*, 96(3), 459-91.
- [142] Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J. B. Long & A. D. Baddeley (Eds.), *Attention and performance* (Hilgard, pp. 187-204). Hillsdale, NJ, USA.
- [143] Juan, C.-H., Shorter-Jacobi, S. M., & Schall, J. D. (2004). Dissociation of spatial attention and saccade preparation. *Proceedings of the National Academy of Sciences of the United States of America*, 101(43), 15541-4. doi:10.1073/pnas.0403507101

- [144] Kahneman, D. (1973). *Attention and effort* (p. 240). Prentice Hall.
- [145] Kawahara, J., Enns, J. T., & Di Lollo, V. (2006). The attentional blink is not a unitary phenomenon. *Psychological Research*, 70(6), 405–413. Springer. doi:10.1007/s00426-005-0007-5
- [146] Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, John J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of neurophysiology*, 95(6), 3844-51. doi:10.1152/jn.01234.2005
- [147] Kelly, S. P., Lalor, E. C., Reilly, R. B., Member, S., & Foxe, John J. (2005). High-Density SSVEP Data for Independent Brain – Computer Communication. *Rehabilitation*, 13(2), 172-178.
- [148] Kepecs, A., Uchida, N., & Mainen, Z. F. (2006). The sniff as a unit of olfactory processing. *Chemical senses*, 31(2), 167-79. doi:10.1093/chemse/bjj016
- [149] Kim, M.-S., & Cave, K. R. (1995). Spatial Attention in Visual Search for Features and Feature Conjunctions. *Psychological Science*, 6(6), 376-380. doi:10.1111/j.1467-9280.1995.tb00529.x
- [150] Kleinfeld, D., Berg, R. W., & O'Connor, S. M. (1999). Anatomical loops and their electrical dynamics in relation to whisking by rat. *Somatosensory & motor research*, 16(2), 69-88.
- [151] Klimesch, Wolfgang, Sauseng, Paul, & Hanslmayr, Simon. (2007). EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, 53(1), 63–88. Elsevier. doi:10.1016/j.brainresrev.2006.06.003
- [152] Kline, K., & Eagleman, D. M. (2008). Evidence against the temporal subsampling account of illusory motion reversal. *Journal of Vision*, 8(4), 13. doi:10.1167/8.4.13.Introduction
- [153] Kline, K., Holcombe, AO, & Eagleman, D. M. (2004). Illusory motion reversal is caused by rivalry, not by perceptual snapshots of the visual field. *Vision Research*, 44(23), 2653–2658. Elsevier.
- [154] Kline, K., Holcombe, AO, & Eagleman, D. M. (2006). Illusory motion reversal does not imply discrete processing: Reply to Rojas et al. *Vision Research*, 46(6-7), 1158–1159. Elsevier. doi:10.1016/j.visres.2005.08.021
- [155] Koch, Christof. (2004). *The Quest for Consciousness : a Neurobiological approach* (p. 432). Englewood, CO: Roberts and Company Publishers.
- [156] Koch, Christof, & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22. Elsevier.
- [157] Koch, Christof, & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, 4, 219-227.
- [158] Kristofferson, A. B. (1967). Attention and psychophysical time. *Acta Psychologica*, 27, 93-100. doi:10.1016/0001-6918(67)90049-2
- [159] Kruglikov, S. Y., & Schiff, S. J. (2003). Interplay of electroencephalogram phase and auditory-evoked neural activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(31), 10122-7.
- [160] Lachaux, J P, Rodriguez, E., Martinerie, J., & Varela, F J. (1999). Measuring phase synchrony in brain signals. *Human brain mapping*, 8(4), 194-208.
- [161] Lakatos, Peter, Karmos, George, Mehta, A. D., Ulbert, I., & Schroeder, C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science*, 320(5872), 110-3. doi:10.1126/science.1154735
- [162] Lakatos, Peter, Shah, A. S., Knuth, K. H., Ulbert, I., Karmos, George, & Schroeder, C. E. (2005). An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *Journal of neurophysiology*, 94(3), 1904-11. doi:10.1152/jn.00263.2005

- [163] Lankheet, M. J., & Verstraten, F. a. (1995). Attentional modulation of adaptation to two-component transparent motion. *Vision research*, 35(10), 1401-12.
- [164] Latour, P. L. (1967). Evidence of internal clocks in the human operator. *Acta psychologica*, 27, 341-8.
- [165] Lewis, J. W., Beauchamp, M. S., & DeYoe, E. a. (2000). A comparison of visual and auditory motion processing in human cerebral cortex. *Cerebral cortex (New York, N.Y. : 1991)*, 10(9), 873-88.
- [166] Lichtenstein, M. (1961). Phenomenal Simultaneity with irregular timing of components of the visual stimulus. *Perceptual and Motor Skills*, 12(1), 47-60.
- [167] Lindsley, D. B. (1952). Psychological phenomena and the electroencephalogram. *Electroencephalography and Clinical Neurophysiology*, 4(4), 443-456. doi:10.1016/0013-4694(52)90075-8
- [168] Lisman, J. (2005). The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme. *Hippocampus*, 15(7), 913-22. doi:10.1002/hipo.20121
- [169] Lisman, J E, & Idiart, M. a. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science (New York, N.Y.)*, 267(5203), 1512-5.
