Towards a " Neuro-Encryption " system : from understanding the influence of brain oscillations in vision to controlling perception

Sasskia Brüers

Université Fédérale

Toulouse Midi-Pyrénées

THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITÉ DE TOULOUSE

Délivré par :

Université Toulouse 3 Paul Sabatier (UT3 Paul Sabatier)

Présentée et soutenue par :
Sasskia BRÜERS
le vendredi 27 octobre 2017

Titre :

Towards a "Neuro-Encryption" system: from understanding the influence of brain oscillations in vision to controlling perception

Vers un système de "Neuro-Encryption": de la compréhension de l'influence des oscillations cérébrales en vision au contrôle de la perception

---

École doctorale et discipline ou spécialité :
ED CLESCO : Neurosciences, comportement et cognition

Unité de recherche :
Centre de Recherche Cerveau et Cognition, CNRS, UMR 5549

Directeur/trice(s) de Thèse :

Dr. Rufin VanRullen

Jury :
Dr. Gregor Thut, University of Glasgow (Rapporteur)
Dr. Edmund Lalor, University of Rochester (Rapporteur)
Dr. Laura Dugué, Laboratoire de Psychologie et Perception (Examinatrice)

"We see things not as they are but as we are–[…] as molded by the individual peculiarities of our minds." (Patrick, 1890)

Pour Simon

ABSTRACT

Our brain activity is inherently rhythmic: oscillations can be found at all levels of organization. This rhythmicity in brain activity gives a rhythm to what we see: instead of continuously monitoring the environment, our brains take "snapshots" of the external world from 5 to 15 times a second. This creates perceptual cycles: depending on the phase of the underlying oscillation, our perceptual abilities fluctuate. Accumulating evidence shows that brains oscillations at various frequencies are instrumental in shaping visual perception.

At the heart of this thesis lies the White Noise Paradigm, which we designed as a tool to better understand the influence of oscillations on visual perception and which ultimately could be used to control visual perception. The White Noise Paradigm uses streams of flashes with random luminance (i.e. white noise) as stimuli, which have been shown to constrain brain oscillations in a predictable manner. The impulse response to WN sequences has a strong (subject specific) oscillatory component at ~10Hz akin to a perceptual echo. Since the impulse response is a model of how our brains respond to one single flash in the sequence, they can be used to reconstruct (rather than record) the brain activity to new stimulation sequences. We then present near-perceptual threshold targets embedded within the WN sequences and extract the time course of these predicted/reconstructed background oscillations around target presentation. Thus, the reconstructed EEG can be used to study the influence of the oscillatory components on visual perception, independently of other types of signals usually recorded in the EEG.

First, we validate the White Noise Paradigm by showing that: 1) the WN sequences do modulate behaviour, 2) the perceptual echoes evoked by these WN sequences are stable in time, 3) they are a (relatively) good model of the subject's recorded brain activity and 4) their neuronal basis can be found in the early visual areas.

Second, we investigate the relationship between these constrained brain oscillations and visual perception. Specifically, we show that the reconstructed EEG can help us recover the true latency at which (theta) phase influences perception. Moreover, it can help us uncover a causal influence of (alpha) power on target detection, independently from any fluctuation in endogenous factors.

Finally, capitalizing on the link between oscillations and perception, we build two algorithms used to control the perception of subjects. First, we build a "universal" forward model which can predict for any observer whether a particular target will be seen or not. Second, we build a subject-dependent model which can predict whether a particular subject (for whom EEG was recorded previously) will perceive a given target or not. Critically, this can be used to present targets optimized to be perceived by one subject only, to the detriment of all other subjects, creating a sort of "Neuro-Encryption" system.

Key-words: visual perception; ongoing oscillation; Impulse Response; White Noise Paradigm; Perceptual Cycles; Neuro-Encryption;

RÉSUMÉ

L'activité de notre cerveau est intrinsèquement rythmique : des oscillations sont observées à tous les niveaux de son organisation. Cette rythmicité de l'activité cérébrale influence notre perception. En effet, au lieu de superviser continuellement notre environnement, notre cerveau effectue de brèves « clichés » du monde extérieur (entre 5 et 15 par seconde). Cela crée des cycles perpétuels : notre perception visuelle fluctue en fonction de la phase de l'oscillation sous-jacente. De nombreuses données témoignent du fait que les oscillations cérébrales à différentes fréquences sont fondamentales à la formation de notre perception visuelle.

Lors de cette thèse, nous avons utilisé le Paradigme de Bruit Blanc comme outil pour comprendre l'influence des oscillations sur la perception visuelle et qui par extension pourra être utilisé pour contrôler cette perception. Le paradigme de bruit blanc visuel utilise des séquences de flashs dont la luminance varie aléatoirement (créant ainsi du « bruit blanc »), comme stimuli, qui contraignent l'activité cérébrale de manière prédictible. Les réponses impulsionnelles à ces séquences de bruit blanc sont caractérisées par une composante oscillatoire forte dans la bande alpha (~10Hz), similaire à un écho perceptuel. Puisque les réponses impulsionnelles sont un modèle de la réponse de notre cerveau à un flash dans la séquence de bruit blanc, elles peuvent être utilisées pour reconstruire (plutôt qu'enregistrer) l'activité cérébrale en réponse à de nouvelles séquences de stimulation. Par ailleurs, des cibles ont été introduites au sein des séquences de bruit blanc à un niveau proche du seuil de perception, et le décours temporel de cette activité reconstruite autour de la présentation des cibles a été extrait. Ainsi, l'EEG reconstruit peut être utilisé pour étudier l'influence de ces oscillations contraintes sur la perception visuelle, indépendamment des autres types de signaux généralement enregistrés dans l'EEG.

Dans un premier temps, nous avons validé le paradigme de bruit blanc en montrant que : 1) les séquences de bruits blancs influencent bien la détection des cibles, 2) les échos perceptuels évoqués par les séquences de bruit blancs sont stables dans le temps, 3) ces échos sont un bon modèle de l'activité cérébrale enregistrée par EEG, et 4) leurs bases neuronales se situent dans les aires visuelles primaires. Dans un second temps, nous avons étudié la relation entre ces oscillations cérébrales contrôlées par la séquence de bruit blanc et la détection des cibles. Ici, nous montrons que l'activité EEG reconstruite nous aide à déterminer la véritable latence à laquelle la phase de l'oscillation (thêta) influence la perception. De plus, nous avons aussi montré que l'amplitude de l'oscillation (alpha) influence la détection des cibles et ce, indépendamment des fluctuations des facteurs endogènes (tel que l'attention). Enfin, tirant parti de ce lien entre oscillation et perception, nous construisons deux algorithmes qui permettent de contrôler la perception des sujets. Tout d'abord, nous mettons au point un modèle « universel » de la perception qui permet de prédire, pour n'importe quel observateur, si une cible dans une séquence de bruit blanc sera vue ou non. Ensuite, nous construisons un modèle individuel qui utilise l'écho perceptuel de chaque sujet comme clé de cryptage et nous permet de présenter des cibles à des moments où la cible sera détectée par un sujet seulement au détriment de tous les autres sujets, créant ainsi une sorte de système de cryptage neuronal (« Neuro-Encryption »).

Mots clés: (8 max) : perception visuelle; oscillations; Réponse impulsionnelle; Paradigme de Bruit Blanc; Cycles perceptuels; Neuro-Encryption

RESUME SUBSTANTIEL

Que nous soyons assis auprès d'un fleuve à regarder les bateaux passer, ou bien au centre d'une foule au cœur de Tokyo, notre expérience du monde externe se fait de manière continue. Nous avons le sentiment qu'un flot ininterrompu d'impressions sensorielles nous arrive. Cependant, est-ce aussi le cas au niveau neuronal : notre cerveau traite-t-il l'information visuelle en continu ? Se pourrait-il que notre cerveau procède à un échantillonnage du monde extérieur ?

En effet, nous savons que l'activité cérébrale (enregistrée en autre par électro-encéphalographie, EEG) est rythmique : les potentiels électriques reflètent l'activité des neurones sous-jacents, et cette activité est naturellement oscillatoire. Ces oscillations reflètent donc des fluctuations dans l'excitabilité des neurones, qui alternent entre des moments de forte excitabilité et de faible excitabilité où la probabilité que le neurone transmette l'information est augmentée ou réduite (ex. : Bishop, 1932; Whittingstall & Logothetis, 2009). Il a été proposé que cette rythmicité dans l'activité cérébrale pourrait entrainer des fluctuations au niveau perceptuel (VanRullen, 2016b; VanRullen & Koch, 2003) : des cycles dans la perception visuelle sont créés en conséquence directe des fluctuations dans les oscillations neuronales. Le cerveau fonctionnerait donc comme un appareil photo en prenant des « instantanés » du monde environnant de 5 à 15 fois par seconde. Deux sources probables ont été proposées comme pouvant créer des cycles perceptuels : soit une source « locale » soit une source « centrale » (Harter, 1967). L'idée des « cycles d'excitabilité corticale » (Lansing, 1957; Lindsley, 1952) suppose que les cycles perceptuels sont créés suite à un mécanisme local (les fluctuations de l'excitabilité corticale décrites ci-dessus) : les fluctuations d'excitabilité des neurones pourraient servir de mécanisme de restriction des entrées visuelles, les informations arrivant aux moments de forte excitabilité seraient traitées de manière plus efficace et rapide. D'un autre côté, l'idée d'un « mécanisme d'échantillonnage central » (*central scanning mechanism)* propose au contraire que les fluctuations perceptuelles seraient le résultats d'un mécanisme central de haut niveau qui échantillonne l'information entrante de manière récurrente (Stroud, 1955). Ceci permet de créer un cadre de référence interne quant à la temporalité des différentes informations sensorielles : leur temps d'arrivée dans le système sera codé par rapport à l'oscillation sous-jacente.

Dans l'activité cérébrale enregistrée en EEG, deux grands types d'activité peuvent être distingués suivant leur relation avec le stimulus présenté : les réponses visuelles qui découlent de la stimulation visuelle (c'est-à-dire traitements neuronaux évoqués et induits par le stimulus) et les oscillations spontanées enregistrées avant (et au moment où) le stimulus arrive.

Nous allons tout d'abord voir comment chacun de ces signaux peut être enregistré. Les réponses visuelles évoquées (ou induites) sont les potentiels électriques qui surviennent immédiatement après la présentation d'une cible (par exemple, une image). Ceux-ci reflètent le traitement de l'information visuelle par les diverses aires neuronales (principalement les aires visuelles primaires). La partie du signal dont la phase est systématiquement la même en tous les essais est appelée le « potentiel visuel évoqué » (VEP pour *visual evoked potential*), qui est traditionnellement différenciée de la partie « induite », c'est-à-dire dont les phases ne sont pas systématiquement alignées par rapport au stimulus. Les réponses oscillatoires spontanées quant à elles sont enregistrées avant la présentation du stimulus et reflètent (comme décrit ci-dessus), des fluctuations dans l'excitabilité neuronale.

Lors de l'étude des corrélats neuronaux de la perception visuelle, plusieurs paradigmes ont été utilisés pour tester la contribution spécifique de ces différents types de signaux à la perception visuelle. Tout d'abord, il a été montré que les VEP peuvent extraire de façon très précise le décours temporel du traitement des informations visuelles (Hillyard, Teder-Sälejärvi, & Münte, 1998; Luck, 2005b).

Ces VEP ont été extraits par deux méthodes. Dans un paradigme de potentiels évoqués « classique », des stimuli visuels sont présentés de manière répétée et isolée avec assez d'espace entre les stimulations pour permettre à l'activité cérébrale de retourner à son niveau de base (Luck, Woodman, & Vogel, 2000). Ensuite, des époques de signaux sont créés autour de chaque stimulus, et l'information est moyennée entre les essais (et parfois entre les sujets) pour enlever le bruit (Luck, 2005b). Ce genre d'approche est cependant relativement long et coûteux au niveau attentionnel pour les patients (Lalor, Pearlmutter, Reilly, McDarby, & Foxe, 2006). La deuxième méthode prend parti des analyses utilisées pour l'identification des systèmes pour accélérer l'acquisition du VEP. Dans les systèmes non linéaires (tel que le cerveau), la présentation d'un seul stimulus (ex. un flash lumineux avec une intensité visuelle donnée) ne permet pas d'extraire une représentation fiable de la réponse du système à toutes les stimulations possibles (ex. des flashes lumineux à différentes intensités) puisque la réponse du système est non-linéaire. De ce fait, il est nécessaire de présenter toutes les

intensités possibles pour extraire la réponse moyenne du cerveau (le VEP). Ceci est fait au moyen des séquences de « bruit blanc », dont l'un des paramètres est contrôlé par l'expérimentateur (dans notre cas, l'intensité de luminance des flashs visuels) pour varier aléatoirement sur une certaine bande de fréquences (Crosse, Di Liberto, Bednar, & Lalor, 2016). Ceci permet d'accélérer grandement le taux d'acquisition du VEP et d'être sûr d'avoir testé le système avec un large éventail de stimuli (P. Z. Marmarelis & Marmarelis, 1978). Dans ce cas, la réponse moyenne du système à ce type de stimulation est appelée la réponse impulsionnelle. Elle est extraite par une cross-corrélation entre la séquence de bruit blanc et la réponse enregistrée dans le système. La partie linéaire de la réponse visuelle évoquée décrite ci-dessus correspond donc à la réponse impulsionnelle. Cette réponse impulsionnelle représente le meilleur modèle linéaire de la réponse moyenne que fera le système à un stimulus correspondant à une unité de stimulation. De ce fait, une fois extraite, elle peut être utilisée pour « reconstruire » l'activité du système à de nouveaux stimuli : il suffit pour cela de pondérer la réponse impulsionnelle par la valeur du stimulus.

Cette approche de stimulation en bruit blanc a été utilisée pour extraire le potentiel évoqué en EEG. Lalor et collègues (2006) ont été parmi les premiers à utiliser cette méthode pour accélérer l'extraction du VEP (Lalor, Pearlmutter, et al., 2006), qu'ils ont appelé « visual evoked potential spread-spectrum activity » (ou VESPA). Ces VESPA durent environ 300ms et sont caractérisées pour une forte amplitude de réponse sur les électrodes occipitales, et sont fortement corrélées avec un VEP extrait par une méthode plus traditionnelle (Lalor, Pearlmutter, et al., 2006). Plus tard, cette même approche a été utilisée, mais pour évaluer la réponse cérébrale à des latences plus longues. VanRullen et Macdonald (2012) ont montré que la stimulation en bruit blanc (i.e. à toutes les fréquences entre 0 et 80 Hz) faisait « résonner » le cerveau des participants dans la bande alpha (c'est-à-dire à environ 10 Hz). Lorsqu'un stimulus est présenté, l'intensité de la stimulation est réverbérée (gardée en « mémoire ») pendant presque 1 seconde (VanRullen & Macdonald, 2012). Ces « échos perceptuels » sont caractérisés par une variabilité inter-sujet : la fréquence de réverbération pour chaque sujet est fortement corrélée avec la fréquence de son activité dans la bande alpha au repos (VanRullen & Macdonald, 2012), et chaque sujet a une relation de phase spécifique avec le stimulus. De plus, les échos perceptuels sont stables dans le temps.

Comme nous pouvons voir, l'activité évoquée par un stimulus peut être utilisée pour mieux comprendre comment notre cerveau traite l'information visuelle. Nous allons

maintenant nous tourner vers l'autre type d'activité décrite ci-dessus : les oscillations spontanées. Dans les dernières années, un regain d'intérêt est visible pour le rôle de ces fluctuations. De nombreuses études ont montré un rôle de la fréquence (ex. : Samaha & Postle, 2015), l'amplitude (ex. : Ergenoglu et al., 2004; Worden, Foxe, Wang, & Simpson, 2000) et la phase (ex. : Busch, Dubois, & VanRullen, 2009; Mathewson, Gratton, Fabiani, Beck, & Ro, 2009) des oscillations spontanées sur la perception visuelle.

Le problème de ces études c'est qu'elles montrent seulement une corrélation entre oscillations et perception. Pour prouver que les oscillations causent des changements dans la perception visuelle, l'état de ces oscillations doit être manipulé par un facteur externe. Ceci peut être fait en utilisant deux méthodes : les méthodes de stimulation cérébrale et les méthodes de stimulation visuelle (Thut, Miniussi, & Gross, 2012). Par exemple, il a été montré que les oscillations spontanées pouvaient être entrainées par une stimulation magnétique trans-crânienne (Thut, Veniero, et al., 2011): des impulsions magnétiques appliquées de manière répétée sur le cortex visuel des sujets créaient une réverbération neuronale dans la bande de fréquence de stimulation. De plus, les oscillations « entrainées » avaient un effet sur la perception des sujets (Romei, Gross, & Thut, 2010). Pour ce qui est des techniques de stimulation visuelle, les oscillations cérébrales peuvent aussi être entrainées par la présentation répétée de flashs lumineux, ce qui a aussi été montré comme modulant la perception de manière phasique (de Graaf et al., 2013; Spaak, de Lange, & Jensen, 2014). Cependant, une des limites de ce type d'approches, c'est qu'elles nécessitent le choix à priori d'une fréquence de stimulation. Ce choix est cependant difficile, puisqu'il est nécessaire, si on veut entrainer les oscillations cérébrales de manière efficace, de choisir la fréquence de stimulation judicieusement pour avoir une correspondance 1 :1 entre fréquence d'entrainement et oscillation naturelle (Thut et al., 2012). Or, il semblerait qu'il existe des différences individuelles fortes dans la fréquence spécifique des oscillations le plus souvent impliquées dans la perception visuelle : les oscillations alpha ; chaque sujet a sa fréquence d'oscillation propre (Haegens, Cousijn, Wallis, Harrison, & Nobre, 2014). Cependant, nous avons vu précédemment que lorsque des séquences de bruit blanc étaient présentées au sujet, leur cerveau répondait spécifiquement à l'information dans la bande alpha : la fréquence de leur écho perceptuel était corrélée avec leur fréquence dans la bande alpha au repos. De ce fait, nous pourrions utiliser ces séquences de bruit blanc pour contraindre l'activité neuronale des sujets de façon individuelle, sans avoir à choisir une fréquence de stimulation spécifique.