- [170] Liu, Y., Yttri, E. a, & Snyder, L. H. (2010). Intention and attention: different functional roles for LIPd and LIPv. *Nature neuroscience*, 13(4), 495-500. Nature Publishing Group. doi:10.1038/nn.2496
- [171] Di Lollo, V., & Wilson, a E. (1978). Iconic persistence and perceptual moment as determinants of temporal integration in vision. *Vision research*, 18(12), 1607-10.
- [172] Lu, Z. L., Lesmes, L. A., & Sperling, G. (1999). Perceptual motion standstill in rapidly moving chromatic displays. *Proceedings of the National Academy of Sciences of the United States of America*, 96(26), 15374-9.
- [173] Lu, Z. L., Lesmes, L. a, & Sperling, G. (1999). The mechanism of isoluminant chromatic motion perception. *Proceedings of the National Academy of Sciences of the United States of America*, 96(14), 8289-94.
- [174] Lu, Z.-L., & Sperling, George. (2001). Three-systems theory of human visual motion perception: review and update. *Journal of the Optical Society of America A*, 18(9), 2331-2370.
- [175] Lu, Z.-lin, & Sperling, George. (1996). Three Systems Perception for Visual Motion. *Current Directions in Psychological Science*, 5(2), 44-53.
- [176] Macdonald, J. S. P., Cavanagh, Patrick, & VanRullen, Rufin. (2011). Attentional sampling of multiple wagon wheels. *Attention, perception, & psychophysics*, in press.
- [177] Mack, A., & Rock, I. (1998). *Inattentive Blindness* (p. 296). Cambridge, MA, USA: MIT Press.
- [178] Makeig, S., & Jung, T. P. (1996). Tonic, phasic, and transient EEG correlates of auditory awareness in drowsiness. *Brain research. Cognitive brain research*, 4(1), 15-25.
- [179] von der Malsburg, C. (1999). The what and why of binding: the modeler's perspective. *Neuron*, 24(1), 95-104, 111-25.
- [180] Marcar, V. L., & Cowey, A. (1992). The effect of removing superior temporal cortical motion areas in the macaque monkey: II. Motion discrimination using random dot displays. *European Journal of Neuroscience*, 4, 1228-1238.
- [181] Mathewson, K. E., Fabiani, M., Gratton, G., Beck, D. M., & Lleras, A. (2010). Rescuing stimuli from invisibility: Inducing a momentary release from visual masking with pre-target entrainment. *Cognition*, 115(1), 186-91. Elsevier B.V. doi:10.1016/j.cognition.2009.11.010
- [182] Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(9), 2725-32. doi:10.1523/JNEUROSCI.3963-08.2009

- [183] Mathewson, K. E., Lleras, A., Beck, D. M., Fabiani, M., Ro, T., & Gratton, G. (2011). Pulsed Out of Awareness: EEG Alpha Oscillations Represent a Pulsed-Inhibition of Ongoing Cortical Processing. *Frontiers in Psychology*, 2(May), 1-15. doi:10.3389/fpsyg.2011.00099
- [184] Matthews, N., Luber, B., Qian, N., & Lisanby, S. H. (2001). Transcranial magnetic stimulation differentially affects speed and direction judgments. *Experimental Brain Research*, 140(4), 397-406. doi:10.1007/s002210100837
- [185] Maunsell, J. H., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual review of neuroscience*, 10, 363-401. doi:10.1146/annurev.ne.10.030187.002051
- [186] Mazaheri, A., & Jensen, Ole. (2008). Asymmetric amplitude modulations of brain oscillations generate slow evoked responses. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 28(31), 7781-7. doi:10.1523/JNEUROSCI.1631-08.2008
- [187] Mazaheri, A., & Jensen, Ole. (2010). Rhythmic pulsing: linking ongoing brain activity with evoked responses. *Frontiers in human neuroscience*, 4(October), 177. doi:10.3389/fnhum.2010.00177
- [188] McAlonan, K., Cavanaugh, J., & Wurtz, R.H. (2008). Guarding the gateway to cortex with attention in visual thalamus. *Nature*, 456(7220), 391-394. Nature Publishing Group. doi:10.1038/nature07382
- [189] McCormick, P. A., Klein, R. M., & Johnston, S. (1998). Splitting versus sharing focal attention: comment on Castiello and Umilta (1992). *Journal of Experimental Psychology*, 24(1), 350-357. APA AMERICAN PSYCHOLOGICAL ASSOCIATION.
- [190] McEree, B., & Carrasco, Marisa. (1999). The temporal dynamics of visual search: Evidence for parallel processing in feature and conjunction searches. *Journal of Experimental Psychology: Human Perception and Performance*, 25(6), 1517-1539. doi:10.1037/0096-1523.25.6.1517
- [191] McKendrick, A. M., Badcock, D. R., Badcock, J. C., & Gurgone, M. (2006). Motion perception in migraineurs: abnormalities are not related to attention. *Cephalalgia: an international journal of headache*, 26(9), 1131-6. doi:10.1111/j.1468-2982.2006.01182.x
- [192] McMains, S. A., & Somers, D. C. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron*, 42(4), 677-686. Elsevier.
- [193] McMains, S. A., & Somers, D. C. (2005). Processing efficiency of divided spatial attention mechanisms in human visual cortex. *Journal of Neuroscience*, 25(41), 9444. Soc Neuroscience. doi:10.1523/JNEUROSCI.2647-05.2005
- [194] McPeck, R. M. (2008). Reversal of a distractor effect on saccade target selection after superior colliculus inactivation. *Journal of neurophysiology*, 99(5), 2694-702. doi:10.1152/jn.00591.2007
- [195] Mcpeek, R. M., & Keller, E. L. (2002). Saccade Target Selection in the Superior Colliculus During a Visual Search Task. *Journal of Neurophysiology*, 88, 2019-2034.
- [196] Miconi, T., & VanRullen, Rufin. (2011). A Feedback Model of Attentional Effects in the Visual Cortex. *IEEE Symposium on Computational Intelligence for Multimedia, Signal and Vision Processing*. Paris.
- [197] Middendorf, M., McMillan, G., Calhoun, G., & Jones, K. S. (2000). Brain-computer interfaces based on the steady-state visual-evoked response. *IEEE transactions on rehabilitation engineering: a publication of the IEEE Engineering in Medicine and Biology Society*, 8(2), 211-4.
- [198] Min, B.-K., & Herrmann, Christoph S. (2007). Prestimulus EEG alpha activity reflects prestimulus top-down processing. *Neuroscience letters*, 422(2), 131-5. doi:10.1016/j.neulet.2007.06.013
- [199] Mitchell, J. F., Sundberg, K. a, & Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron*, 63(6), 879-88. Elsevier Ltd. doi:10.1016/j.neuron.2009.09.013

- [200] Monosov, I. E., & Thompson, K. G. (2009). Frontal eye field activity enhances object identification during covert visual search. *Journal of neurophysiology*, *102*(6), 3656-72. doi:10.1152/jn.00750.2009
- [201] Monto, S., Palva, S., Voipio, Juha, & Palva, J. M. (2008). Very slow EEG fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *28*(33), 8268-72. doi:10.1523/JNEUROSCI.1910-08.2008
- [202] Moore, C. M., Egeth, H., Berglan, L. R., & Luck, S. J. (1996). Are attentional dwell times inconsistent with serial visual search? *Psychonomic Bulletin & Review*, *3*(3), 360-365. doi:10.3758/BF03210761
- [203] Moore, T., & Fallah, M. (2001). Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(3), 1273-6. doi:10.1073/pnas.021549498
- [204] Moore, Tirin, Armstrong, K. M., & Fallah, Mazyar. (2003). Visuomotor origins of covert spatial attention. *Neuron*, *40*(4), 671-83.
- [205] Moore, Tirin, & Armstrong, K. M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*(January), 370-373. doi:10.1038/nature01285.
- [206] Morgan, S. T., Hansen, J. C., & Hillyard, S. a. (1996). Selective attention to stimulus location modulates the steady-state visual evoked potential. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(10), 4770-4.
- [207] Morrone, M., Ross, J., & Burr, David C. (2005). Saccadic eye movements cause compression of time as well as space. *Nature Neuroscience*, *8*(7), 950-954. Nature Publishing Group.
- [208] Murphree, O. D. (1954). Maximum rates of form perception and the alpha rhythm: an investigation and test of current nerve net theory. *Journal of experimental psychology*, *48*(1), 57-61.
- [209] Müller, H. J., & Rabbitt, P. M. (1989). Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *Journal of experimental psychology. Human perception and performance*, *15*(2), 315-30.
- [210] Müller, J. R., Philiastides, M. G., & Newsome, William T. (2005). Microstimulation of the superior colliculus focuses attention without moving the eyes. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(3), 524-9. doi:10.1073/pnas.0408311101
- [211] Müller, M M, Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature neuroscience*, *1*(7), 631-4. Nature Publishing Group. doi:10.1038/2865
- [212] Müller, M M, Malinowski, P., Gruber, T., & Hillyard, S. A. (2003). Sustained division of the attentional spotlight. *Nature*, *424*(6946), 309-12. Nature Publishing Group. doi:10.1038/nature01812
- [213] Müller, M M, Picton, T. W., Valdes-Sosa, P., Riera, J., Teder-Sälejärvi, W. a, & Hillyard, S. a. (1998). Effects of spatial selective attention on the steady-state visual evoked potential in the 20-28 Hz range. *Brain research. Cognitive brain research*, *6*(4), 249-61.
- [214] Müller, Matthias M, & Hübner, R. (2002). Can the spotlight of attention be shaped like a doughnut? Evidence from steady-state visual evoked potentials. *Psychological science : a journal of the American Psychological Society / APS*, *13*(2), 119-24.
- [215] Nakayama, K., & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision research*, *29*(11), 1631-47.
- [216] Newsome, W., & Pare, E. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience*, *8*(6), 2201. Soc Neuroscience.

- [217] Nikulin, V. V., Linkenkaer-Hansen, K., Nolte, G., Lemm, S., Müller, K. R., Ilmoniemi, R. J., & Curio, G. (2007). A novel mechanism for evoked responses in the human brain. *The European journal of neuroscience*, 25(10), 3146-54. doi:10.1111/j.1460-9568.2007.05553.x
- [218] Nishida, S., & Ashida, H. (2000). A hierarchical structure of motion system revealed by interocular transfer of flicker motion aftereffects. *Vision research*, 40(3), 265-78.