Dans cette thèse, nous avons cherché à mieux comprendre la relation existant entre oscillations et perception visuelle. Pour cela, nous avons utilisé le 'paradigme du bruit blanc' : des séquences de bruit blanc (décrites ci-dessus) ont été utilisées pour contraindre l'activité cérébrale des sujets (enregistrée par EEG) de manière spécifique. Les réponses impulsionnelles (c'est-à-dire les échos perceptuels) ont été extraites pour caractériser la relation spécifique entre stimulation visuelle et activité cérébrale pour ce sujet. Cette réponse impulsionnelle représente le meilleur modèle linéaire de l'activité cérébrale d'un sujet donné en réponse à un flash dans la séquence. Par conséquent, nous pouvons ensuite l'utiliser pour reconstruire l'activité cérébrale en réponse à de nouvelles séquences de bruit blanc sans avoir à enregistrer l'EEG. Au contraire, nous utilisons la réponse impulsionnelle comme modèle de l'activité cérébrale, qui sera donc « reconstruite ». C'est cet EEG reconstruit qui a été utilisé tout au long de la thèse pour évaluer comment la détection de cibles imbriquées dans la séquence de luminance est influencée par les oscillations contraintes par le bruit blanc.

L'utilisation de ce paradigme nous permet de tester des hypothèses spécifiques quant au lien entre oscillations et perception, en particulier parce qu'il nous permet de modéliser seulement une partie (très spécifique) de l'activité cérébrale. La plupart des études liant les oscillations cérébrales à la perception visuelle ont utilisé des méthodes d'enregistrement de l'activité cérébrale au niveau global (EEG ou MEG) permettant ainsi de capturer avec une excellente résolution temporelle l'activité simultanée d'un grand nombre de réseaux de neurones. Cependant, cette force de l'EEG est aussi sa limite : de nombreux types d'activités sont enregistrés en même temps, ce qui ne nous permet pas de dire quelle activité est responsable des effets mentionnés ci-dessus. Pour mieux comprendre quels sont les signaux qui pourraient être modélisés par les séquences de bruit blanc, nous allons créer une taxonomie (relativement arbitraire, mais utile dans notre cas spécifique) des activités cérébrales enregistrées lors d'un EEG. Pour cela nous allons poser trois questions :

- est-ce que l'activité cérébrale répond aux séquences de bruits blancs.
- est-ce que cette activité cérébrale joue un rôle dans la visibilité des cibles.
- si oui, la phase de cette réponse est-elle déterminée (« phase-locked ») ou induite (« non-phase-locked ») ?

Ceci nous permet de distinguer six types d'activités cérébrales différentes. Tout d'abord, nous pouvons distinguer les activités cérébrales qui jouent un rôle dans la visibilité des cibles (signaux auxquels nous nous sommes intéressés pendant cette thèse et que nous avons appelés 'pertinents') de celles qui n'en jouent pas (et que nous avons donc appelés 'non pertinents'). De plus, chacun de ces signaux peut être subdivisé en trois parties, suivant la manière dont leur activité est contrainte par la séquence de bruit blanc. Les signaux dont l'état n'est pas contraint par la séquence de bruit blanc seront appelés 'spontanés'. Parmi les signaux dont l'activité est dépendante de la stimulation visuelle, nous distinguerons deux types de signaux : ceux dont la phase est 'déterminée' et ceux dont la phase est 'induite' par la séquence de bruit blanc. Cette taxonomie nous permet donc de différentier les activités suivantes:

- signaux « non-pertinents spontanés » (appelés par la suite S1)
- signaux « pertinents spontanés » (S2)
- signaux « non-pertinents déterminés » (S3)
- signaux « pertinents déterminés » (S4)
- signaux « non-pertinents induits » (S5)
- signaux « pertinents induits » (S6)

Dans le **chapitre 2,** nous nous sommes attachés à valider les hypothèses faites par le paradigme du bruit blanc. Tout d'abord, puisque nous souhaitons utiliser les échos perceptuels comme modèle de l'activité cérébrale, nous devons évaluer pendant combien de temps ce modèle est valide. Est-ce que ce modèle change dans le temps ou bien reste-t-il stable ? De plus, nous voulons savoir si c'est un « bon » modèle de l'activité cérébrale enregistrée : y-a-t-il une corrélation entre les oscillations contraintes enregistrées par EEG et celles reconstruites avec l'IRF. Nous montrons que les échos perceptuels sont stables dans le temps à des périodes plus longues que ce que l'on attendait précédemment : il faudrait presque 12 ans pour que nous ne puissions plus distinguer si deux échos perceptuels enregistrés proviennent d'un même sujet (qui a vieilli entre-temps) ou de deux sujets différents. De plus, les échos perceptuels sont un relativement bon modèle de l'activité cérébrale contrainte enregistrée par EEG : la corrélation entre les deux signaux peut atteindre $r = 0.09$ (ou même $r=0.16$ dans la bande alpha). Cette corrélation est certes relativement faible, mais elle est de valeur égale à la corrélation utilisant le VEP comme modèle de l'activité cérébrale ; modèle qui est

globalement accepté comme contenant de l'information, compte tenu du bruit inhérent de l'EEG comme méthode d'enregistrement de l'activité cérébrale (Picton et al., 2000). Enfin, nous voulons évaluer quels sont les effets de la séquence de bruit blanc sur la détection des cibles présentées à l'intérieur de la stimulation. Dans un premier temps, nous démontrons que la luminance moyenne présente autour de la présentation de la cible a un effet direct sur la perception. Des luminances plus fortes vont avoir tendance à masquer la cible alors que des luminances relativement plus faible vont améliorer sa perception. Puisque nous cherchons à relier les oscillations à la perception, nous avons décidé d'enlever les fluctuations de luminance pendant 150ms autour de la cible pour éviter qu'un lien direct entre séquence de bruit blanc et perception d'un côté, et séquence de bruit blanc et oscillation de l'autre ne nous mène à trouver une corrélation triviale entre oscillations et perception. Enfin, nous montrons que ces séquences de bruit blanc peuvent influencer la perception des cibles présentées de façon systématique. La probabilité pour qu'une cible présentée dans une même séquence de bruit blanc soit perçue de la même manière entre sujets ou pour un même sujet (qui a vu la séquence deux fois) est faible, mais significative. Cependant, la majorité de l'effet de la séquence de bruit blanc sur la perception est déterminé par un effet indépendant des sujets.

Dans le **chapitre 3**, les bases cérébrales des réponses impulsionnelles ont été examinées. En particulier, nous savons que la partie tardive des échos est constituée d'une réverbération dans la bande alpha. De ce fait, nous avons pris avantage du paradigme de bruit blanc. Lors d'une première session, la réponse impulsionnelle des sujets était extraite lors d'une session d'enregistrement EEG. Lors de la deuxième session, l'activité fonctionnelle des sujets en réponse à des séquences de bruit blanc était extraite lors d'une session d'enregistrement en imagerie à résonance magnétique fonctionnelle. Les réponses d'impulsions étaient extraites pour reconstruire la réponse cérébrale spécifique de chaque sujet à la stimulation de bruit blanc présentées dans l'IRM. L'avantage de cette méthode c'est qu'elle nous permet d'extraire le décours temporel spécifique aux échos perceptuels (et non toute l'activité cérébrale si nous avions enregistré l'EEG). C'est cette activité reconstruite qui reflète le mieux les oscillations alpha qui ont donné lieu aux échos perceptuels De ce fait, nous avons utilisé l'enveloppe de l'EEG reconstruit comme marqueur des fluctuations contraintes par les séquences de bruits dans l'IRM. Nous avons trouvé que les réverbérations neuronales avaient une base dans les aires corticales primaires (V1 et V2). Cette activité était spécifique à la bande alpha (comparée aux autres bandes de stimulation) et n'était pas

spécifique aux régions stimulée de manière rétinotopique : elle s'étendait à l'ensemble de l'aire visuelle primaire et secondaire.

Une fois que le paradigme a été validé, et que nous avons trouvé quelles sont les régions cérébrales qui sous-tendent les échos perceptuels, nous avons cherché à démontrer que l'activité cérébrale contrainte par la séquence de bruit blanc avait une influence sur la perception. Dans le **chapitre 4**, nous avons montré que la phase de l'EEG reconstruit dans la bande thêta (sur les électrodes frontales et occipitales) détermine la perception des cibles : une cible présentée à une certaine phase avait 11 % plus de chance d'être détectée qu'une cible présentée à la phase opposée. De plus, l'utilisation du paradigme de bruit blanc nous a permis de démontrer que la phase des oscillations influence la détection d'une cible *au moment même* où cette cible est traitée par le cerveau (c'est-à-dire après sa présentation). De plus, nous démontrons par le moyen de simulations, que les effets rapportés dans la littérature d'un impact de la phase pré-stimulus peuvent être expliqués par un masquage par le potentiel évoqué (en post-stimulus) des différences de phase entre cibles vues et non vues. Dans le **chapitre 5**, nous avons montré que l'amplitude de l'EEG reconstruit dans la bande alpha (sur les électrodes occipitales) avait une influence sur la détection de cibles : lorsque l'amplitude est plus basse lorsque la cible est présentée, elle est mieux perçue. Cet effet supporte le rôle inhibiteur de l'oscillation alpha sur la perception visuelle. De plus, il nous permet de démontrer qu'un lien causal existe entre amplitude des oscillations et détection de la cible. Bien que des effets similaires aient déjà été démontrés par l'utilisation de méthodes de stimulation cérébrales et visuelles, l'avantage du paradigme de bruit blanc c'est qu'il permet de contraindre la phase des oscillations de manière spécifique au sujet, sans avoir à faire d'hypothèse sur la « bonne » fréquence de stimulation.

Sur la base des effets de ces deux derniers chapitres, nous avons voulu employer ce lien entre la phase et l'amplitude des oscillations contraintes et la détection de cible pour construire des algorithmes qui nous permettraient de prédire la perception de nouvelles cibles et pour de nouveaux sujets. En d'autres termes, nous voulions être capables de prédire, pour une cible donnée dans une séquence de bruit blanc donnée, si elle sera perçue ou non. Nous avons pris deux approches. Lors de la première approche (**chapitre 6**), la prédiction de la perception était faite de manière « indépendante » des sujets. Le but était de prédire la

perception de n'importe quelle nouvelle cible (présentée dans une séquence de bruit blanc) pour n'importe quel nouveau sujet. En d'autres termes, nous avons relié les patterns d'EEG reconstruit « universels » à la perception d'une cible ; de ce fait, il nous est possible de prédire, pour toute nouvelle séquence, si les cibles seront vues ou non. Le prédicteur basé sur la phase des oscillations dans la bande thêta avait la meilleure capacité de prédiction : nous avons pu correctement prédire si une cible serait vue ou non basé sur la phase de l'oscillation. Une cible présentée à la phase prédite comme « bonne » pour la perception avait 15% de chance de plus d'être perçue qu'une cible présentée à la phase prédite comme « mauvaise » pour la perception. La deuxième approche (**chapitre 7**) se base sur une prédiction dépendante des sujets. Ici, le but était de manipuler la perception de chaque sujet spécifiquement. Nous avons extrait les patterns d'activation qui relient la phase et l'amplitude de l'EEG reconstruit à la perception de manière à ce que, quand une nouvelle cible est présentée, il est possible de prédire (basé sur la phase de l'oscillation reconstruite de ce sujet) si elle va être perçue ou non par un sujet en particulier. Puisque cet algorithme utilise l'EEG reconstruit, une fois que la réponse impulsionnelle a été extraite, il est possible de construire des séquences de bruits blancs avec des cibles spécifiquement présentées pour qu'un sujet puisse la percevoir alors qu'un autre sujet qui regarde exactement la même séquence ne pourra pas la voir. Ceci permet de créer un système de cryptage neuronal, où la relation spécifique entre séquences de bruit blanc et réponses cérébrales (i.e. la réponse impulsionnelle) peut être utilisée comme une clé de cryptage permettant de créer des séquences de bruit blanc spécifiques au sujet pour qu'il perçoive le message intégré alors qu'un autre sujet regardant la même séquence ne verra rien. Le prédicteur basé sur la phase des oscillations était statistiquement significatif, même s'il était moins bon que celui construit dans le chapitre précédent : expliquant tout au mieux 3% de variabilité pour la phase de l'oscillation thêta. Cet effet n'est pas surprenant : le chapitre 2 a démontré que l'effet principal des séquences de bruit blanc sur la détection de cible était indépendant des sujets.

Au travers de la thèse, nous avons utilisé le paradigme de bruit blanc pour questionner la nature de la relation entre oscillations et perception visuelle. Dans un premier temps, nous avons montré que la partie (restreinte) des oscillations présentes dans l'EEG qui est contrainte et modélisée par le paradigme de bruit blanc peut être reliée par sa phase et son amplitude à la détection de cibles présentées à l'intérieur des séquences d'entrainement. La taille des effets de l'état instantané de la phase (11%) et l'amplitude (7%) de l'EEG

reconstruit sur le comportement des sujets est du même ordre de grandeur que celles rapportées dans la littérature (entre 10 et 20%). Dans un deuxième temps, nous avons montré que ces oscillations pouvaient être utilisées pour construire des algorithmes capables de prédire la perception. Ceci suggère que le signal oscillatoire contraint par la séquence de bruit blanc pourrait être à la source de certains effets de la phase et de l'amplitude des oscillations spontanées rapportées dans la littérature. De plus, ces oscillations pourraient refléter le mécanisme d'échantillonnage de l'information visuelle qui donne lieu à des fluctuations rythmiques de la perception visuelle. En effet, les effets de la phase contrainte par la séquence de bruit blanc sur la perception de cible ont été trouvés en post-stimulus, à des latences qui sont plausibles au niveau physiologique. La fonction de ces fluctuations rythmiques reste encore à explorer, et pourrait donner lieu à de nouvelles expériences utilisant le paradigme de bruit blanc.

REMERCIEMENTS

Thank you/Merci …

…Rufin pour ton aide au cours de ces 4 années. Tu étais toujours disponible pour répondre à toutes mes questions scientifiques, techniques, professionnelles… Merci de m'avoir fait confiance, il y a 4 ans, et de m'avoir permis de faire un doctorat. Ça a été très enrichissant de travailler avec toi.

… Leila pour ton aide inestimable avec les analyses de données IRM. Merci de m'avoir fait bénéficier de ton expertise, et d'avoir bien voulu répondre à mes questions.

… Isabelle pour votre disponibilité et flexibilité quant aux inclusions des participants pour l'expérience en IRM.

… Zoé, Claire, Maxime, Carmen, Joël, Damien: c'est grâce à vous que la magie opère. Sans vous pour m'aider à naviguer les eaux troubles de l'administration et les grands mystères de l'informatique, faire un doctorat aurait été encore plus difficile.

… Nathalie, et toute l'équipe du plateau IRM pour votre aide lors des enregistrements en IRM. Ce fut un plaisir de travailler avec une équipe aussi efficace et sympathique. Merci aussi à Yoann pour son calme et les discussions sans fin.

… Théophile, d'avoir été mon premier étudiant. Je suis contente que nous ayons pu apprendre l'un de l'autre.

… Grace, Seb, Rasa, Manu, Doug, and Benedikt. I have pestered each of you at one point or another with my endless questions about EEG recordings, time-frequency analyses and various questions about science and MATLAB. Thank you for taking the time to help me.

… à toute l'équipe "Crème" et PAF, ainsi qu'à tous les membres du CerCo qui ont participé à créer une atmosphère de travail agréable. Merci pour ces discussions intéressantes partagées autour d'un repas ou d'un café.

… to the CogCineL participants. Thank you for the shared movies, laughter and wine. These evenings were mind opening and sharpening.

… Cécile, Mehdi, Nicolas, Kévin, Marie, Damien, Ed, Yseult, Marina, Doug, Manu, Rasa, Seb, Grace, et Samy. Merci pour les moments de détente et les discussions passionnantes. Un merci tout particulier à Cécile et Nicolas, qui ont su doser avec justesse la patience, le courage et la sagesse lors de ces derniers mois. A charge de revanche.

… Pascale & Detlef, Mathias, Elodie & Tobias, Bernadette & Bernard, Martin & Mathilde pour les moments partagés tout au long de ces 4 années. Un merci tout particulier à mes parents pour votre soutien, et pour votre optimisme.

… Simon, d'avoir été à mes côtés pendant toutes ces années. Finalement, à nous deux, on y est arrivé ! Merci de m'avoir épaulée à chaque instant : sans toi, je n'y serais pas arrivée. Merci d'avoir tout pris en charge pour me laisser le temps de me concentrer sur l'écriture de ce manuscrit. Merci pour ta gentillesse et ton soutien… Merci, milles fois MERCI !

TABLE DES MATIERES

Chapter 1. LITERATURE REVIEW

1.1 GENERAL INTRODUCTION

Imagine yourself in the middle of a busy city, about to cross a road to meet up with your friend to go to the cinema. As you prepare to step off the pavement, your brain is receiving a large amount of visual information about your goal (i.e. the cinema), but also about the various other relevant (and irrelevant) visual information such as the incoming traffic, the passer-by picking up his phone, your friend waving hello, the cinema posters… All of this, we experience as a continuous flow of information. Our visual experience seems an uninterrupted stream of sensory impressions: we see cars zip by and our friends waving in continuous motions. However, is it also the case in our brains? Do they constantly process the incoming sensory information? At first glance, this seems a reasonable assumption: our continuous experience arises as a consequence of an uninterrupted processing of the external visual information. However, there is mounting evidence that this might be a misconception.

A simple example can give an intuition for this. When we watch a movie, the actors on the screen appear to be continuously animated, yet, it is made up of 24 distinct images presented in 1 seconds. Despite the sub-sampling of information, we still have the illusion of unbroken movements and flowing actions. Thus, continuous sensory information is not strictly necessary for a continuous sensory experience. This raises the question whether our brains might use a similar mechanism. Could it be that, instead of continuously recording the incoming sensory information, our brains sample the sensory environment in a discrete manner?