- [219] Noudoost, B., Chang, M. H., Steinmetz, N. a, & Moore, Tirin. (2010). Top-down control of visual attention. *Current opinion in neurobiology*, 20(2), 183-190. Elsevier Ltd. doi:10.1016/j.conb.2010.02.003
- [220] Oksama, L., & Hyönä, J. (2008). Dynamic binding of identity and location information: a serial model of multiple identity tracking. *Cognitive psychology*, 56(4), 237-83. doi:10.1016/j.cogpsych.2007.03.001
- [221] Osipova, D., Hermes, D., & Jensen, Ole. (2008). Gamma power is phase-locked to posterior alpha activity. *PLoS one*, 3(12), e3990. doi:10.1371/journal.pone.0003990
- [222] Otero-millan, J., Troncoso, X. G., Macknik, S. L., Serrano-pedraza, I., & Martinez-conde, S. (2008). Saccades and microsaccades during visual fixation , exploration , and search : Foundations for a common saccadic generator. *Journal of Vision*, 8, 1-18. doi:10.1167/8.14.21.Introduction
- [223] Ozen, S., Sirota, a, Belluscio, M. a, Anastassiou, C. a, Stark, E., Koch, C., & Buzsaki, G. (2010). Transcranial Electric Stimulation Entrain Cortical Neuronal Populations in Rats. *Journal of Neuroscience*, 30(34), 11476-11485. doi:10.1523/JNEUROSCI.5252-09.2010
- [224] O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, 401(6753), 584-7. doi:10.1038/44134
- [225] O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3), 317-30. doi:10.1002/hipo.450030307
- [226] Palmer, J. (1995). Attention in Visual Search: Distinguishing Four Causes of a Set-Size Effect. *Current Directions in Psychological Science*, 4(4), 118-123. doi:10.1111/1467-8721.ep10772534
- [227] Palmer, S. E. (1999). *Vision science : from Photons to Phenomenology* (p. 810). Cambridge, MA: MIT Press.
- [228] Palva, S., & Palva, J. M. (2007). New vistas for alpha-frequency band oscillations. *Trends in neurosciences*, 30(4), 150-8. doi:10.1016/j.tins.2007.02.001
- [229] Pascual-Leone, Alvaro, & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, 292(5516), 510-2. doi:10.1126/science.1057099
- [230] Peng, H., Long, F., & Ding, C. (2005). Feature selection based on mutual information: criteria of max-dependency, max-relevance, and min-redundancy. *IEEE transactions on pattern analysis and machine intelligence*, 27(8), 1226-38. doi:10.1109/TPAMI.2005.159
- [231] Penttonen, M., & Buzsáki, György. (2003). Natural logarithmic relationship between brain oscillators. *Thalamus & Related Systems*, 2(2), 145-152. doi:10.1016/S1472-9288(03)00007-4
- [232] Petersen, S. E., Robinson, D. L., & Morris, J. D. (1987). Contributions of the pulvinar o visual spatial attention. *Neuropsychologia*, 25(1A), 97-105.
- [233] Piantoni, G., Kline, K. A., & Eagleman, D. M. (2010). Beta oscillations correlate with the probability of perceiving rivalrous visual stimuli. *Journal of vision*, 10(13), 18. doi:10.1167/10.13.18
- [234] Pieron, H. (1952). *The sensations, their functions, processes, and mechanisms*. New Haven, CT, USA: Yale University Press.
- [235] Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, 32(1), 3-25.

- [236] Purves, D., Paydarfar, J. a, & Andrews, T. J. (1996). The wagon wheel illusion in movies and reality. *Proceedings of the National Academy of Sciences of the United States of America*, 93(8), 3693-7.
- [237] Rajkai, C., Lakatos, Peter, Chen, C.-M., Pincze, Z., Karmos, Gyorgy, & Schroeder, C. E. (2008). Transient cortical excitation at the onset of visual fixation. *Cerebral cortex*, 18(1), 200-9. doi:10.1093/cercor/bhm046
- [238] Reddy, L., Rémy, F., Vayssière, N., & VanRullen, R. (2011). Neural correlates of the continuous Wagon Wheel Illusion: a functional MRI study. *Human brain mapping*, 32(2), 163-70. doi:10.1002/hbm.21007
- [239] Reeves, A., & Sperling, George. (1986). Attention gating in short-term visual memory. *Psychological Review*, 93(2), 180-206.
- [240] Regan, D. (1977). Steady-state evoked potentials. *Journal of the Optical Society of America*, 67(11), 1475-1489.
- [241] Reichardt, W. E. (1961). Autocorrelation, a principle for the evaluation of sensory information by the central nervous system. In W. A. Rosenblith (Ed.), *Sensory Communication* (pp. 303-317). Cambridge, MA, USA: MIT Press.
- [242] Rensink, R. A., Regan, J. K. O., & Clark, J. J. (1997). TO SEE OR NOT TO SEE : The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, 8(5), 368-373.
- [243] Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61(2), 168-185. Elsevier. doi:10.1016/j.neuron.2009.01.002
- [244] Rezec, A., Krekelberg, B., & Dobkins, K. R. (2004). Attention enhances adaptability: evidence from motion adaptation experiments. *Vision research*, 44(26), 3035-44. doi:10.1016/j.visres.2004.07.020
- [245] Rihs, T. a, Michel, C. M., & Thut, Gregor. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *The European journal of neuroscience*, 25(2), 603-10. doi:10.1111/j.1460-9568.2007.05278.x
- [246] Rihs, T. a, Michel, C. M., & Thut, Gregor. (2009). A bias for posterior alpha-band power suppression versus enhancement during shifting versus maintenance of spatial attention. *NeuroImage*, 44(1), 190-9. Elsevier Inc. doi:10.1016/j.neuroimage.2008.08.022
- [247] Rizzo, M., Nawrot, M., & Zihl, J. (1995). Motion and shape perception in cerebral akinetopsia. *Brain*, 118, 1105-27.
- [248] Rizzolatti, G., Riggio, L., Dascola, I., & Umiltà, C. (1987). Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia*, 25(1A), 31-40.
- [249] Rojas, D., Carmona-Fontaine, C., López-Calderón, J., & Aboitiz, F. (2006). Do discreteness and rivalry coexist in illusory motion reversals? *Vision Research*, 46(6-7), 1155-1157. Elsevier. doi:10.1016/j.visres.2005.07.023
- [250] Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. (2008). Spontaneous Fluctuations in Posterior -Band EEG Activity Reflect Variability in Excitability of Human Visual Areas. *Cerebral Cortex*, 18(9), 2010-2018. doi:10.1093/cercor/bhm229
- [251] Romei, Vincenzo, Rihs, T., Brodbeck, Verena, & Thut, Gregor. (2008). Resting electroencephalogram alpha-power over posterior sites indexes baseline visual cortex excitability. *Neuroreport*, 19(2), 203-8. doi:10.1097/WNR.0b013e3282f454c4
- [252] Roopun, A. K., Kramer, M. a, Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, Roger D, Kopell, N. J., et al. (2008). Temporal Interactions between Cortical Rhythms. *Frontiers in neuroscience*, 2(2), 145-54. doi:10.3389/neuro.01.034.2008
- [253] Rossi, A. F., & Paradiso, M. A. (1995). Feature-specific effects of selective visual attention. *Vision research*, 35(5), 621-34.

- [254] Saalmann, Y. B., Pigarev, I. N., & Vidyasagar, T. R. (2007). Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science (New York, N.Y.)*, *316*(5831), 1612-5. doi:10.1126/science.1139140
- [255] Sack, A. T., Kohler, A., Linden, D. E. J., Goebel, R., & Muckli, L. (2006). The temporal characteristics of motion processing in hMT/V5+: combining fMRI and neuronavigated TMS. *NeuroImage*, *29*(4), 1326-35. doi:10.1016/j.neuroimage.2005.08.027
- [256] Sacks, O. (2004). In the River of Consciousness. *The New York Review of Books*, *51*(1).
- [257] Sadaghiani, S., Hesselmann, G., & Kleinschmidt, A. (2009). Distributed and antagonistic contributions of ongoing activity fluctuations to auditory stimulus detection. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *29*(42), 13410-7. doi:10.1523/JNEUROSCI.2592-09.2009
- [258] Sadaghiani, S., Hesselmann, G., Friston, Karl J, & Kleinschmidt, A. (2010). The relation of ongoing brain activity, evoked neural responses, and cognition. *Frontiers in systems neuroscience*, *4*(June), 20. doi:10.3389/fnsys.2010.00020
- [259] Saenz, M., Buračas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*(7), 631-632.
- [260] Saenz, Melissa, Lewis, L. B., Huth, A. G., Fine, I., & Koch, Christof. (2008). Visual Motion Area MT+/V5 Responds to Auditory Motion in Human Sight-Recovery Subjects. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *28*(20), 5141-8. doi:10.1523/JNEUROSCI.0803-08.2008
- [261] Sanford, J. A. (1971). A periodic basis for perception and action. In W. Colquhoun (Ed.), *Biological Rhythms and Human Performance* (pp. 179-209). Academic Press.
- [262] van Santen, J. P., & Sperling, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America. A, Optics and image science*, *2*(2), 300-21.
- [263] Sauseng, P, Klimesch, W, Stadler, W., Schabus, M., Doppelmayr, M, Hanslmayr, S, Gruber, W. R., et al. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *The European journal of neuroscience*, *22*(11), 2917-26. doi:10.1111/j.1460-9568.2005.04482.x
- [264] Schroeder, C, & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, *32*(1), 9-18. doi:10.1016/j.tins.2008.09.012
- [265] Schroeder, Charles, Lakatos, Peter, Chen, C.-ming, Radman, T., & Barczak, A.-marie. (2009). Aligning the Brain in a Rhythmic World.
- [266] Seizova-cajic, T., Smith, J. L., Taylor, J. L., & Gandevia, S. C. (2007). Proprioceptive Movement Illusions Due to Prolonged Stimulation: Reversals and Aftereffects. *PLoS One*, *2*(10), e1037. doi:10.1371/journal.pone.0001037
- [267] Shallice, T. (1964). The detection of change and the perceptual moment hypothesis. *British Journal of Statistical Psychology*, *17*(2), 113-135.