This "discrete processing" makes sense from a physiological point of view. Sampling the environment would be a more parsimonious and efficient use of resources, especially if we are unconscious of it. Continuously updating consciousness creates redundant information: visual events are unlikely to change randomly and very rapidly (i.e. at the millisecond scale). Subsampling the visual information would still allow us to extract meaning from a scene, while lowering the energy demands of the neurons. Furthermore, "discrete processing" also makes sense when looking at our brain activity. Ever since the first recording, we have known that our brain activity fluctuates in a rhythmic manner: it oscillates! Could this rhythmicity in

brain activity have a functional role? Do oscillations have a causal influence on our visual perception? Could it be that oscillations are the mechanisms by which the environment is sampled?

These questions were at the heart of my thesis. In this thesis, we will examine how brain oscillations are related to perception and whether they causally influence our visual experience. Our hope is that, once we understand this influence well enough we might use it to take control of visual perception. That is, once we know how oscillations shape what we see, can we find a way to manipulate their state so that we can control what people are conscious of? In fact, I will show that the rhythmic activity in our brains has a direct causal influence on perception and that it is possible to control this rhythmic activity to influence what people see.

I will start with an introduction in three parts. The first part will give a very brief outline of the anatomy of the visual pathways, illustrating the presence of oscillations at all stages of the hierarchy. The second part will give the necessary introduction and definition to the measurement of brain activity and oscillations. This will set up the stage for the third part of this introduction, which will review the evidence showing how oscillations shape what we see. I will conclude this introductory chapter by outlining the general framework used in the remainder of the thesis.

1.2 A BRISK WALK DOWN THE VISUAL PATHWAYS: FROM THE RETINA TO THE VISUAL CORTEX

Our brains are not faithful recording machines, which display to our "mind's eye" an unbiased view of the outside world. Rather, through their very functional and anatomical organization, our brains actively shape what we see.

1.2.1 The retina

As light enters the eye, it falls on the retina, a collection of about 100 million cells, ready to capture the projected image of the world. Once the photoreceptors have transformed the incoming light into an electrical signal through photo-transduction, the information is transferred to the next cells in the hierarchy and travels throughout the visual pathways in the form of an action potential. Cells along the visual pathway form a strict anatomical and functional hierarchy, with a retinotopic organisation which holds true from the retina through to the visual cortex. The retinotopic organization means that events happening in neighbouring areas of the visual field are processed by neighbouring cells in the brain (see Figure 1-1). As mentioned before, oscillations are everywhere in our brains: already at the level of the retina, ganglion cells start to oscillate: both spontaneous and stimulus driven oscillations in a variety of frequencies have been recorded (Neuenschwander, Castelo-Branco, & Singer, 1999).

As information flows in the system and is fed forward to different areas, a hierarchical organisation is created. Each cell at a given stage of the hierarchy receives and aggregates information from cells one level lower (see the representation of this hierarchy on Figure 1-1). Consequently, as we move up the hierarchy, cells respond to more and more complex stimuli (Livingstone & Hubel, 1988) and respond to larger patches of the external world (i.e. they have larger receptive fields).

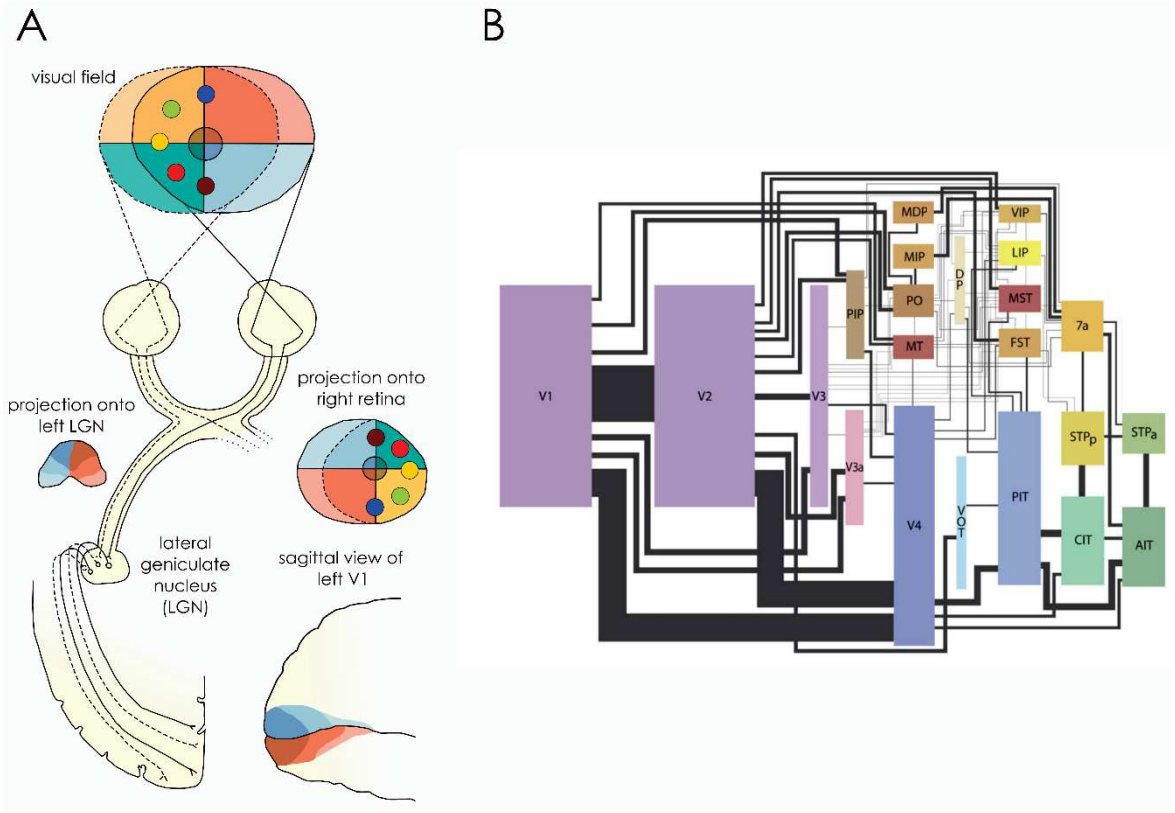

Figure 1-1. Visual pathways and retinotopic mapping **A**. Illustration of how the retinotopic mapping is preserved from the retina until the visual cortex. The coloured points and shaded areas in the visual field illustrate that the information is flipped up down and that each hemisphere processes information about the contralateral hemi-field. **B.** Representation of the cortical visual areas (coloured boxes) and the connections (lines) between them in the macaque. Warm colours (red and browns) highlight areas belonging to the dorsal stream and cold colour (blue and greens) represent areas of the ventral stream. The size of boxes are scaled depending on the cortical surface area estimated in the literature. The thickness of the lines represents an estimate of the relative number of fibres connecting the different brain areas based on their relative sizes. Reproduced from Wallisch & Movshon (2008)

1.2.2 Parallel processing pathways

Based on their response properties, specific types of ganglion cells in the retina create two main pathways of information transfer to the visual cortex: one subtending the processing of depth, motion information and the other one of colour and object recognition, which run in parallel from the retina to the cortex (Livingstone & Hubel, 1988). Of course, this is a generalisation, there are many more smaller streams that transfer information from the retina to the brain (Wässle, 2004): for example, some projections directly go from the retina to the superior colliculus, a small region in the centre of the brain. The superior colliculus contains

sensory maps and has been involved in the control of eye movements and spatial attention (Krauzlis, Lovejoy, & Zénon, 2013). Here however, we will only focus on the visual pathway subtending conscious visual perception, which goes through the lateral geniculate nucleus in the thalamus. The information is transferred via the optic nerve and the optic chiasm, which splits and "sorts" information: neurons processing information about the right visual field are sent to the left hemisphere, while the information about the left visual field is sent to the right hemisphere (see Figure 1-1).

1.2.3 Thalamus

Sitting at the centre of the brain, the thalamus is the entry point for virtually all sensory information en route to the cortex. The parallel pathways emerging from the retina project to different layers of the dorsal lateral geniculate nucleus of the thalamus. The retinotopic organisation still holds true and the pathways stay segregated. Until recently, it was thought that its role was merely that of a relay centre, transmitting information to various cortical regions. This notion arose mainly from the virtually unchanged response properties of the receptive fields in the lateral geniculate nucleus (Sherman, 2005). However, recent evidence suggests that it plays an active role in shaping the processing of incoming sensory information through the modulatory role of feedback connections from brain areas higher in the visual processing hierarchy. It is thought that these cortico-thalamic feedback connections create spatio-temporal windows of enhanced processing (Cudeiro & Sillito, 2006), which modulate the shape and firing strength of specific population of LGN cells (Briggs & Usrey, 2008). Another important role of the thalamus is that of gating incoming information (Steriade & Llinas, 1988): during sleeping state, the thalamic neurons are characterized by synchronized activity which inhibits the transfer of information to the cortex (e.g. McCormick & Bal, 1994; Steriade, McCormick, & Sejnowski, 1993). Finally, as we will see later, it has been implicated in the genesis of many cortical rhythms.

1.2.4 Primary visual cortex

The visual information then continues its journey in parallel streams from the lateral geniculate nucleus to the cortex and the primary (striate) visual area. Also known as V1, it is located in the occipital areas at the back of our heads. It spreads around the calcarine fissure (see Figure 1-1), on the inside of each hemisphere.

The architecture of V1 is very complex: it is composed of 6 distinct layers of cells with different functions roles, organisation reflecting that of the whole cortical architecture (Douglas & Martin, 2004). Because it is not directly relevant to the topic of this thesis, the exact architectural organisation of V1 will stay unaddressed, however, I will highlight two important points. First, the feed-forward parallel pathways carrying information from the thalamus arrive in layer 4 of V1 (yet they project to segregated patches of cells). In fact, feedback connections project to different layers, targeting superficial and deep layers (Felleman & Van Essen, 1991). This is relevant insofar as it has been proposed that distinct types of oscillations might underlie the feed-back and feed-forward connections (van Kerkoerle et al., 2014). Higher oscillatory frequencies (beta and gamma band) have been proposed to carry the feed-forward information while feed-back information from higher areas are reflected in the activity of lower neural oscillations (such as theta and alpha band; Bastos et al., 2015; Michalareas et al., 2016; van Kerkoerle et al., 2014). This augurs the possibility for oscillatory activity to be used as a probe of the functional organisation of the cortex, using this difference of oscillatory frequency as a marker for different types of connections. Secondly, it bears mentioning that connections from other brains areas (in the form of feed-forward and feed-back connection) only make up a very small amount of connections in the cortex: the majority of connections within a layer comes from the recurrent (i.e. local or "lateral") connections (Douglas & Martin, 2012).

As for the functional organisation of V1, the retinotopic organisation is kept here too: cells located near each other on the retina will project to neighbouring cells in V1. This leads to an increased complexity in the properties the cells respond to. The ganglion cells transform the pixel-like information of photoreceptors with the help of various other intermediate cells (Priebe & Ferster, 2012) to an antagonist centre-surround organisation (Kuffler, 1953). V1 simple cell integrates the information from several receptive fields and respond preferentially to stimuli with a given orientation, shape and position (Hubel & Wiesel, 1959). This has fostered the idea that these cells are feature detector, with RF whose shape is a Gaussian filter (Spillmann, 2014). The increasing complexity of the properties cells respond to along the hierarchy is a key mechanism of the visual pathway.

1.2.5 Organisation of higher visual areas

From V1, the information is transferred to V2 and then V3 through feed-forward connections. At each step in the hierarchy, the information is transformed and processed in increased levels of complexity, each brain regions responding to more and more complex properties of the objects in the environment. It has been proposed that the elaboration of these properties is split between the ventral and the dorsal pathways, respectively specialized for the "what" and the "how"/"where" of vision (Livingstone & Hubel, 1988; Milner & Goodale, 2008; Ungerleider & Mishkin, 1982). Our "vision-for-perception" (i.e. the "what") is thought to be mediated by the ventral pathway, which processes all information about shapes, colours, and object identity. It goes from the early visual areas to temporal regions (such as the antero-inferior temporal cortex or AIT, Figure 1-1). On the other side of the brain, the dorsal pathway goes from the primary visual areas to higher visual areas located in the parietal cortex, ending up in motion specific centres (such as MT, the LIP and MST, see Figure 1-1). It is thought to underpin our ability to use vision for action (i.e. the "how"/"where" of vision), computing information about depth and movement to prepare the relevant motor programs. From then on, the information travels further on to the associative, motor and frontal cortices and is used eventually for action.

We have now finished our brisk walk down the visual pathway, which hopefully, has painted a (admittedly very rudimentary) picture of the "where?" of the visual brain. We will now turn to what goes on inside these anatomical structures.

1.3 INTRODUCTION TO OSCILLATIONS AND EEG

As mentioned in the general introduction, oscillations are ubiquitous in our brains and can be recorded at many levels. In this section, I will focus on brushing a broad picture of what oscillations are, giving a general introduction to their analysis and classification and how they can be recorded using electro-encephalography.

Hans Berger is commonly accepted as the forefather of electro-encephalography (EEG). In 1929, he described the electrical potentials he recorded from the scalp of, amongst others, his young son. In this seminal paper, he described two oscillatory signals: the first was a relatively slow oscillation (about 10 cycles per seconds), strongest at the back of the head and when subjects closed their eyes, which he called the alpha rhythm (Berger, 1929). The second was faster and only appeared when subjects had their eyes opened, and which he termed the beta rhythm (Berger, 1929).

1.3.1 What is an oscillation? Phase, amplitude and frequency

While some oscillations are strong enough to be visible with the naked eye (see Figure 1-2), we have seen that some oscillatory activity require the use of time-frequency decompositions (see Figure 1-2). An oscillation can be described using 3 characteristics (see Figure 1-2): its speed or frequency (in Hertz), its strength or amplitude (usually in microvolts), and its position in the cycle or phase (in radians or degrees). These three parameters are not independent from one another. For example, the phase of an oscillation is inherently linked to its frequency: in fact, the phase's instantaneous time course has recently been used to retrieve frequency changes in time (Cohen, 2014). Moreover, as the frequency of the oscillation increases, its strength (power) decreases. The signal that is recorded through EEG has a typical '1/f' power profile: lower frequencies have more amplitude (Roopun et al., 2008).

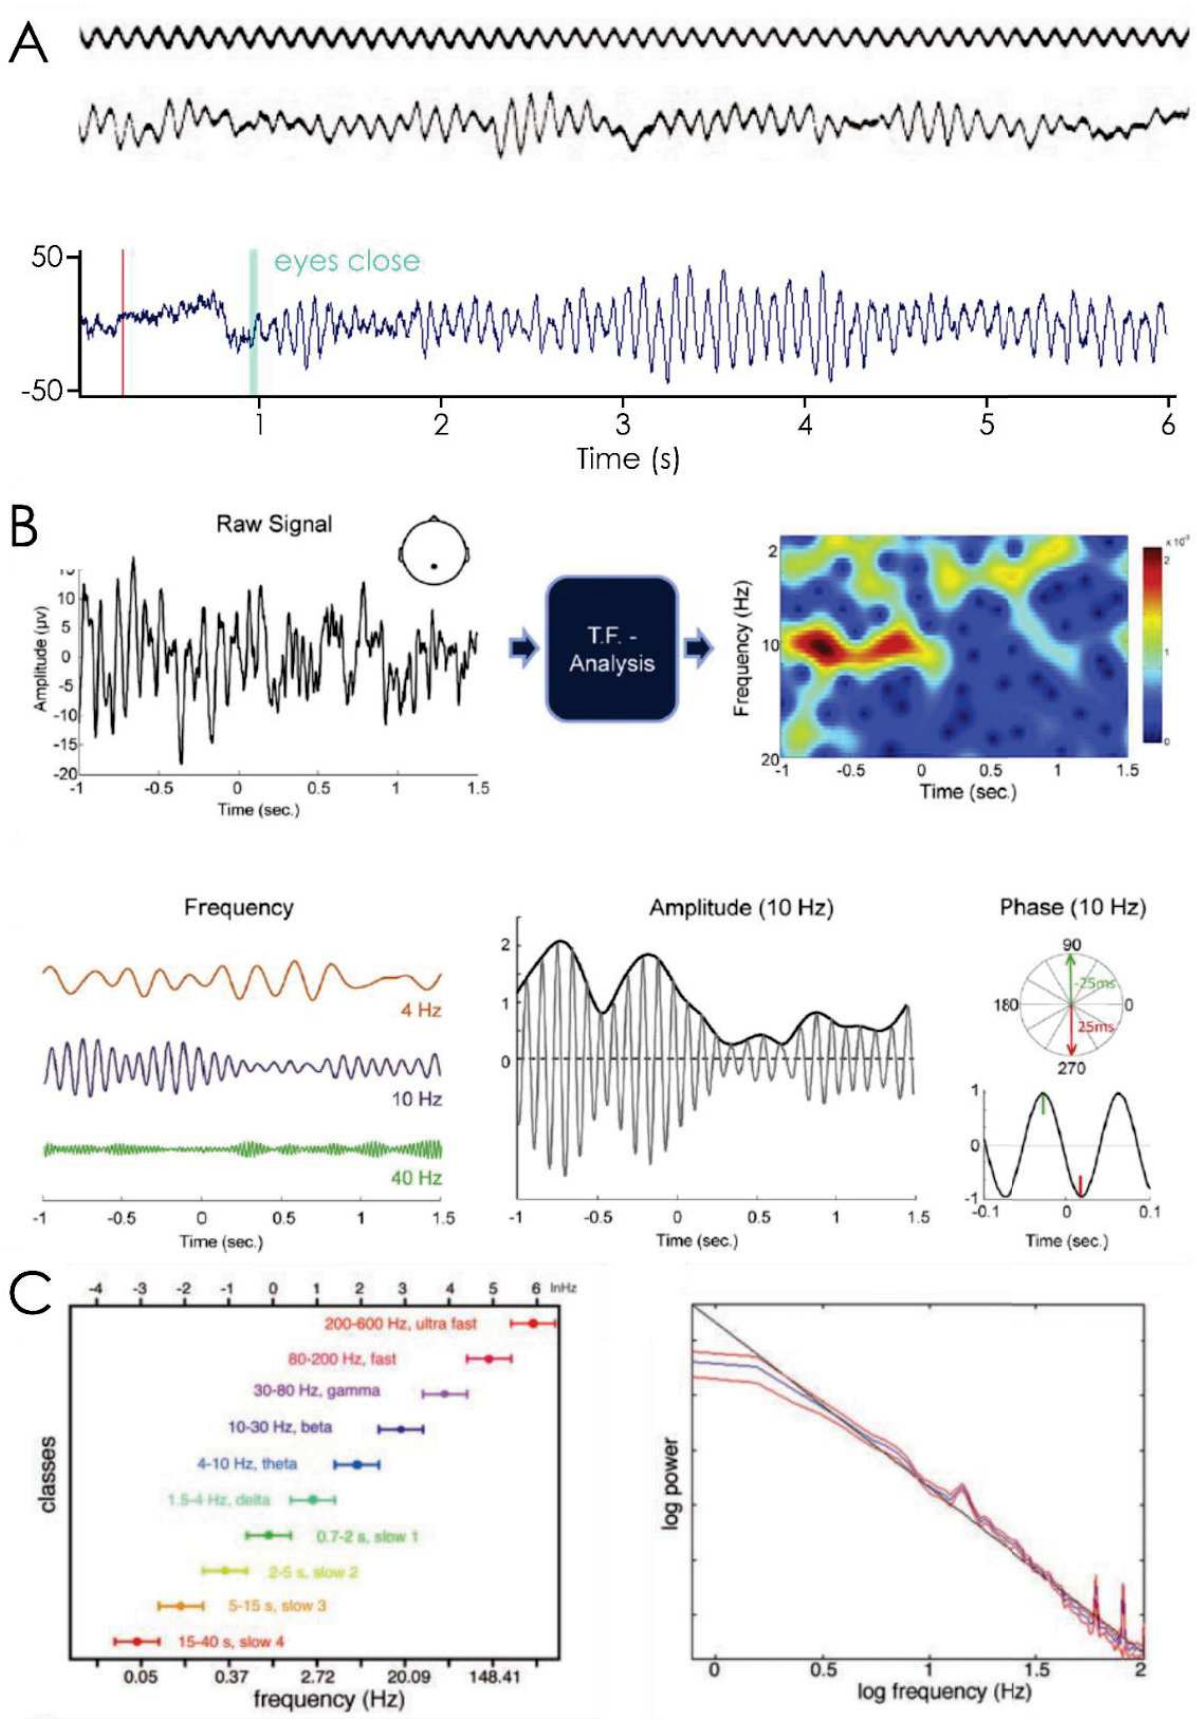

Figure 1-2. Illustration of spontaneous brain oscillations in the raw EEG signal (**A**) and how to extract their spectral component (**C**) using time frequency information (**B**). **A**. Examples of spontaneous alpha activity: recorded by Berger with a 10 Hz trace (respectively

top and middle trave, reproduced from Millett, 2001) and from one of my own subjects on an occipital channel (Oz). **B**. The "raw" EEG signal can be decomposed into its oscillatory components by means of a time frequency analysis. These oscillations vary in 1) frequency (i.e. its speed usually expressed in hertz, or number of cycles per seconds), 2) amplitude or power (i.e. its strength), phase (i.e. instantaneous state of the signal in time, usually represented in angles or radian). The phase of a signal is often represented in a circular plot: here two phases of the alpha oscillation are represented (reproduced from Hanslmayr, Gross, Klimesch, & Shapiro, 2011). **C**. **Left** – Oscillations can be classified depending on their frequency: here, 10 separate frequency bands are created based on the observation of the naturally occurring frequency bands in animal research and computations. **Right** – Power spectrum of the human EEG follows the typical 1/f shape (reproduced from Buzsáki & Draguhn, 2004).

Oscillations have traditionally been grouped in bands depending on their frequencies: from the delta oscillations (δ, 0.1 Hz to 4Hz), theta (θ, 4 to 8 Hz), alpha (α, 8 to 12 Hz), beta (β, 12 to 30 Hz), to the gamma (γ, 30 Hz and plus) oscillations. The specific bounds of frequency bands tend to be highly dependent on the specific field of study (i.e. different bands in animals and in sleep research for example). Buzsaki and colleagues (Buzsáki & Draguhn, 2004; Penttonen & Buzsáki, 2003) offered a classification based on results from animal research which showed the presence of at least 3 different frequency bands to uncover functionally separate frequencies bands. Using computations, they hypothesized the presence of at least 10 independent frequency bands (Penttonen & Buzsáki, 2003), 5 of which fall within the traditional bands of EEG recording (from 1 to about 100 Hz). These frequency bands range from the very slow to the ultra-fast (see Figure 1-2) and are linearly-spaced in the natural 'e' log scale (Buzsáki & Draguhn, 2004; Penttonen & Buzsáki, 2003; however also see Roopun et al., 2008).

Each of these frequency bands has been linked with specific cognitive processes which have given topographical representations. Slower oscillations synchronize over large spatial scale (Buzsáki & Draguhn, 2004), for example delta oscillations have been linked with large-scale cortical integration and language processes (Sauseng & Klimesch, 2008). Theta oscillations have been linked to memory (Lisman & Idiart, 1995) and executive functions underlying control and inhibition (e.g. Cohen & Donner, 2013; Gulbinaite, van Rijn, & Cohen, 2014). Alpha oscillations have been linked with various cognitive functions (some of which will be reviewed later), such as memory (e.g. Klimesch, 1999) spatial attention (e.g. Foxe & Snyder, 2011), inhibition (e.g. Jensen & Mazaheri, 2010). Faster oscillations such as

beta and gamma are likely to be generated in the cortex and synchronize over smaller brain areas (Sauseng & Klimesch, 2008). Beta oscillations have been involved in somato-sensory functions (e.g. Baumgarten, Schnitzler, & Lange, 2015; Neuper & Pfurtscheller, 2001). Finally, gamma oscillations have been linked to a wide range of mechanisms, from visual awareness (Engel & Singer, 2001), to neuronal communication (Fries, 2015) and object representation (Catherine Tallon-Baudry & Bertrand, 1999). This is a very succinct, biased and incomplete review of the wealth and breadth of cognitive functions which have been linked with the various brain oscillations. Very good reviews have been written on this subject, and it is beyond the scope of this thesis to attempt such a feat (see amongst many others, the following papers for reviews Herrmann, Strüber, Helfrich, & Engel, 2016; Sauseng & Klimesch, 2008; Wang, 2010). Crucially, here I want to hint at the richness of evidence for the importance of these brain oscillations. In fact, their central role in cognition has been emphasized by looking at the "broken brain" (Başar et al., 2016): specific disorders (such as schizophrenia, attention deficit or bipolar disorders) have marked alterations in oscillations during rest and on a range of cognitive functions compared to controls (see Başar & Güntekin, 2013 for a review).

1.3.2 What are we measuring?

Brain oscillations arise from the electrical activity that can be recorded outside of neurons in the extra-cellular space (Buzsáki, Anastassiou, & Koch, 2012). Depending on the scale at which these signals are recorded, they get different names, from local field potentials (LFPs) when intra-cranial electrodes are used, to electro-corticogram (ECoG) when using subdural grids and EEG when scalp electrodes are used (Thut et al., 2012).

EEG oscillations are the result of the collected activity from thousands or more pyramidal cells: the excitatory input (i.e. the action potential) to these cells create a positive ion flow inside the cell, depolarizing their membranes, and thus creating a local negativity in the extracellular space (Luck, 2005b, Chapter 1). This opposition between local negativity in the post-synaptic dendrites and local positivity in the axons creates a local dipole (Luck, 2005b). Crucially, only the synchronized activity of large, aligned populations of cells can be recorded by electrodes on the scalp level (Steriade, Gloor, Llinás, Lopes da Silva, & Mesulam, 1990).

As mentioned, oscillations reflect the collected activity of many cells at the same time. As such, the amplitude of the oscillation is proportional to the number of cells that are active (Pfurtscheller & Lopes Da Silva, 1999). Furthermore, the instantaneous state of the oscillation reflects changes in excitability of the underlying neuronal populations (see Figure 1-3A): the phases of oscillations in various frequency bands have been shown to correlate with the underlying spiking activity at various spatial and temporal scale and in various sensory modalities (Fries, Nikolić, & Singer, 2007; Haegens et al., 2015; Haegens, Nácher, Luna, Romo, & Jensen, 2011; Jacobs, Kahana, Ekstrom, & Fried, 2007; Lakatos et al., 2005; Montemurro, Rasch, Murayama, Logothetis, & Panzeri, 2008). For example, recordings in the human brain through intra-cranial electrodes have shown that the phase of the LFP ongoing oscillation in the theta band influenced the spiking activity: it was increased by two fold at the "preferred" phase (Jacobs et al., 2007). In fact, this link has also been shown with scalp recordings in monkeys: spiking activity (from multi-unit recordings) was strongest when the gamma amplitude was high at the descending delta phase (Whittingstall & Logothetis, 2009).

The fact that spikes are more likely to occur at a given phase of the oscillation has been proposed as a mechanism through which effective communication between neuronal populations is achieved (Fries, 2005, 2015). Through phase synchronization in the gamma band, two neurons can synchronize the moment of peak excitation leading to an optimized information exchange (see red and blue neuron on Figure 1-3). On the other hand, a neuron can become "deaf" to incoming information from another cell by synchronizing its inhibitory phase to the other cell's excitatory phase (represented by the blue and black cells on Figure 1-3). This communication through coherence allows for the creation of flexible effective communication architecture from the anatomical connections (Fries, 2005).

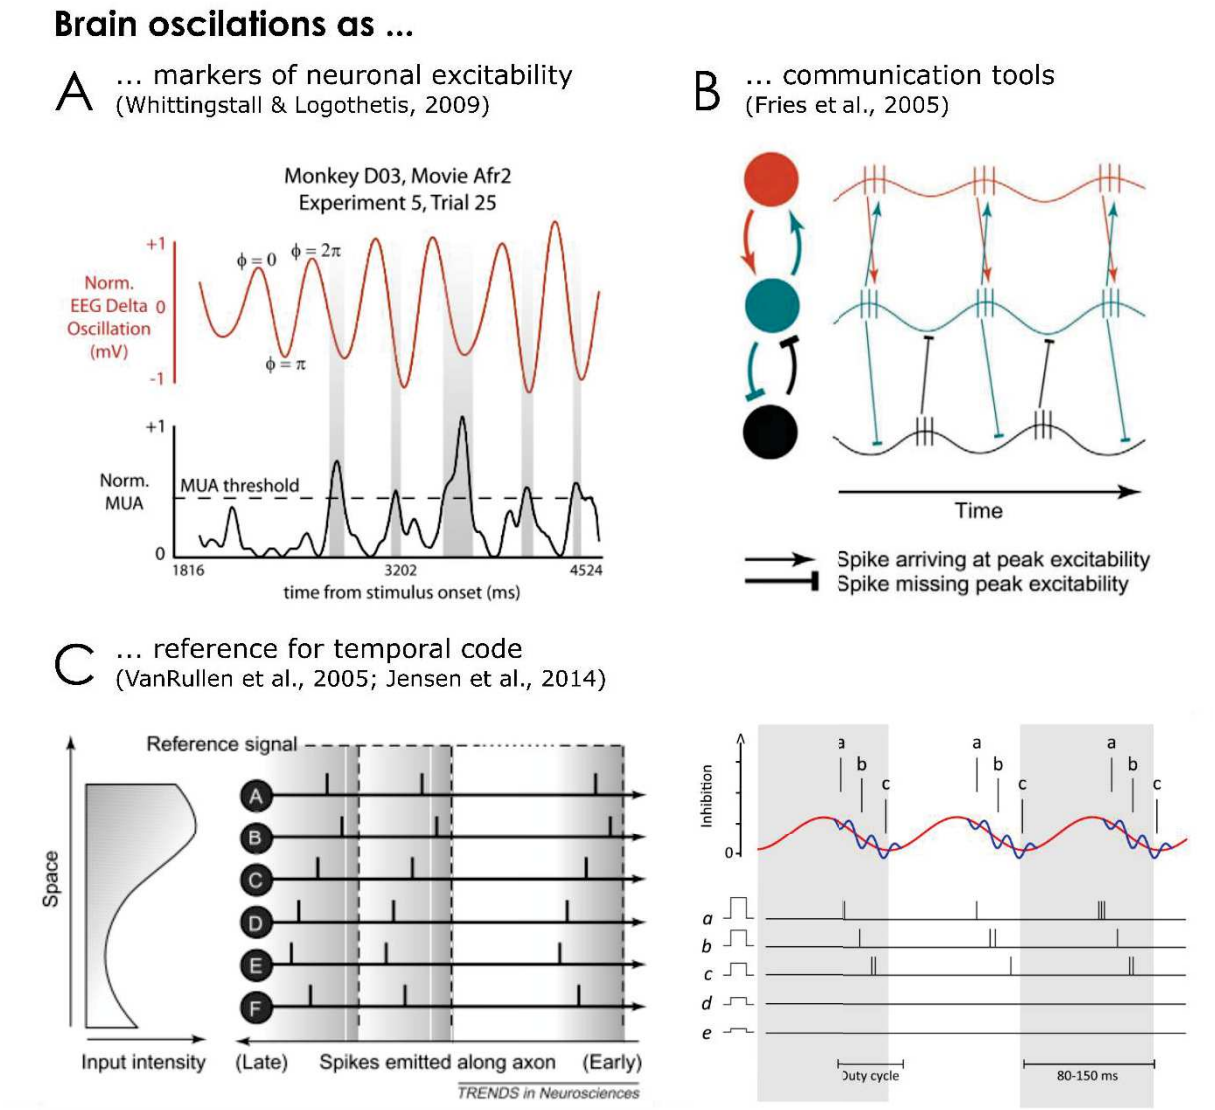

Figure 1-3. Illustration of a few functions in which oscillations have been involved. **A**. Both the phase and the amplitude of neuronal oscillations have been shown to reflect activity of the neuronal population. Here the delta phase of EEG recording is linked with population spiking activity measured through multi-unit activity (MUA) during a movie (Whittingstall & Logothetis, 2009). **B.** At the heart of the Communication through Coherence Theory, the synchronization of the excitatory phase of neuronal ensembles allows for the exchange of information (e.g. between the red and blue neuronal group), while inhibition is achieved by synchronizing the phase so that the lowest excitability phase is present when information comes in (e.g. between the blue and black neuronal population, reproduced from Fries, 2005). **C.** Neuronal oscillations could act as a reference frame for the creation of a temporal code. Left – information could be coded in the relative timing of spikes. For example, stimulus intensity could be coded in the timing of the discharge: strongly activated neurons fire at shorter latency leading to that stimulus to be processed faster (reproduced from VanRullen, Guyonneau, & Thorpe, 2005). Right - The reference for this relative timing could be the alpha oscillations, which imposes (as will be described later) a gated-inhibition on neuronal firing. As phase changes, inhibition is released and the most excitable presentation, a, will discharge first, followed by b and c. While d and e won't be activated due to the phasic inhibition. Here, the model supposes that these representations, coded within the gamma oscillatory activity

(nested within the alpha oscillation) will keep the stimuli's representation separated through the inter-neuron inhibitory network (Jensen, Gips, Bergmann, & Bonnefond, 2014).

In fact, this link between neuronal excitability and oscillatory phase could be instrumental in brain function: oscillations could act as a reference signal for spikes. This question is germane to how neurons encode information: rather than the quantity of spikes (i.e. rate coding first proposed by Adrian, 1928), the temporal coding theory holds that the time of spike itself contains information (Thorpe, 1990). Oscillations could provide the necessary reference mechanism through which spike time is judged (Fries et al., 2007; Panzeri, Brunel, Logothetis, & Kayser, 2010; Singer, 1999; VanRullen et al., 2005). In fact, evidence supports this idea: using the time of spiking with regards to the phase of the lower frequency LFP (between 1 and 4 Hz) could explain up to 54% more information than the spike rate alone (Montemurro et al., 2008). The strongest evidence for this type of coding comes from studies of the "theta phase processing" in the rat hippocampus: the activity of place cells (neurons coding for specific locations in space) will be activated at earlier time of the LFP theta cycle as the animal moves in its environment (Burgess & O'Keefe, 2011; Mehta, Lee, & Wilson, 2002). This temporal code segregates information based on various stimuli property (e.g. intensity, see Figure 1-3C): it can be used to achieve several functions such as sequential processing, figure-ground contrast, visual contrast or even in higher cognitive functions such as attentional allocation (Jensen, Bonnefond, & VanRullen, 2012; Jensen et al., 2014) and working memory (Lisman & Jensen, 2013).

1.3.3 Interaction between frequency bands

These various frequencies are not always independent. It has been suggested that neighbouring frequencies compete with each other (Buzsáki & Draguhn, 2004). One way in which these frequencies might be kept separate is that they are generated by different networks and the fact that they never fully synchronize (Roopun et al., 2008). On the other hand, synchronization could be a mechanism through which two or more rhythms could actually co-exist (spatially and temporally), interacting with one another through the cross-frequency coupling (see Figure 1-4). Various types of interaction between the two signals can

be imagined (see Figure 1-4): the most commonly reported is the phase to power modulation whereby the phase of the slower oscillation modulates the power of the faster oscillation.