- [268] Shiffrin, R. M., & Gardner, G. T. (1972). Visual processing capacity and attentional control. *Journal of experimental psychology*, *93*(1), 72-82.
- [269] Shipp, S., Jong, B., Zihl, J., Frackowiak, R., & Zeki, S. M. (1994). The brain activity related to residual motion vision in a patient with bilateral lesions of V5. *Brain*, *117*(5), 1023. Oxford Univ Press.
- [270] Shulman, G L, Remington, R. W., & McLean, J. P. (1979). Moving attention through visual space. *Journal of experimental psychology. Human perception and performance*, *5*(3), 522-6.
- [271] Shulman, Gordon L, & Wilson, J. (1987). Spatial frequency and selective attention to local and global information. *Perception*, *16*(1), 89-101.

- [272] Shulman, Gordon L. (1991). Attentional modulation of mechanisms that analyze rotation in depth. *Journal of Experimental Psychology: Human Perception and Performance*, 17(3), 726-737. doi:10.1037/0096-1523.17.3.726
- [273] Silvanto, J., Cowey, A., Lavie, N., & Walsh, V. (2005). Striate cortex (V1) activity gates awareness of motion. *Nature Neuroscience*, 8(2), 143-144. New York, NY: Nature America Inc., c1998-
- [274] Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24(1), 49-65, 111-25.
- [275] Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., & Buzsáki, György. (2008). Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron*, 60(4), 683-97. doi:10.1016/j.neuron.2008.09.014
- [276] Snowden, R. J., & Hess, R F. (1992). Temporal frequency filters in the human peripheral visual field. *Vision research*, 32(1), 61-72.
- [277] Snyder, A. C., & Foxe, John J. (2010). Anticipatory attentional suppression of visual features indexed by oscillatory alpha-band power increases: a high-density electrical mapping study. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 30(11), 4024-32. doi:10.1523/JNEUROSCI.5684-09.2010
- [278] Soltesz, I., & Deschênes, M. (1993). Low- and high-frequency membrane potential oscillations during theta activity in CA1 and CA3 pyramidal neurons of the rat hippocampus under ketamine-xylazine anesthesia. *Journal of neurophysiology*, 70(1), 97-116.
- [279] Spearman, C. (1937). *Psychology down the ages*. London, England: McMillan.
- [280] Sperling, George, & Weichselgartner, E. (1995). Episodic theory of the dynamics of spatial attention. *Psychological review*, 102(3), 503-532. [Washington, etc.] American Psychological Association [etc.].
- [281] Steriade, M. (1999). Coherent oscillations and short-term plasticity in corticothalamic networks. *Trends in neurosciences*, 22(8), 337-45.
- [282] Steriade, Mircea, McCormick, D. A., & Sejnowski, T. J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science (New York, N.Y.)*, 262(5134), 679-85.
- [283] Stetson, C., Fiesta, M. P., & Eagleman, D. M. (2007). Does Time Really Slow Down during a Frightening Event? *PLoS One*, 2(12), 10-12. Public Library of Science. doi:10.1371/journal.pone.0001295
- [284] Stroud, J. M. (1967). The fine structure of psychological time. *Annals of the New York Academy of Sciences*, 138(2), 623-631. doi:10.1111/j.1749-6632.1967.tb55012.x
- [285] Tallon-Baudry, C, Bertrand, O, Delpuech, C., & Pernier, J. (1997). Oscillatory gamma-band (30-70 Hz) activity induced by a visual search task in humans. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 17(2), 722-34.
- [286] Tallon-Baudry, C, Bertrand, O, Peronnet, F., & Pernier, J. (1998). Induced gamma-band activity during the delay of a visual short-term memory task in humans. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 18(11), 4244-54.
- [287] Tallon-Baudry, Catherine, & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, 3(4), 151-162. Elsevier Ltd.
- [288] Tass, P., Rosenblum, M. G., Weule, J., Kurths, J., Pikovsky, A., Volkman, J., Schnitzler, A., et al. (1998). Detection of n:m Phase Locking from Noisy Data: Application to Magnetoencephalography. *Physical Review Letters*, 81(15), 3291-3294. doi:10.1103/PhysRevLett.81.3291
- [289] Taylor, K., Mandon, S., Freiwald, W. a, & Kreiter, a K. (2005). Coherent oscillatory activity in monkey area v4 predicts successful allocation of attention. *Cerebral cortex (New York, N.Y. : 1991)*, 15(9), 1424-37. doi:10.1093/cercor/bhi023

- [290] Teuber, H.-L. (1959). Some alterations in behavior after cerebral lesions in man. In A. D. Bass (Ed.), *Evolution of nervous control from primitive organisms in man* (pp. 157-194). Washington, DC, USA: Amer. Ass. Adv. Sci.