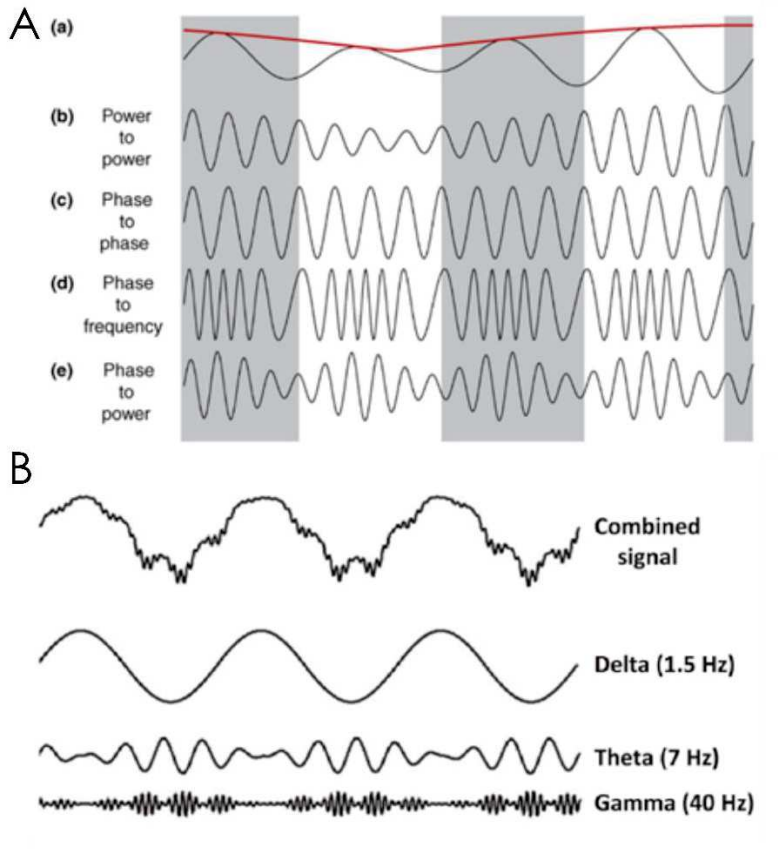

Figure 1-4. Cross-Frequency coupling. **A.** With regards to the slower (8 Hz) reference oscillation (a), different parameters of the faster oscillations (in the gamma band) could be modulated. These modulations could occur together (they are not mutually exclusive). **(b)** The power of both oscillations could be co-modulated. **(c)** The two oscillations could be phase-locked: here there are systematically 4 cycles of the faster oscillation within one cycle of the slower one. **(d)** The phase of the slower oscillation could influence the frequency of the faster oscillation. **(e)** The phase of the slower oscillation could influence the power of the faster oscillation. (reproduced from Jensen & Colgin, 2007) **B.** Example of phase/amplitude coupling the combined signal (top) is actually composed of 3 oscillations at different frequencies. The amplitude of the gamma oscillations is modulated by the phase of the theta band oscillation, whose amplitude is itself modulated by the phase of the delta oscillation. (Reproduced from Calderone, Lakatos, Butler, & Castellanos, 2014)

Thus cross-frequency coupling allows the transfer of more information by multiplexing (i.e. grouping together) information into one simple combined signal. This mechanism has been implicated in many cognitive functions such as memory (for a review see Canolty & Knight, 2010; Watrous, Fell, Ekstrom, & Axmacher, 2015): where, as mentioned earlier, theta oscillations are thought to provide a reference mechanism for gamma

oscillation in the hippocampus (Lisman & Jensen, 2013). Another example of the presence of cross-frequency coupling comes from the coding of facial expressions. Schyns and colleagues (2011) computed the quantity of information in each of the phase, power and both taken together. They found that more information was present in the phase of the oscillation (compared to the power alone), and that the combination of low beta (12 Hz) power and the theta (4Hz) phase codes information about specific facial features, thus "increases coding capacity in the brain" (Schyns, Thut, & Gross, 2011).

1.3.4 What mechanisms give rise to oscillations in the brain?

Finally, we will briefly turn to the mechanisms giving rise to neural oscillations. Some neurons have inherent (membrane and synaptic) properties that make them oscillate (Buzsáki & Draguhn, 2004). In vitro, thalamic neurons have been shown to spontaneously oscillate at a frequency of 6 to 10 Hz (Lopes da Silva, 1991). Moreover, oscillations are likely to occur as a result of the properties of the circuit they are embedded in (Buzsáki & Draguhn, 2004). For example, the pyramidal-interneuron network gamma model suggests that the gamma oscillations are a result of the interconnection between local excitatory pyramidal cells, which when they fire, activate inhibitory interneurons (Fries et al., 2007; Tiesinga & Sejnowski, 2009). Long range connections are also important: the interplay between feed-back and feed-forward connections coupled with fluctuating inhibition promotes the emergence of oscillations (Lopes da Silva, 2013). Although the genesis of the alpha rhythm is still a matter of debate, it is thought to result from both thalamo-cortical and cortico-cortical loops (Bollimunta, Chen, Schroeder, & Ding, 2008; Foxe & Snyder, 2011; Lopes da Silva, Vos, Mooibroek, & van Rotterdam, 1980; Lorincz, Kékesi, Juhász, Crunelli, & Hughes, 2009).

1.3.5 Evoked, induced, ongoing … different signals in the EEG

In the literature, different types of signals have been described in the recorded EEG depending on when they arose with regards to the presentation of the target. First, there are the signals that arise as a function of visual stimulation: the evoked and induced responses.

The **evoked responses** are any signal whose phases are locked to the presentation of the stimulus such as the event related potential or the ERP (or the visually evoked potential, or

VEP[1]) or any oscillatory component. Systematically after the target is presented, these evoked responses have the same phase across trials and can thus be extracted by a simple averaging of the signal in the time domain (see Figure 1-5B). Whether the VEP reflect a resetting of the ongoing oscillation or an additional signal superimposed on the ongoing signal is still a matter of debate (e.g. Becker, Ritter, & Villringer, 2008; Makeig et al., 2002; Mazaheri & Jensen, 2010), but as this debate is not directly addressed in this thesis we will not go into any more detail. Targets can also give rise to not phase-locked activity: the **induced signals**. Their exact latency can vary from trial to trial (Catherine Tallon-Baudry & Bertrand, 1999). Therefore time-frequency transforms are needed to uncover their time course: time based averaging will cancel them out (see Figure 1-5B). These induced oscillations have been linked with various cognitive functions, for example gamma band induced oscillations are thought to reflect object representations (see Tallon-baudry, Bertrand, Delpuech, & Pernier, 1996 for a review), processing of incoming sensory information (Baldauf & Desimone, 2014) and awareness (Wyart & Tallon-baudry, 2008)

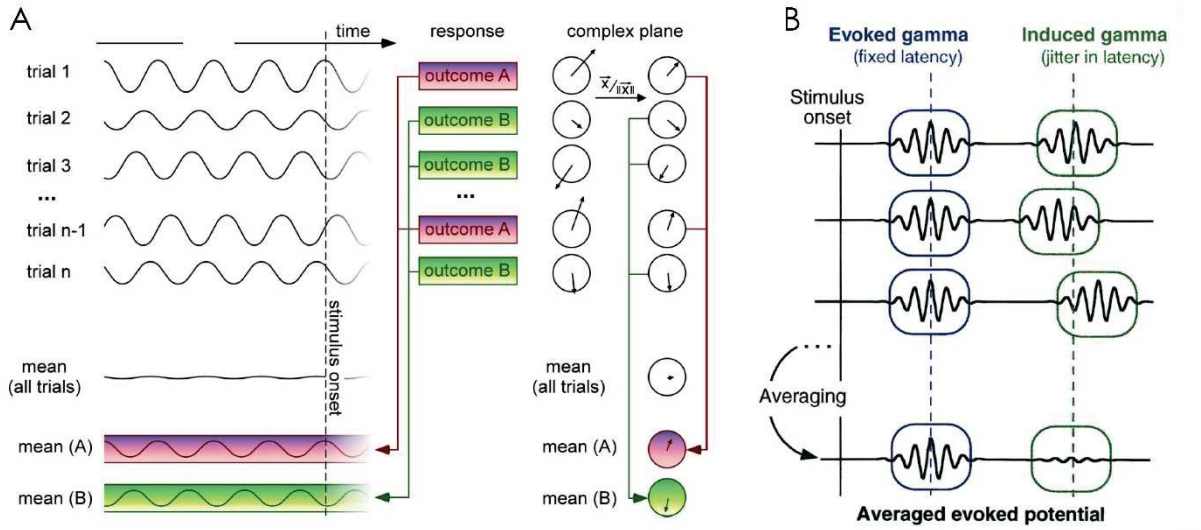

Figure 1-5. Example of evoked, induced and ongoing oscillations. **A.** On each trial, the phase of ongoing pre-stimulus signals causally influence visual perception (see section 1.5). If the simple average across all trials is taken, these systematic relationship cancel each other out. However, the phase relationship can be extracted by computing the inter-trial phase coherence for each condition separately.(adapted from VanRullen, Busch, Drewes, & Dubois, 2011) **B**. Examples of an evoked and induced oscillatory activity in response to a stimulus. When averaged in the time domain, the evoked (phase locked) response is recovered while the induced response is cancelled out. To extract the induced gamma oscillation, a time-frequency transform before averaging (reproduced from Tallon-Baudry & Bertrand, 1999).

[1] In this thesis, evoked responses and event related potentials will both be used indiscriminately to refer to the activity specifically evoked by a stimulus or related to its processing.

Finally, there are signals which do not directly reflect the processing of the presented target: these have been called **ongoing (or spontaneous) oscillations**. Because of their small amplitude (relative to the evoked response), these oscillations are more easily visible (in the EEG) before stimulus onset (see Figure 1-5A). In recent years, they have been linked with visual outcome (see section 1.5 below). These signals can occur in the absence of sensory inputs, however, it has been suggested that their state can be modified by stimulation (Nunez & Srinivasan, 2006, p. 10). Steriade and colleagues call "spontaneous oscillations" a sort of tonic background activity that occurs after evoked/induced oscillations and is linked with brain excitability (Steriade et al., 1993). We will come back to this at the end of the introduction to set up the general framework of the thesis (see section 1.6).

1.4 EVENT RELATED POTENTIALS AND VISUALLY EVOKED RESPONSES

This (and the next) section evaluate how brain activity has been linked to visual perception in the literature. Let us first look at the evoked responses and how they have been used to evaluate the time course of visual perception.

1.4.1 Extracting the visually evoked potentials (VEP)

One of the most common approach to understanding how our brains see the world has been to extract the visually evoked potentials (VEP). It represents the stereotypical response of the brain to the presentation of the stimulus. The VEP can give information about the unfurling of information processing in the brain with an excellent temporal resolution (on the order of the millisecond). This approach has been extensively used to study attention for example (Luck, 2005a, pp. 75–95; Luck et al., 2000; Woodman, 2010) as it allows for the careful evaluation of the time course of cognitive functions and specific sensory processes (Hillyard et al., 1998).

The VEP is characterized by a series of positive and negative deflections following a stereotypical pattern (see Figure 1-6b) named after their peak latency and the sign of the deflection (Luck, 2005a). First, there is the C1, which starts early, from about 60 to 80 milliseconds (ms), followed by the first positive component (P1) from about 80ms to 120ms and then finally the first negative component (N1), occurring from about 120ms to 180ms

(Regan, 1966). Interestingly, the exact shape of the evoked response is very variable across subjects (Ciganek, 1969). Various other components have been reported in the literature, which are thought to reflect specific cognitive functions such as the selection negativity or the N170 linked with stimulus categorization (Hillyard et al., 1998).

In traditional paradigms, single stimuli (e.g. single flashes of light at various intensity) are presented in isolation while the brain activity is recorded using EEG (see Figure 1-6a). The VEP is extracted by creating epochs of data around each target and averaging across trials (Dawson, 1954) in order to reduce both the neural and the recording noise (Luck, 2005b).

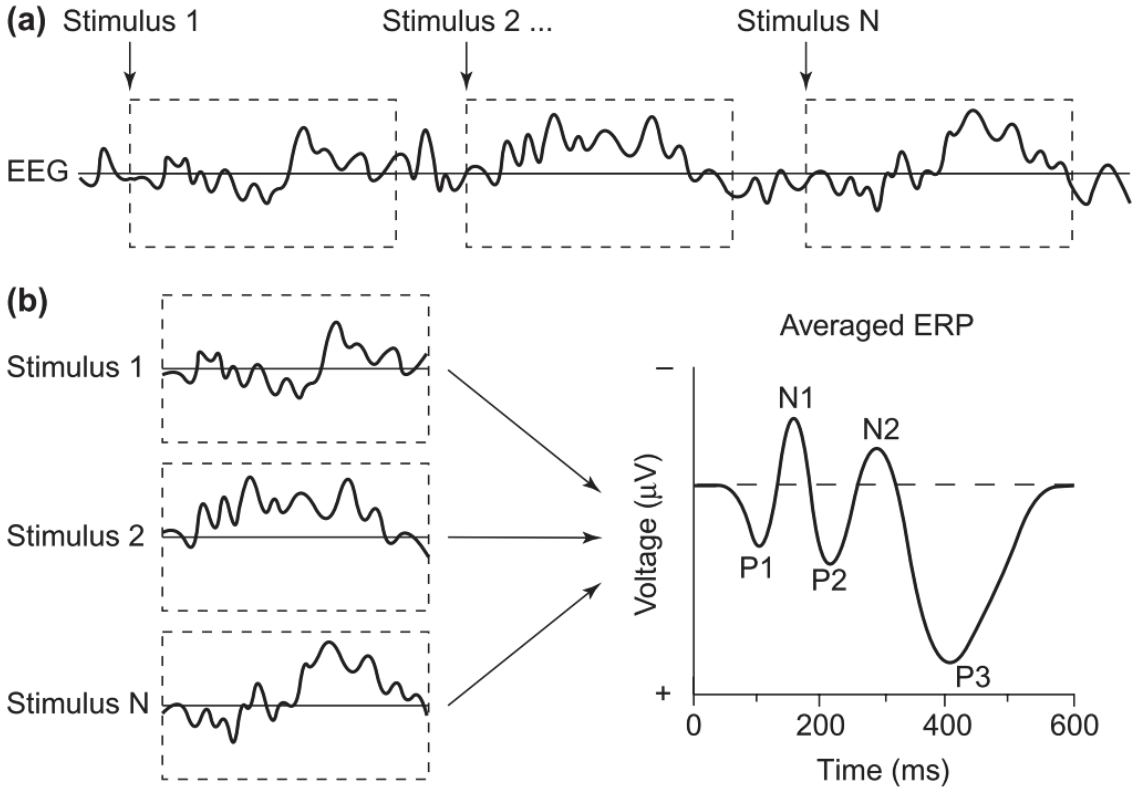

Figure 1-6. Extraction of the ERP waveform from the ongoing EEG signal. **A.** Stimuli are presented in isolation while the ongoing EEG is recorded. The inter-stimulus interval is long enough for the evoked responses to return to baseline. The event related potentials are too small to be seen from the single trial data. **B.** The ERP is extracted from the ongoing EEG by averaging the trials together (reproduced from Luck, Woodman, & Vogel, 2000)

However, the benefit of increasing the signal to noise ratio comes with a cost: ERP studies are usually long and boring. Many trials are needed for a good signal to noise ratio (although the amount of noise is not a direct linear function of the number of trials) and

stimuli need to be separated by at least 500ms to allow the previous activity to die down (Luck, 2005, p. 99).

1.4.2 Evoked responses as Impulse Response Functions of the brain

While the "traditional" method of extracting the ERP (described above) is useful in understanding how our brains process visual information, it is too time consuming. An alternative method has come from the analysis of linear systems and in particular from methods used in system identification. It has been used in the field of electrophysiological recordings to extract the response properties of cells (Marmarelis, 1980; Ringach & Shapley, 2004; Sakai, Naka, & Korenberg, 1988).

As an example, let us take a very simple linear system, with an answer (output) to a stimulus (input) as a function of some internal transformation (i.e. its impulse response; see Figure 1-7A). Here, the systematic relationship between input and output is called the impulse response: it represents how the system responds to one input (i.e. an impulse). If we simplify, what we are trying to do when extracting the VEP, is trying to extract the impulse response of the brain.

If we want to characterize how this linear system responds to an input, we can use two approaches (Marmarelis & Marmarelis, 1978).

First, isolated stimuli can be repeatedly presented (i.e. traditional method described above, see Figure 1-7A). However, this approach has limitations. As mentioned, it is very time consuming. Moreover, it is not very ecological: single stimuli are rarely presented in isolation in the external world, rather information changes at various time scales. Finally, problems arise when we consider non-linear systems (such as our brains). In such a case, presenting a single stimulus will only tell us how the system responds to that particular stimulus: it could be that a single flash will produce one type of response, but a sinusoidal stimuli might give rise to yet another response. Thus, many different types of stimuli need to be presented before generalization can be made about how the non-linear system responds to inputs (P. Z. Marmarelis & Marmarelis, 1978). In a simple linear system, the impulse response can be fully extracted from a single pulse: no information can be gained by presenting different types of pulses since the behaviour is wholly linear. In non-linear systems however, this is not the case.

Thus, the alternative method is to present many stimuli at (nearly) the same time by using White noise (WN, i.e. random luminance) sequences as a stimulus: they have a very rich spectral content that stimulates the brain at all frequencies within a given frequency range (P. Z. Marmarelis & Marmarelis, 1978). The advantage of presenting WN stimuli becomes readily apparent in this case: they can drive the brain of the subjects with many different types of stimuli in a very short amount of time thus "maximiz[ing] the information-gathering rate about the stimulus-response behaviour of the system" (P. Z. Marmarelis & Marmarelis, 1978). In the case of white noise sequences, the impulse responses are usually extracted by doing a cross-correlation between the input and output (Ringach & Shapley, 2004). They represent the average correlation between each flash presented (at time T) and the recorded activity at the same time and various delays after that (at time T + lags).