- [291] Thiele, a, Henning, P., Kubischik, M., & Hoffmann, K.-P. (2002). Neural mechanisms of saccadic suppression. *Science (New York, N.Y.)*, 295(5564), 2460-2. doi:10.1126/science.1068788
- [292] Thomas, N. W. D., & Paré, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. *Journal of neurophysiology*, 97(1), 942-7. doi:10.1152/jn.00413.2006
- [293] Thompson, B., Aaen-Stockdale, C., Koski, L., & Hess, Robert F. (2009). A double dissociation between striate and extrastriate visual cortex for pattern motion perception revealed using rTMS. *Human brain mapping*, 30(10), 3115-26. doi:10.1002/hbm.20736
- [294] Thompson, K. G., Biscoe, K. L., & Sato, T. R. (2005). Neuronal basis of covert spatial attention in the frontal eye field. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 25(41), 9479-87. doi:10.1523/JNEUROSCI.0741-05.2005
- [295] Thut, Gregor, Nietzel, A., Brandt, S. a, & Pascual-Leone, Alvaro. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 26(37), 9494-502. doi:10.1523/JNEUROSCI.0875-06.2006
- [296] Thut, Gregor, Schyns, P., & Gross, J. (2011). Entrainment of perceptually relevant brain oscillations by non-invasive rhythmic stimulation of the human brain. *Frontiers in Psychology, in press*.
- [297] Tononi, G., & Edelman, G. M. (2000). Schizophrenia and the mechanisms of conscious integration. *Brain research. Brain research reviews*, 31(2-3), 391-400.
- [298] Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like tweedledum and tweedledee but they can (and should) be distinguished. *Psychological Science*, 1(1), 46-54.
- [299] Traub, R D, Cunningham, M. O., Gloveli, T., Lebeau, F. E. N., Bibbig, A., Buhl, E. H., & Whittington, M. A. (2003). GABA-enhanced collective behavior in neuronal axons underlies persistent gamma-frequency oscillations. *Proc Natl Acad Sci U S A*, 100(19), 11047-11052.
- [300] Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological review*, 76(3), 282-99.
- [301] Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, 12(1), 97-136.
- [302] Trimble, J. L., & Potts, A. M. (1975). Ongoing occipital rhythms and the VER. I. Stimulation at peaks of the alpha-rhythm. *Investigative ophthalmology*, 14(7), 537-46.
- [303] Tsal, Y. (1983). Movements of attention across the visual field. *Journal of Experimental Psychology*, 9(4), 523-530.
- [304] Tseng, C.-huei, Gobell, J. L., Lu, Z.-lin, & Sperling, George. (2006). When motion appears stopped : Stereo motion standstill. *Proceedings of the National Academy of Sciences of the United States of America*, 103(40), 14953-14958.
- [305] Tsotsos, J. K. (1990). Analyzing vision at the complexity level. *Behavioral and brain sciences*, 13(3), 423-469.
- [306] Tsotsos, J. K. (2005). Computational foundations for attentive processes. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of attention* (p. 744). New York, NY, USA: Academic Press.
- [307] Ungerleider, L. G., & Mishkin, M. (1982). Two Cortical Visual Systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior* (pp. 549-586). Cambridge, MA, USA: MIT Press.
- [308] Vaina, L. M., Cowey, A., Eskew, R. T., LeMay, M., & Kemper, T. (2001). Regional cerebral correlates of global motion perception: Evidence from unilateral cerebral brain damage. *Brain*, 124, 310-321.

- [309] Le Van Quyen, M., Foucher, J, Lachaux, J., Rodriguez, E., Lutz, a, Martinerie, J., & Varela, F J. (2001). Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *Journal of neuroscience methods*, 111(2), 83-98.
- [310] VanRullen, Rufin. (2003). Visual saliency and spike timing in the ventral visual pathway. *Journal of physiology, Paris*, 97(2-3), 365-77. doi:10.1016/j.jphysparis.2003.09.010
- [311] VanRullen, Rufin. (2006). The continuous Wagon Wheel Illusion is object-based. *Vision research*, 46(24), 4091-5. doi:10.1016/j.visres.2006.07.030
- [312] VanRullen, Rufin. (2007). The continuous Wagon Wheel Illusion depends on, but is not identical to neuronal adaptation. *Vision research*, 47(16), 2143-9. doi:10.1016/j.visres.2007.03.019
- [313] VanRullen, Rufin, & Koch, Christof. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, 7(5), 207–213. Elsevier.
- [314] VanRullen, Rufin, & Thorpe, S. J. (2002). Surfing a spike wave down the ventral stream. *Vision research*, 42(23), 2593-615.
- [315] VanRullen, Rufin, Carlson, T., & Cavanagh, Patrick. (2007). The blinking spotlight of attention. *Proceedings of the National Academy of Sciences*, 104(49), 19204. National Acad Sciences.
- [316] VanRullen, Rufin, Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *Trends in neurosciences*, 28(1), 1-4. doi:10.1016/j.tins.2004.10.010
- [317] VanRullen, Rufin, Pascual-Leone, Alvaro, & Battelli, L. (2008). The Continuous Wagon Wheel Illusion and the “When” Pathway of the Right Parietal Lobe: A Repetitive Transcranial Magnetic Stimulation Study. *PLoS One*, 3(8), e2911. doi:10.1371/journal.pone.0002911
- [318] VanRullen, Rufin, Reddy, L., & Koch, Christof. (2005). Attention-driven discrete sampling of motion perception. *Proceedings of the National Academy of Sciences*, 102(14), 5291. National Acad Sciences.