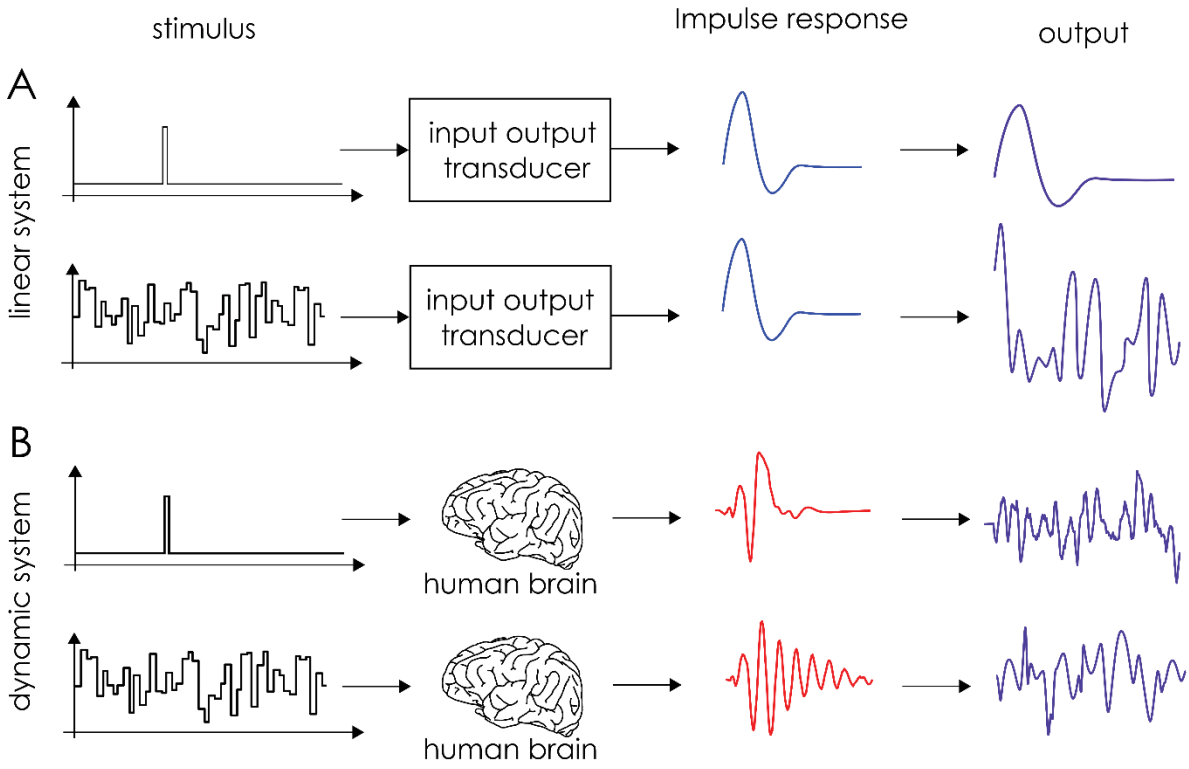

Figure 1-7: Extraction of Impulse Responses (IRF) in a linear (**A**) and a dynamic (**B**) system. IRFs can be extracted in two ways. First, by repeatedly presenting brief pulses of light at various intensities and averaging the resulting brain activity to extract the IRF. Second, white noise (WN) inputs can be used: these are equivalent to presenting an indefinite number of pulses of light of various shapes and frequency (Marmarelis & Marmarelis, 1978). The IRF is extracted by doing a cross-correlation between the WN sequence and the EEG signal. In a strictly linear system, the impulse response will be identical, whether extracted from the first or second method. However, in a non-linear system (such as our brains), WN sequences can

extract more sensitive IRF (Lalor et al., 2006) (figure partially reproduced from Ringach & Shapley, 2004 and Zaehle, Lenz, Ohl, & Herrmann, 2010).

Since IRF are the best (linear) model of how the system responds to each of the impulses in the WN sequence, once extracted IRF provide us with a means of reconstructing the output of the system to any new WN sequence. This is the consequence of three properties of a linear system: stability, homogeneity and superposition (Ringach & Shapley, 2004). Stability means that given an input intensity, the linear system's response will always be the same (Ringach & Shapley, 2004). Homogeneity means that, as the strength of the input increases, the strength of the response will also increase linearly, commensurate with the increase in stimulus strength. Finally, superposition means that when two different stimulus intensities are added, the response strength of the system to this combined stimulus will be equal to the response to each of these stimuli presented in isolation. All of this boils down to a simple (yet crucial) conclusion: in a linear system, once the impulse response has been extracted, the output doesn't need to be recorded, it can be successfully reconstructed based on the stimulation pattern (Lalor, 2009; P. Z. Marmarelis & Marmarelis, 1978; Ringach & Shapley, 2004). The reconstruction of the output is simple: the input is decomposed into a succession of various pulses of different intensities and the IRF is applied to each of these pulses with the appropriate scaling (depending on the input strength) and translated to be finally summed across all pulses (Ringach & Shapley, 2004). This process of scaling, translating and summing is equivalent to performing a convolution of the stimulus sequence and the impulse response (Ringach & Shapley, 2004). It bears mentioning that since the IRF wholly describes the behaviour of a linear system, the signal reconstructed in this case will be identical to the "recorded" signal. This, however, will probably not be the case in a non-linear system such as our brains. In fact, Ringach and Shapley (2004) argue that a necessary step in extracting the system's IRF is to evaluate how well it explains the data by computing the percentage of variance in the original signal explained by the reconstructed one. This reconstruction can be helpful in understanding the behaviour of the system under study.

1.4.3 Visual evoked potentials as IRF

The white noise approach described above has been used in the past to extract the EEG impulse response, this was predominantly done in two research groups.

The extraction of the VEP by white noise was first described by Lalor, Foxe and colleagues. Their original motivation was to increase the extraction speed for the VEP. Since the impulse response is the linear part of the event related dynamic, they used the linear system analysis method presented above to extract the visual evoked spread spectrum activity (VESPA) (Lalor, Pearlmutter, et al., 2006). As stimuli, they smoothly modulated the contrast of two patterns (a checkerboard or a snowflake of 5° of visual angle) presented at fixation using a random luminance sequence (with flat power between 0 and 30 Hz), while concurrently recording the brain activity using EEG. The VESPA was extracted by doing a linear regression of the EEG signal to the luminance sequence (Lalor et al., 2006). The VESPA is most prominent over the occipital channels and shows negative peaks at 75ms, 125ms, 175ms and positive peaks at 100ms and 150ms (see Figure 1-8A, left). While it is highly correlated ($r = 0.91$) with the VEP, there is a lot of variability between subjects (Lalor, Pearlmutter, et al., 2006). In a later study, they investigated the correlation more finely and found that the amplitude of the early component (C1) of both VEP and VESPA are highly correlated (Murphy, Kelly, Foxe, & Lalor, 2012). This was not the case for the P1 amplitudes, rather, the VESPA P1 was more correlated with the C1 measured from the VESPA, suggesting a common cortical generator (Murphy et al., 2012). Using source analysis of high density EEG data, they confirmed this results, and found that both components likely shared a cortical generator in the calcarine sulcus (Lalor, Kelly, & Foxe, 2012).

They have published a large number of studies looking at the VESPA in response to various stimuli: for example stimuli modulated in terms of spatial frequency (Lalor, Lucan, & Foxe, 2009) or modulated in terms of coherence in motion (Gonçalves, Whelan, Foxe, & Lalor, 2014). The stimuli can even be designed to target specific cells (Lalor & Foxe, 2009) or probe the specific frequency characteristics of the visual system (Lalor, Reilly, Member, Pearlmutter, & Foxe, 2006). Similarly to the VEP, the VESPA seem to be modulated by attention: the IRF to attended sequences show larger amplitudes in the P1 component (Lalor, Kelly, Pearlmutter, Reilly, & Foxe, 2007; confirmed by Frey, Kelly, Lalor, & Foxe, 2010; Jia, Liu, Fang, & Luo, 2017).

They also evaluate how good of a model of brain activity the VESPA was. To test the goodness of fit of the VESPA as a model of the brain activity, they systematically computed the VESPA using 80% of their data and used this to predict the EEG response for the remaining 20% of data (Lalor, 2009). The reconstructed EEG extracted from the VESPA had

a mean correlation of r = 0.084 across 7 subjects with the recorded EEG, although there was quite a lot of inter-subject variability (from 0.033 to 0.228).

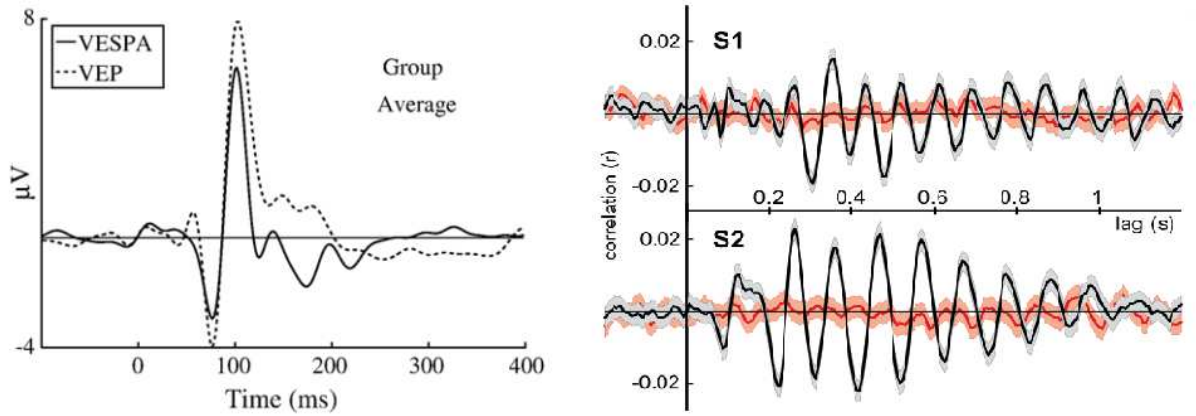

Figure 1-8. Examples of impulse response functions measured in the EEG. Left. The Visual Evoked Spread Spectrum Response Potential or VESPA (full black line) averaged across subjects extracted from central checkerboard and the Visual Evoked Potential (dotted line) extracted from pattern-reversals (Reproduced from Lalor et al., 2006). Right. Perceptual Echoes recorded from periphery (in black) compared to surrogates (i.e. the cross-correlation was also performed on shuffled trials, in red). The shaded area represent the SEM across trials. (Reproduced from VanRullen & Macdonald, 2012).

A similar approach has also used by VanRullen & Macdonald (2012): the difference was that the relationship between the WN (random luminance) stimuli presented in the periphery and EEG was evaluated at longer latencies lags. In addition to the VEP (e.g. seen for example in S2 from 0 to 0.2s on Figure 1-8), they found a ringing pattern at longer latencies (VanRullen & Macdonald, 2012). They termed this ringing a "perceptual echo", as it seems that the visual information (i.e. flash intensity) present in the stimulus sequence at time *t* was reverberated at a rate of ten times per second in the visual system. Interestingly, this ringing pattern has also been uncovered in single trial responses to a single sudden flash (Barlow, 1960; Makeig et al., 2002). The "perceptual echoes" are characterized by a strong oscillatory component mainly present over occipital channels with, on average across subjects, a peak in the alpha band at about 10 Hz. Interestingly, each subject has their own specific reverberation frequency (see Figure 1-8), which is correlated to the individual alpha peak frequency recorded with eyes closed over the occipital cortex (VanRullen & Macdonald, 2012). Moreover, subjects reported perceiving an illusory flickering in the WN sequences at a rate correlated with their own peak frequency of the echo (VanRullen & Macdonald, 2012). Furthermore, these perceptual echoes were stable in time: measured at 6 months interval, the

echoes were still very highly correlated (we will re-assess this question more systematically in the next chapter, with even longer intervals).

What is the functional significance of these echoes? As we will see later, they might play a role in the rhythmic sampling of the environment: they could reflect a mechanisms through which the brain carries information over between the different snapshots of perception (VanRullen, 2016c). Recently, Gulbinaite and colleagues (2017) showed that the triple-flash illusion could be explained by perceptual echoes. They presented two flashes with various stimulus onset asynchronies (SOA). They reported that the SOA at which subjects reported perceiving a third illusory flash was correlated with a peak frequency of the echo (Gulbinaite, İlhan, & VanRullen, 2017).

Finally, it bears mentioning that while the auditory evoked spread-spectrum response was successfully extracted using the same approach as described above but using auditory stimuli instead (Lalor, Power, Reilly, & Foxe, 2009), later attempts failed to find the presence of auditory "echoes" (Ilhan & VanRullen, 2012). This hints to a potential difference in the role of oscillations in both modalities, which might have evolved as a result of the differences in the characteristics of the auditory and visual stimuli (VanRullen, Zoefel, & Ilhan, 2014; Zoefel & VanRullen, 2017).

1.5 Experimental Evidence for the Influence of oscillations on visual perception

Alpha oscillations are the hallmark of the visual cortex: from the very first recording by Berger, a wide range of empirical findings has linked the alpha oscillation to the visual perception. In this section, I will review the evidence showing the functional role of alpha oscillations in perception.

. We will briefly review experimental evidence falling into 4 categories. First I will look at studies showing that the alpha rhythm is the preferred rhythm of the visual cortex, to then turn to studies linking the alpha frequency, amplitude and phase to perception. The following section is heavily inspired from previous reviews on the topic (VanRullen, 2016b, 2016c, VanRullen et al., 2011, 2014; Zoefel & VanRullen, 2017).

1.5.1 Alpha resonance in the visual cortex

Alpha oscillations are a prominent part of the EEG activity recorded over the occipital channels: subjects need just close their eyes and the alpha amplitude will visibly increase.

This preponderance of alpha activity is also seen in response to visual stimulation. Early studies of the sensory responses reported the response to single stimuli was characterised by a rhythmic after discharge (Adrian, 1941; G H Bishop & O'Leary, 1938). In fact, Barlow showed that in humans, the late part of the evoked response was characterized by a "ringing" component in the alpha band (Barlow, 1960; Childers & Perry, 1971), which was closely related to the ongoing alpha oscillation (Lansing & Barlow, 1972). Higher ongoing oscillatory alpha power led to stronger post-stimulus ringing (Makeig et al., 2002). However, this was discussed by Childers & Perry who argued that the ringing was only an alpha-like response, which was functionally separate from the ongoing alpha oscillations (Childers & Perry, 1971). As we have just seen, using broad-band stimulation reveals a reverberation of visual information in the alpha band over the occipital channels (VanRullen & Macdonald, 2012). The "steady-state visual evoked potential" in response to flickering stimuli is characterized, at all frequencies of the stimulation between 1 and 100 Hz with a response in the alpha range (Herrmann, 2001). Moreover, a resonance phenomenon has been shown in the alpha band: visual stimuli flickered at 10 Hz give rise to a stronger response than neighbouring flicker frequencies (Herrmann, 2001), although resonance was not specific to the alpha band.

Finally, alpha activity has also been shown to arise after a single TMS pulse to the left visual cortex (Herring, Thut, Jensen, & Bergmann, 2015).

Taken together, these studies highlight the preponderance of alpha oscillations in the visual cortex. They are the hallmark of its functioning, but do they play a functional role in visual perception or are they just the "epiphenomenal wiggling of the jelly brain" as eloquently put by György Buzsáki (2006)?

1.5.2 Frequency dependent perception

Many studies of frequency dependent perception have proposed that the alpha rhythm could be a stable marker for a wide range of cognitive functions. It has been correlated with various cognitive factors such as intelligence scores, information processing speed, and even memory performance (Bazanova & Vernon, 2013). The exact frequency of the resting state

alpha oscillation is stable over longer periods of times, although it does decrease with age (Klimesch, Fellinger, & Freunberger, 2011). The individual alpha peak frequency (IAPF) is a relatively stable trait within subjects: across different tasks, the within subjects variability in the IAPF is 3 times smaller than the between subject variability (Haegens et al., 2014). Importantly, this study shows that limiting the alpha frequency band to the 8 to 12 Hz is too narrow: in some conditions, some subjects had peak frequencies which were as low as 7Hz or even higher than 12 Hz (Haegens et al., 2014). Thus, it seems important to take into account the subject's IAPF when analysing data. Creating subject specific frequency bands could increase the sensitivity of the analysis (for a method on how to compute the IAF, see Klimesch, Fellinger, & Freunberger, 2011), thereby increasing statistical power by taking individual differences into account (Cohen, 2016).

This subject specific frequency has been linked to visual perception too: the reaction time to the sound was correlated with the alpha frequency for each subject on a single trial level (Survillo, 1961). However, only 13 people (with few trials each) were recorded in this study. Furthermore, it is not possible to completely rule out that age could act as a confounding variable: the older participants had slower alpha oscillations and slower reaction times on average: a central mechanism related to aging could influence both the speed and the alpha frequency.

Recently, Samaha and Postle (2015) showed that the IAF frequency was critical for temporal parsing, that is for a good temporal resolution in visual perception. The phase of the alpha rhythm has been suggested to act as a parsing mechanism: when two flashes are presented briefly one after the other they will be perceived as simultaneous (i.e. fused) or sequential depending on where they fall with regards to the ongoing alpha phase cycle. The two-flash fusion threshold was correlated with the IAF: subjects with higher IAF (i.e. shorter cycle duration) had significantly lower thresholds (Samaha & Postle, 2015).

One limit to these effects of the frequency-dependent performance is that they require the use of between subject comparisons where other factors might also co-vary along with the alpha peak frequency such as age. This has to be carefully taken into account when designing and interpreting the results from experiments.

1.5.3 Alpha oscillation amplitude and visual perception

We continue our tour of the causal role of oscillations by looking at the evidence showing a link between alpha amplitude and visual perception. Until recently, the amplitude of the alpha oscillations were thought to reflect the "idling" rhythm of the brain (Pfurtscheller & Lopes Da Silva, 1999; Pfurtscheller, Stancák, & Neuper, 1996). However, two strands of evidence have been unfurled in the past 20 years, which speak to the contrary.

1.5.3.1 Alpha amplitude is linked to target detection

First, the development of new analysis techniques as well as the increasing power of computers has allowed the study of single trial variation in oscillatory components. These have revealed that the pre-stimulus amplitude of the ongoing alpha oscillation is correlated with target detection. Ergenoglu and colleagues (2004) used weak stimuli (i.e. near perceptual threshold) while recoding the EEG. They found that on a trial by trial basis the perception of the stimuli was linked to the amplitude of the alpha oscillations: higher alpha amplitudes led to poorer perceptual outcomes (Ergenoglu et al. 2004). This finding has since then been replicated (Busch et al., 2009) as well as expended using contrast masking (Mathewson et al., 2009), spatial attention tasks (van Dijk, Schoffelen, Oostenveld, & Jensen, 2008; Wyart & Tallon-baudry, 2008), illusory perception (Lange, Keil, Schnitzler, van Dijk, & Weisz, 2014), discrimination (Hanslmayr et al., 2005, 2007). Using single pulse TMS applied over the visual cortex, Romei and colleagues showed that the ongoing alpha amplitude was correlated within subjects with the likelihood of perceiving a phosphene (Romei, Brodbeck, et al., 2008) and between subjects with the individual threshold for triggering phosphenes (Romei, Rihs, Brodbeck, & Thut, 2008). Using tACS either in the dark (or illumination) at a wide range of frequency revealed that in darkness, alpha band stimulation was the most effective at eliciting phosphenes (Kanai, Chaieb, Antal, Walsh, & Paulus, 2008). In a recent paper, Iemi and colleagues (2017) evaluated what the functional role of these fluctuations could be. Using two experiments rooted in signal detection theory, they show that the fluctuations in alpha amplitude were better explained by a "baseline" model, where the alpha power reflected a global shift in baseline excitability (i.e. criterion shift), leading to higher likelihood of answering "yes", whether a stimuli is presented or not (Iemi, Chaumon, Crouzet, & Busch, 2017).

1.5.3.2 Alpha amplitude as a marker of spatial attention

As well as being involved in the target perception, the amplitude of alpha oscillations is also modulated by (covert) spatial attention. The orienting of attention has usually been

investigated using the Posner paradigm (Posner, 1980), where (covert spatial) attention is cued (validly or invalidly) to the side of the presentation of a peripheral target to be detected. It allows for a detailed investigation as the cue validity, cue-target asynchrony and the presence of distractors (among many others) can be manipulated to probe the time course of attention (Carrasco, 2011, 2014; Posner, 1980).

Two effects in alpha amplitude have been reported as a function of the deployment of attention towards the stimulus. First, Worden and colleagues reported an increase in the amplitude of the ongoing alpha oscillations over *ipsilateral* (to the attended side, contralateral to the inhibited hemifield) occipital-parietal channels (Worden et al., 2000). This was later replicated by many studies (Cosmelli et al., 2011; Doesburg, Roggeveen, Kitajo, & Ward, 2008; Kelly, Lalor, Reilly, & Foxe, 2006). It has been taken as a sign for the inhibition of task-irrelevant brain areas (Hanslmayr et al., 2011; Rihs, Michel, & Thut, 2007). Moreover, this effect is retinotopic: the topography of the alpha increase over ipsilateral cortex was specific to each of 8 attended locations (Rihs et al., 2007). In fact, the effect of spatial attention was specific enough in terms of topography so that the locus of spatial attention could be reconstructed (Samaha, Sprague, & Postle, 2016). Finally, it was shown to be independent from the event related synchronisation (ERS, i.e. increased alpha amplitude). The ERS can occur in response to stimuli offset (Pfurtscheller & Lopes Da Silva, 1999) or even during passive states when the subjects had their eyes closed (Pfurtscheller & Aranibar, 1977). Kelly and colleagues (2006) used flickering stimuli to constrain brain activity and thus remove any effect of the ERS due to the cue. They confirmed that the effect was due to an increased alpha amplitude over the contralateral cortex above and beyond that explained by ERS (Kelly et al., 2006).

However, the opposite effect has also been reported, namely that the deployment of spatial attention was linked with a decrease in alpha amplitude over the hemisphere *contra-lateral* to the attended side (Sauseng et al., 2005; Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Wyart & Tallon-Baudry, 2009). The effects were shown to be independent from the systematic alpha band event related desynchronisation (ERD) evoked by the presence of a stimulus (Pfurtscheller & Lopes Da Silva, 1999) by using auditory cues instead of visual ones (Thut et al., 2006).

The causal role of alpha oscillation amplitude was shown by using repetitive TMS in the theta, alpha or the beta band over the occipital and parietal cortex to entrain the oscillations (Romei et al., 2010). The authors report that increased alpha amplitude, due to the

rTMS, specifically impaired the detection of contralateral (near-perceptual threshold) visual targets. This was not the case for the other frequency bands (Romei et al., 2010). A later study confirmed that this effect was due to the entrainment of the natural occipital alpha rhythm (Thut, Veniero, et al., 2011).

This raises the question whether this alpha lateralization effect is due to a decreased alpha amplitude due to the visual attention engagement or to an increased alpha amplitude in the unengaged hemisphere? Both mechanisms could act hand in hand to bias perception (Thut & Miniussi, 2009). In fact, a study has shown that both effects could be present at different time points. Rihs et al. (2009) reported an alpha amplitude *decrease* in the "early" (<700ms) phase of processing, reflecting the preparation for target detection, and a "late" (>700ms) alpha amplitude *increase* reflecting attention driven inhibition of the unattended regions.

1.5.4 Alpha amplitude as a tool for inhibition

It has been suggested that ongoing EEG alpha amplitude reflects an active inhibitory mechanism (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, & Hanslmayr, 2007; Mathewson et al., 2011). Task irrelevant brain areas are silenced by increased alpha amplitude, allowing for the information to be routed to task relevant areas (Jensen & Mazaheri, 2010). In other words, ongoing alpha amplitude tune occipital-parietal areas responsible for the processing of visual information for the upcoming stimuli: it serves as a regulation mechanism for the incoming information flow, up and down regulating the activity in the occipito-parietal regions as necessary (Thut & Miniussi, 2009). In fact, the phase of the alpha oscillations has been shown to modulate the amplitude of gamma band oscillations (Osipova, Hermes, & Jensen, 2008), which reflect the active processing of information in these task relevant areas (Jensen & Mazaheri, 2010). This cross-frequency coupling of alpha and gamma activity could also explain how un-attended salient stimuli can be quickly brought under the spotlight of attention when necessary (Jensen et al., 2012).

This inhibitory role of alpha oscillations has been reported not only in vision, but also in other sensory modalities as well as in multi-sensory tasks (Banerjee, Snyder, Molholm, & Foxe, 2011; Foxe & Snyder, 2011).

Recently, this gating information flow by inhibitory alpha activity was integrated with the idea of communication through coherence presented above (Bonnefond, Kastner, & Jensen, 2017). In this framework, functional connectivity between regions is created by phase

synchrony of alpha oscillations and allows the communication of representational information reflected in the coupled gamma band activity. This is supported, among other evidence, by the specific oscillatory frequency supporting feed-back and feed-forward connections (e.g. Bastos et al., 2015; van Kerkoerle et al., 2014)

1.5.5 Phase dependent perception

We have seen that both the frequency and amplitude can have an impact on visual perception. However, can oscillations have moment by moment influence on our visual perception? In the next section we will show that the instantaneous phase of ongoing oscillations modulates visual perception at a very fine time scale.

Usually, the effects of phase are investigated by contrasting the phase of two type of stimuli (for example detected versus missed) at different time points in the pre-stimulus time window. The effects cannot be evaluated after stimulus onset as the evoked potential has mainly similar phases across all trials. Studies on phase dependent performance in the VanRullen lab have been reviewed in a recent paper have reported that about 10 to 20% of the variability of visual performance could be explained by the instantaneous, trial-by-trial variability in the pre-stimulus phase (VanRullen, 2016b).The effects of phase on visual performance have been reported in two different frequency bands.

First, there seems to be an influence of the theta band (~7 Hz) over the frontal cortices. For example, Busch and colleagues showed the likelihood of perceiving a very small and brief target in the periphery was dependent on the frontal theta (~7Hz) oscillatory phase just before stimulus onset (about -100ms), accounting for about 16% of performance variability (Busch et al., 2009). This effect was replicated with attentional manipulation (Busch & VanRullen, 2010; Dugué, Marque, & VanRullen, 2015). In an EEG-fMRI experiment, the phase of the theta band oscillation (~7Hz) was found to gate the information flow between frontal “higher level” areas and occipital “lower level” areas in a contour detection task (Hanslmayr, Volberg, Wimber, Dalal, & Greenlee, 2013).

Second, some effects are reported in the alpha band. For example, Mathewson and colleagues showed an effect of phase on the detection of masked targets but this time in the alpha range over the occipital cortex (Mathewson et al., 2009). This is reminiscent of the study of the galvanic skin response to the masked word “danger” done by Nunn and Osselton over 40 years ago: they found that the evoked response was larger at the descending and

trough than at the ascending and peak phase of the ongoing alpha oscillations (Nunn & Osselton, 1974). Using EEG to record the ongoing oscillation, Dugué and colleagues took advantage of single pulse TMS to probe the excitability of the cortex (Dugué, Marque, & VanRullen, 2011). They showed that "phosphenes" (illusory visual percept evoked by magnetic pulses) were more likely to occur at a certain phase of the ongoing alpha band (~10 Hz) oscillation measured through EEG than at the opposite phase (Dugué et al., 2011). Alpha oscillatory phase has been shown to influence eye-movement initiation (Drewes & VanRullen, 2011) and mis-localisation (McLelland, Lavergne, & VanRullen, 2016), and a spatio-temporal illusion (Chakravarthi & VanRullen, 2012). Using four flashes of light to entrain the oscillations, de Graaf and colleagues (2013) showed that entrainment of the alpha oscillations (compared with other frequencies) resulted in a specific impairment of the usual cueing benefit (de Graaf et al., 2013). In separate task, they also found that the discrimination of a target presented after the entraining rhythm was modulated during 3 cycles of the alpha rhythm (de Graaf et al., 2013). This effect was also replicated by Spaak and colleagues (2014) who showed that this rhythmic modulation of performance was supported by a neural entrainment of alpha oscillations. A recent study has shown that tACS could successfully entrain alpha oscillations over the parietal occipital activity, which successfully modulated behaviour (Helfrich et al., 2014).

In fact, the presence of (at least) two rhythms whose phase influences perception can be clearly seen from the meta-analysis of 10 different studies (see Figure 1-9) (VanRullen, 2016b).

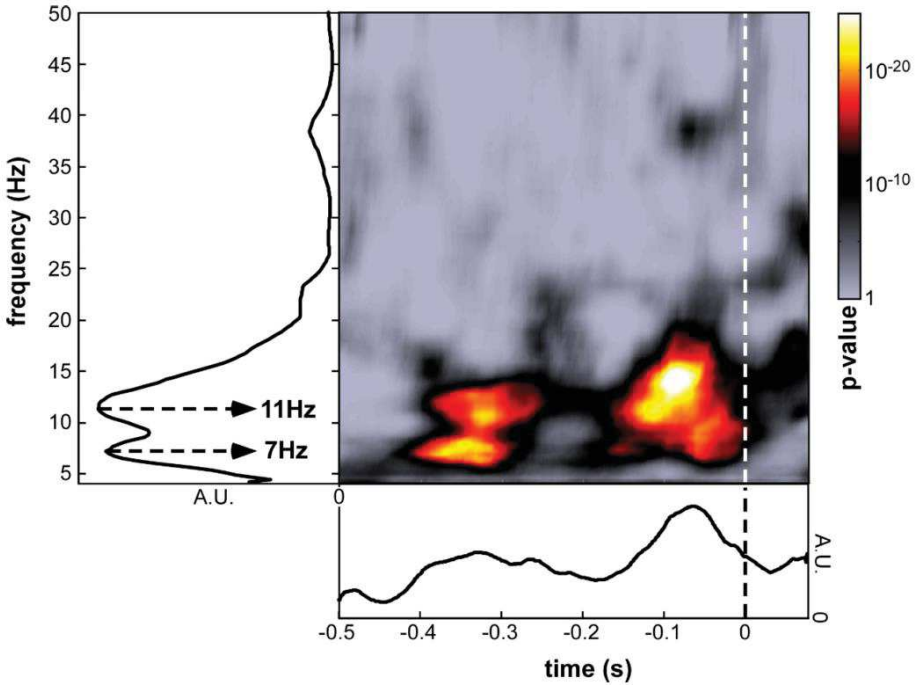

Figure 1-9. Meta-analysis from 10 datasets showing the dependence of visual perception (various paradigms) on the pre-oscillatory phase. The phase dependent effects are predominantly found in the pre-stimulus time window, between -400 to -100ms before the target onset. Moreover, two frequency peaks seem to emerge, one at ~7Hz and ~10 Hz. (reproduced from VanRullen, 2016a).

The largest phase differences between conditions (reflected in the significance of the effects) are found in the pre-stimulus time window from -400 to about -100ms before the target is presented (at time 0, see Figure 1-9). The spectral representation of the effect clearly shows the presence of 2 frequency peaks, one in the lower alpha/theta range at ~7 Hz and one in the alpha range at ~11 Hz. While speculative, the pattern of evidence suggests two different roles for these rhythms: lower alpha/theta oscillations tend to occur over frontal cortices and seem to reflect oscillations which are involved in an attentional task while alpha oscillations could be involved in more sensory information processing over occipital channels (VanRullen, 2016b). For the remainder of this thesis, these oscillations will be referred to by "occipital alpha" and "frontal theta", for convenience and distinction. However, this is not to suggest that these frontal theta band oscillation have the same functional role as that usually associated with "frontal midline theta oscillations" (i.e. executive functions and cognitive control, Cavanagh & Frank, 2014).

1.5.6 Perceptual Cycles

"a lifetime is roughly 20 billion moments"(Stroud, 1955)

The "perceptual cycles" hypothesis is directly related to the evidence just reviewed: it offers that these rhythmic modulation in perception are a direct result of the underlying brain oscillations. The next section presents the key concepts of the theory; as such, it is inspired from recent reviews on the topic (Dugué & VanRullen, 2017; VanRullen, 2016b; Zoefel & VanRullen, 2017).

As mentioned, rather than involving a single central scanning mechanism supported by a single oscillatory band, the literature described above points towards the existence of several rhythms with different functional roles (VanRullen, 2016b).In particular, there seem to be two prominent rhythms in visual perception, which could reflect separate functions and have complementary roles. The alpha perceptual rhythm could carry out a sensory sampling mechanism over the incoming information in the occipital cortex while the theta band oscillations could represent a central attentional scanning mechanism involved in inter-areal communication (Dugué & VanRullen, 2017; Zoefel & VanRullen, 2017).

This distinction is somewhat reminiscent of the two mechanisms reported by Harter (1967) as to the source of periodicity in perception. The oscillatory mechanism creating perceptual cycles could be the results of two mechanisms: the "cortical excitability cycles" and the "central scanning mechanisms".

In the "cortical excitability cycles" hypothesis, rhythmicity in perception is driven by a local sampling of incoming visual information. This sampling is a direct consequence of the link between LFP phase and excitability, first demonstrated by Bishop in the visual cortex of the rabbit (Bishop, 1932). As we have seen, many studies have provided evidence for the gating of neuronal excitability by the phase of local oscillations in various frequencies (Fries et al., 2007; Haegens et al., 2015; Lorincz et al., 2009; Montemurro et al., 2008; Whittingstall & Logothetis, 2009). Under this mechanism, the sampling of the visual information is a direct result of the fluctuations between moments of more or less excitability in the neuronal population which gate incoming information flow (Lansing, 1957; Lindsley, 1952). Information arriving at the moments of high neuronal excitability would be processed in a fast

and more efficient manner, while information arriving at moments of low neuronal excitability would be less well processed, or even missed altogether (see Figure 1-10A).

The "central scanning mechanism" on the other hand proposes that rhythmicity in perception arises as a function of a global oscillatory mechanism. This central scanning mechanism creates a frame of reference for incoming sensory information, it acts as the clocking device of the brain, giving rise to "perceptual moments" (Stroud, 1955). These perceptual moments are stable (i.e. no temporal information is present within one moment), the processing of information occurs in units (or moments) of ~100ms, which are not divisible (Stroud, 1955). In other words, they are like a video frame: time doesn't exist within it, only in between frames can the story exist (see Figure 1-10B). Thus, this central scanning mechanism creates "gaps" between the different frames and consequently in our perception. Information arriving at moments of lower excitability does not have to be completely lost from perception, it could simply be deferred until the next cycle (VanRullen, 2016b). This property is directly related to the temporal parsing ability described above: stimuli falling within the same discrete frame will be perceived as occurring together while the stimuli falling into two separate frames will be see as sequential (see section 1.5.2 on frequency dependent perception).

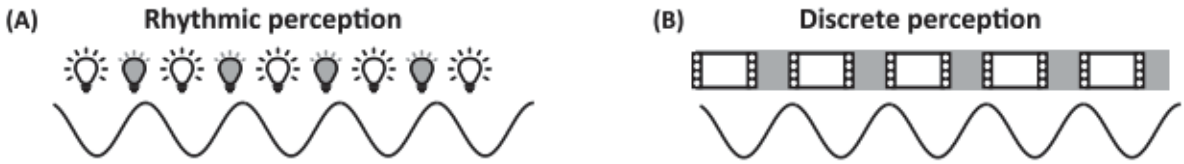

Figure 1-10. Difference between rhythmic and discrete perception. A) Rhythmic perception entails that efficiency of processing fluctuates as a function of the underlying phase: it will be optimal at one phase (brighter light bulbs) and sub-optimal at the opposite, it does not, however necessitate a gap in between frames as suggested by the idea of discrete perception (B). (reproduced from VanRullen, 2016b)

In other words, the idea is that the visual systems functions as a camera would, taking snapshots of the outside world every 100ms or so. From this metaphor, two different properties of the theory can be unfolded: first, that the shutter of the camera opens rhythmically, and second, that the picture taken is itself static. This distinction reflects the dissociation between the idea of "rhythmic" and "discrete" perception (see Figure 1-10): the

first “metaphorically relates to the shutter of the camera opening and closing periodically” (VanRullen, 2016b) while the latter “refers to the static property of each snapshot, namely that no temporal information can ever unfold within one snapshot” (VanRullen, 2016b). These properties (“snapshots” and “gaps”) can be linked back to the two mechanisms described above for the source of the oscillatory sampling rhythm (namely, that a “cortical excitability cycles” and a “central mechanisms”). The existence of a central scanning mechanism, which acts as the clocking mechanism of the brain, will create “gaps” in our perception and thus render our perception discrete. This discretization of perception would create “chunks” of mental and sensory events, between which nothing occurs (VanRullen, 2016b). This is similar to the potential role of the frontal “theta” oscillations presented above, which samples the sensory information and reflects a global attentional scanning mechanism of visual information by higher level areas (Dugué & VanRullen, 2017). As such, these frontal theta oscillations would speak in favour of a central “cortical scanning” mechanism, and thus reflect rhythmicity in attention. On the other hand, the sampling of environmental information through cycles in cortical excitability will create “snapshots” in perception and giving rise to a rhythmicity in perception. This is similar to the role of the occipital alpha oscillations described above. They reflect the presence of a local perceptual sampling mechanism related to the processing of sensory information and are a direct result of local fluctuation in excitability.