- [319] VanRullen, Rufin, Reddy, L., & Koch, Christof. (2006). The continuous wagon wheel illusion is associated with changes in electroencephalogram power at approximately 13 Hz. *Journal of Neuroscience*, 26(2), 502-507.
- [320] Vanhatalo, S., Tallgren, P., Andersson, S., Sainio, K., Voipio, J, & Kaila, K. (2002). DC-EEG discloses prominent, very slow activity patterns during sleep in preterm infants. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology*, 113(11), 1822-5.
- [321] Varela, F J, Toro, a, John, E. R., & Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia*, 19(5), 675-86.
- [322] Varela, F., Lachaux, J.P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4), 229–239. Nature Publishing Group.
- [323] Venables, P. H. (1960). Periodicity in reaction time. *British Journal of Psychology*, 51(1), 37-43. doi:10.1111/j.2044-8295.1960.tb00722.x
- [324] Vinck, M., Lima, B., Womelsdorf, T., Oostenveld, R., Singer, W., Neuenschwander, S., & Fries, P. (2010). Gamma-phase shifting in awake monkey visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(4), 1250-7. doi:10.1523/JNEUROSCI.1623-09.2010
- [325] Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., & Pennartz, C. M. a. (2010). The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization. *NeuroImage*, 51(1), 112-22. Elsevier Inc. doi:10.1016/j.neuroimage.2010.01.073
- [326] Vion-Dury, J., & Blanquet, F. (2008). *Pratique de l'EEG - Bases neurophysiologiques; principes d'interpretation et de prescription* (p. 208). Masson.

- [327] Voytek, B., Canolty, R. T., Shestyuk, A., Crone, N. E., Parvizi, J., & Knight, R. T. (2010). Shifts in gamma phase-amplitude coupling frequency from theta to alpha over posterior cortex during visual tasks. *Frontiers in human neuroscience*, 4(October), 191. doi:10.3389/fnhum.2010.00191
- [328] Wardak, C., Ibos, G., Duhamel, J.-R., & Olivier, E. (2006). Contribution of the monkey frontal eye field to covert visual attention. *The Journal of neuroscience: the official journal of the Society for Neuroscience*, 26(16), 4228-35. doi:10.1523/JNEUROSCI.3336-05.2006
- [329] Wardak, C., Olivier, E., & Duhamel, J.-R. (2004). A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron*, 42(3), 501-8.
- [330] Warren, J. D., Zielinski, B. A., Green, G. G. R., Rauschecker, J. P., & Griffiths, T. D. (2002). Perception of sound-source motion by the human brain. *Neuron*, 34(1), 139-48.
- [331] Weichselgartner, E., & Sperling, George. (1987). Dynamics of automatic and controlled visual attention. *Science*, 238(4828), 778. AAAS.
- [332] Wesson, D. W., Verhagen, J. V., & Wachowiak, M. (2009). Why sniff fast? The relationship between sniff frequency, odor discrimination, and receptor neuron activation in the rat. *Journal of neurophysiology*, 101(2), 1089-102. doi:10.1152/jn.90981.2008
- [333] White, C T. (1963). Temporal numerosity and the psychological unit of perception. *Psychological Monographs*, 77(12), 1-37.
- [334] White, Carroll T, & Cheatham, Paul G. (1953). Temporal numerosity : II . Evidence for central factors influencing perceived number. *Journal of Experimental Psychology*, 46(4), 283-287.
- [335] White, Carroll T, & Cheatham, Paul G. (1959). Temporal numerosity : IV . A comparison of the major senses. *Journal of Experimental Psychology*, 58(6), 441-444.
- [336] Wolfe, J. M. (1998). Visual Search. In Harold Pashler (Ed.), *Attention* (pp. 13-74). East Sussex, UK: Psychology Press Ltd.
- [337] Worden, M. S., Foxe, J J, Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 20(6), RC63.
- [338] Wurtz, Robert H, & Goldberg, Michael E. (1972). The primate superior colliculus and the shift of visual attention. *Investigative Ophthalmology*, 11(6), 441-449.
- [339] Wyart, V., & Sergent, C. (2009). The Phase of Ongoing EEG Oscillations Uncovers the Fine Temporal Structure of Conscious Perception. *Journal of Neuroscience*, 29(41), 12839-12841.
- [340] Wyart, V., & Tallon-Baudry, Catherine. (2009). How Ongoing Fluctuations in Human Visual Cortex Predict Perceptual Awareness: Baseline Shift versus Decision Bias. *Journal of Neuroscience*, 29(27), 8715-8725.
- [341] Yamagishi, N. (2003). Attentional modulation of oscillatory activity in human visual cortex. *NeuroImage*, 20(1), 98-113. doi:10.1016/S1053-8119(03)00341-0
- [342] Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic bulletin & review*, 7(3), 424-65.
- [343] Zeki, S. M. (1991). Cerebral akinetopsia (visual motion blindness). A review. *Brain*, 114 (Pt 2, 811-24.
- [344] Zihl, J., Von Cramon, D., & Mai, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, 106(2), 313. Oxford Univ Press.
- [345] Zirnsak, M., Beuth, F., & Hamker, F. H. (2011). Split of spatial attention as predicted by a systems-level model of visual attention. *European Journal of Neuroscience*, 33(June), 2035-2045. doi:10.1111/j.1460-9568.2011.07718.x