Note that, it is not necessary that the “gaps” and “snapshots” properties described above arise as a function of different mechanisms. These two properties are inherently linked: the presence of a fluctuations in the neural processes relating to excitability will lead to fluctuations in the timing of information that will effectively give the illusion of “gaps” (and thus of discrete perception) as a function of a rhythmicity in perception (VanRullen, 2016b).

1.5.7 Perceptual echoes as a tool for studying the perceptual cycles?

Most of the evidence reviewed above presented evidence for the existence of perceptual cycles by looking at the pre-stimulus fluctuations in spontaneous oscillations. The perceptual echoes, on the other hand, could be evidence that this sampling of visual information continues *during* visual processing (VanRullen, 2016b). As mentioned above, the perceptual echoes reflect a long lasting correlation between the input luminance and the brain response. However, the functional role of these perceptual echoes is still unresolved. Are they

just the by-product of the natural preference of the visual system for alpha oscillations or do they have a functional role? Evidence presented above suggests that they do in fact play a role in perception: Gulbinaite and colleagues (2017) showed that the triple flash illusion could be explained, on a subject by subject basis, by a superposition of the perceptual echoes in such a way as to create a third "illusory" flash. Since the perceptual echoes occur *during* stimulus processing, they could be the direct reflection of the cyclic perception. In particular, they could allow the coalescence of sensory information in time, transferring information from one perceptual cycle to the next (VanRullen, 2016b).

In a nutshell, this review of the literature has highlighted the strong evidence in favour of the causal role of oscillations in shaping perception. Lower frequency oscillations in particular seem to be crucial in visual processing: the phase, amplitude and frequency of theta and alpha band oscillations directly influence our experience of the external world. They create perceptual cycles, of which the perceptual echoes are a direct measure.

1.6 General Framework

Broadly speaking, the overarching aim of this thesis is to better understand the relationship between oscillatory signals and visual perception. In particular, we wish to confirm that perceptual echoes can be used as a tool for studying perceptual cycles in vision. If the perceptual echoes are found to directly reflect the intrinsic cyclic nature of perception, they could be a valuable tool in controlling perception (as we will see later). Thus, we used the White Noise Paradigm throughout this thesis. Namely, we use white noise sequences to constrain brain oscillations giving rise to the perceptual echoes. Once these perceptual echoes are extracted, we reconstruct the time course of the background oscillations. We then used a target, embedded within these white noise sequences, to probe the functional role of these oscillations.

1.6.1 The problem with recording EEG

In recent years, converging evidence has emerged for the functional role of oscillatory signals on sensory neural processing: the instantaneous state of oscillations (whether measured in terms of phase, amplitude or frequency) has been found to causally influence the

processing of incoming visual information. This relationship has often been studied using M/EEG method to record the ongoing oscillations.

EEG has been a very useful tool in studying the link between ongoing oscillations and visual perception: it was this method and the development of powerful analyses techniques that allowed the discovery of such a link. However, the problem with EEG is also its strength: when we record the brain activity, all signals co-occurring in the brain at the same time are recorded. While this is extremely useful to study the temporal development of various cognitive functions, this also means that the activity from these various sources cannot always be easily teased apart. This problem is circumvented to some extent by rigorous experimental paradigms and controls. Nevertheless, even with the most rigorous experimental paradigms, I would like to argue that several types of oscillatory signals are consistently mixed in an EEG recording. For example as mentioned above, at least two oscillatory mechanisms are implicated in the creation of rhythmicity in vision. How can we functionally differentiate them? We suggest that this can be done using white noise sequences to control the time course of the "occipital alpha" oscillation related to sensory selection and processing. Arguably, the perceptual echoes, which arise as a function of the white noise stimulation, are a direct reflection of this mechanism.

To better understand the types of signals which are constrained by the white noise sequence and their relative contribution to the effect of interest (how oscillations influence visual perception), we present a taxonomy of brain signals

1.6.2 White noise sequences and brain signals

Here, I present the taxonomy of brain signals which will be used in the remainder of the thesis. Note that, as such, the taxonomy is artificial: it directly relates to our main experimental paradigm. In order to tease apart the different kinds of signals recorded in response to white noise sequences, we ask three questions: 1) Is the signal influenced by the white noise sequence? 2) If yes, does it respond in a phase-locked (or "evoked") or non-phase-locked manner (or "induced")? 3) Does the signal play a functional role in the perception of the target? Note that here the perception of the target is specifically relating to the presence of a target presented within the white noise sequences. The answer to each of these questions taken together allows us to differentiate six potential types of oscillatory signals recorded in the EEG at any one time (see table 1).

		Is the signal influenced by the white noise sequence?		
		No	Yes	
			Does the signal respond in a phase-locked manner?	
			Yes	No
Does the signal play a role in the target's perception?	No	*Signal S1* **Irrelevant spontaneous oscillation** *Neural noise, recording noise and any signal (potentially oscillatory) related to other cognitive sensory function*	*Signal S3* **Irrelevant phase locked oscillation** *Activity (oscillatory or not) that arises as a consequence of visual processing in a phase locked manner. It has no influence on the perception of upcoming targets*	*Signal S5* **Irrelevant non-phase locked oscillation** *Similar to S3, only it has no phase relationship with the stimuli sequence*
	Yes	*Signal S2* **Relevant spontaneous oscillation** *Oscillations which occur independently from the sequence (i.e. are not dependent on visual stimulation) but which can still influence the detection of upcoming targets*	*Signal S4* **Relevant phase locked oscillation** *Oscillations of interest: they are both constrained in a phase locked manner by the white noise sequence and directly influence the processing of the upcoming target*	*Signal S6* **Relevant non phase locked oscillation** *An oscillatory response that would not be phase-locked to the luminance sequence would "average out" in the cross-correlation*

Table 1. Taxonomy of signals recorded in our White Noise EEG experiments. Three questions can help us differentiate the signals. First, we ask whether the signal plays a functional role in the perception of the targets: if it does not, it is called "irrelevant" and if it does, it is called "relevant" (blue outline in the table). Second, we asked whether the signals are influenced by the WN sequence: if they aren't they are called "spontaneous", if they are, we ask a third question, to distinguish between "phase-locked" and "non-phase locked" responses. In total, 6 different categories can be created. Although, some of these signal types might not actually be found in the brain, it is a useful distinction as it allows to easily isolate the signals we can model using the WN sequences (shaded in yellow).

First, we distinguish two types of signals depending on whether they play any functional role in visual targets detection. Since we are interested in oscillatory signals which do have a role to play in visual perception, we will simply refer to these signals as either "relevant" (if they do play a role) or "irrelevant" (if they don't). Second, we distinguish three types of signals depending on how they respond to white noise sequences. If they are blind to the stimulation, we will call these signals "spontaneous", that is signals whose state is not

determined by the WN sequence. While this is not a traditional definition of spontaneous oscillations, it will serve well insofar as these signals are not constrained by the visual stimulation presented. On the other hand, for the signals which are influenced by the WN sequence, we will further distinguish the "phase locked" signals from the "not-phase locked" ones. In total this creates 6 stimulus categories:

1. irrelevant spontaneous oscillations (signal S1)
2. relevant spontaneous oscillations (signal S2)
3. irrelevant phase locked oscillations (signal S3)
4. relevant phase locked oscillations (signal S4)
5. irrelevant non phase locked oscillations (signal S5)
6. relevant non phase locked oscillations (signal S6)

This does not necessarily mean that all these 6 types of signals exist in the brain. However, we will see that most of them do. Crucially, this allows us to define the type of signal which will be accessible to our IRF analysis of the WN sequence: the phase locked signals (see Table 1, shaded area) and those that are of interest when trying to explain the perception of targets (see Table 1, cells contoured in blue).

The easiest signals to functionally differentiate are the "spontaneous" signals which do not respond to the white noise sequences. Of these, two different types of signals can be distinguished. First, there are the signals which do not play any functional role in the processing of the visual target: these we will call the "irrelevant spontaneous signals" or in other words the "Noise" signals (signal S1, table 1). In these signals, we will regroup any signal recorded in the EEG that does not respond to the WN sequence at all and which are irrelevant for the perception of targets. This category could include signals which make up the recording noise, or neural noise and might include signals that are related to information processing in other sensory modalities and cognitive functions for example. The second part of the spontaneous signals is the ones which are "relevant" for visual processing and thus would fall under the purview of "relevant spontaneous oscillations" (signal S2, table 1). While the instantaneous state of these oscillations might not be influenced by the WN stimulation, they do play a role in the processing of a target stimulus. This might, for example, include signals from the auditory cortex which have been shown to influence visual perception (Romei, Gross, & Thut, 2012). It might also be signals that originate in the visual cortex and

which represent spontaneous fluctuations in excitability as part of an attentional system (Boncompte, Villena-Gonzalze, Cosmelli, & Lopez, 2016; Romei, Brodbeck, et al., 2008). In fact, ongoing oscillations whose state is modulated, not by the visual stimulation but by "top-down" or endogenous factors (such as spatial attention) would also fall in this category. Thus, these signals might also include the signals which reflect the attentional sampling of visual information by the frontal theta oscillations. This is by no means an exhaustive list of all possible signals satisfying this condition.

We now turn to the signals whose state is constrained by the white noise sequence. The influence of the WN sequence can be realised in two ways: either in a phase locked or non-phase locked manner. The "Phase Locked" signals (signal S3 and S4 in table 1) will be the ones whose phase is directly influenced by our manipulation. Of these signals, we can differentiate two types of signals: the "relevant phase locked signals", and the "irrelevant phase locked signals". In the irrelevant signals, I will classify any evoked response such as the ERP or VEP: signals which directly reflect the processing of an incoming visual stimulus (such as the luminance fluctuations in the background sequence), without playing any further role in the processing of a subsequent target. In the relevant signals, I will classify signals whose state is constrained by the WN sequence and thereby will play a role in the perception of any target presented. Of course, this distinction is subtle and a direct consequence of our choice of taxonomy. The distinction between the signals S3 and S4 is the difference between the evoked response that is simply a consequence of visual processing of the background WN sequence (in this case, it is S3) and one that also causes *further* changes in visual processing of a subsequent target (in which case it will be S4). Note that while WN sequences will constrain their state, it does not follow that WN sequences are necessary for these relevant phase locked signals to exist. It is feasible to imagine that S4 signals are still present in the brain when there is no stimulus/visual input to constrain their state. The type of signals described here is reminiscent of the oscillatory signal described by Steriade and colleagues as "spontaneous oscillations": a sort of tonic background activity that occurs after evoked/induced oscillations and is linked with brain excitability (Steriade et al., 1993). Constrained oscillations represent the readiness of the brain for new stimulation after stimuli have been presented. In this thesis I will use the terms "relevant phase locked oscillations" or "constrained oscillations" (signals S4) to talk about these signals to mark a clear distinction

with the “spontaneous signals” (signal S1 and S2) which, under our taxonomy cannot be modified by visual stimulation.

Finally, we have the signals which are constrained by the WN sequence but in a “non-phase locked” manner (signals S5 and S6 in table 1). Of these signals, we here again differentiate between the “irrelevant” and the “relevant” signals. An example of the latter could be the induced gamma band oscillations involved in sensory representation and awareness (Baldauf & Desimone, 2014; Wyart & Tallon-baudry, 2008) mentioned above.

In a nutshell, using three questions, we create a taxonomy of brain oscillations which can be recorded in EEG in response to the white noise sequences stimulation. We define six types of signals: Irrelevant Spontaneous Oscillations (signal S1), Relevant Spontaneous Oscillations (signal S2), Irrelevant Phase Locked Oscillations (signal S3), Relevant Phase Locked Oscillations (signal S4), Irrelevant non Phase Locked Oscillations (signal S5) and Relevant non Phase Locked Oscillations (signal S6). Of course, any taxonomy is by definition biased and partial, insofar as all attempts at classifying any sort of object must make choices about which features to emphasize and which to ignore. However, it is a useful classification because it allows us to bring to light key limitations of M/EEG recording in light of our paradigm.

In the case of the EEG recording to white noise sequences, all six signals presented above are recorded at the same time. This has serious implications for understanding the relationship between perception and oscillations.

First, constraining the oscillations which give rise to the perceptual echoes, we can directly ask what their functional relevance is, outside of the influence from any other type of oscillation described above.

Second, we can re-examine the findings in the literature of the phase and amplitude dependent modulation of performance and ask whether these can be found in the absence of any attentional modulation to which they are usually ascribed. Most studies linking alpha oscillations to visual perception interpret the effects as reflecting the effect of attentional sampling or a top-down gating/biasing of visual oscillations, which would fall within the “spontaneous oscillations” (signal S2). More specifically, the effects of any attentional

sampling would fall within this category of signals (although there might be attentional effects driven by the stimuli sequence itself, which will be further discussed in the general discussion).

Third, if we find an influence of these constrained oscillations on visual perception, this would be (another) proof that oscillations causally shape perception. Since the visual stimulation allows us to control the state of the oscillation, any link made with visual perception will be necessarily causal.

Fourth, recording all signals together might occlude some effects of oscillations on perception. In particular, it is well known that phase effects cannot be investigated in the post-stimulus time-window due to the presence of an evoked response.

1.6.3 Aim of the thesis

In this thesis, we use the White Noise Paradigm to overcome the limitations presented above. Shortly stated (a precise description will be given in the next section), the white noise paradigm uses white-noise sequences to constrain the state of these background oscillations (i.e. signal S3 and S4). We extract a model (i.e. the perceptual echoes) of how the "constrained" oscillations are affected by the white noise sequences. This model is then used to predict the brain activity to new WN sequences instead of recording it. Effectively this predicted brain signal only contains phase locked signals, while all the other signals are not present in the prediction (though they may still be present in the brain, of course). Of these two phase-locked signals S3 and S4, only one is relevant for perception of the target (S4). Thus, this allows us to functionally differentiate it from the other two "relevant" signals for visual perception (S2 and S6). In the remainder of this thesis, we use this functional differentiation to answer various questions in the literature about the relative role of these different types of signals. In particular, we wish to understand what the functional role of the perceptual echoes is. I will now briefly present the aims and questions answered throughout this manuscript.

First, since the perceptual echoes are at the heart of this thesis, we ask what their neuronal bases are using an fMRI paradigm. In particular, this allows us to directly search for the brain regions supporting the rhythmic sampling of incoming sensory information. The preliminary results are reported in chapter 3.

Second, there has been an outstanding issue in the literature about the latency at which phase effects are found. If phase has a causal impact on visual perception, then we should be able to see it in the pre-stimulus time window. This issue is investigated in chapter 4, using simulations and the WN paradigm. It is especially suited to study this question as it allows us to disentangle the relative contribution of the evoked responses (signal S3) and the background oscillations (signal S4).This also allows us to show (as we will see later) that most of the phase effects reported in previous studies might actually be due to this S4 signal type.

Third, the link between the amplitude of ongoing oscillations and perception is thought to be causal. However, when the EEG is recorded, both the constrained oscillations and endogenous factors (such as attention) are recorded together. Thus, there is always the possibility that the link found between the amplitude of alpha oscillations and target detection could be explained by mitigating effects of endogenous factors. In chapter 5, we use the WN paradigm to demonstrate the causal role of alpha oscillations. This is possible because only background oscillations are modelled (signal S4) but not the endogenous factors (signal S2).

Finally, while understanding how visual perception and oscillations are linked is interesting in itself, the WN paradigm allows us to go beyond this. Since perceptual echoes are a stable model of the brain's response to visual stimulation, we can use them to query the universal features that link visual perception and oscillations. In other words, we build a classifier that links luminance values (from a white noise sequence) with spectral features to extract a prediction about the perception of a target embedded within the white noise. In chapter 6 we determine the subject independent (i.e. universal) features of perception. In chapter 7 we take this one step further by building a subject specific classifier of perception. Building upon the finding from the previous chapters, we investigate the possibility of building forward models of perception which would allow us to control perception in a subject specific manner.

First however, we need to verify that the approach we want to use in this thesis is valid. Chapter 2 will be dedicated to presenting and evaluating the White Noise Paradigm.

Chapter 2. GENERAL APPROACH AND VALIDATION

In this chapter, we will describe and evaluate the experimental paradigm used throughout this thesis. I will first present the White Noise paradigm. Then, I will appraise its validity by asking 4 questions concerning the underlying assumptions of the paradigm.

2.1 THE WHITE NOISE PARADIGM

The white noise (WN) paradigm is based on linear system analyses, which uses white noise sequences as a means of identifying systems (Marmarelis & Marmarelis, 1978). WN (i.e. random luminance) sequences are used to drive the system to identify the stable relationship between inputs and outputs that is the linear relationship between visual stimulation and brain activity (see introduction, section 1.4.2).

The white noise paradigm is composed of two sessions: during the first session, brain activity is recorded (using EEG) while WN sequences are presented. Then, the Impulse Response Functions (i.e. stable linear relationship, IRF) between stimulation and brain response is extracted by doing a cross-correlation of the two signals (see Figure 2-1; Lalor, Pearlmutter, Reilly, McDarby, & Foxe, 2006; VanRullen & Macdonald, 2012). The IRF can be conceptualized as the brain's filter of the WN sequences luminance information: they measure how the brain responds, on average, to one unit increment in light (Ringach & Shapley, 2004). During the second session, subjects perform a task (in our case, a target detection task), which is embedded inside the white noise sequence. This allows us to directly constrain the background oscillations, while measuring performance too. We use the IRF (recorded in session 1) to reconstruct (rather than record) this background oscillatory activity to the new WN sequences presented (see Figure 2-1) by performing a convolution between the WN sequences and the IRF. It is this reconstructed EEG which is then tested in the different types of analyses, linking for example the phase or the amplitude of the reconstructed EEG to perception. Note that, while we chose to present visual WN (random luminance) sequences here, other types of stimuli could also be used too (such as coloured stimuli for example), as long as the statistical properties of the sequence are preserved. Moreover, the task used during session 2 could also be changed from a detection task to a discrimination task for example.

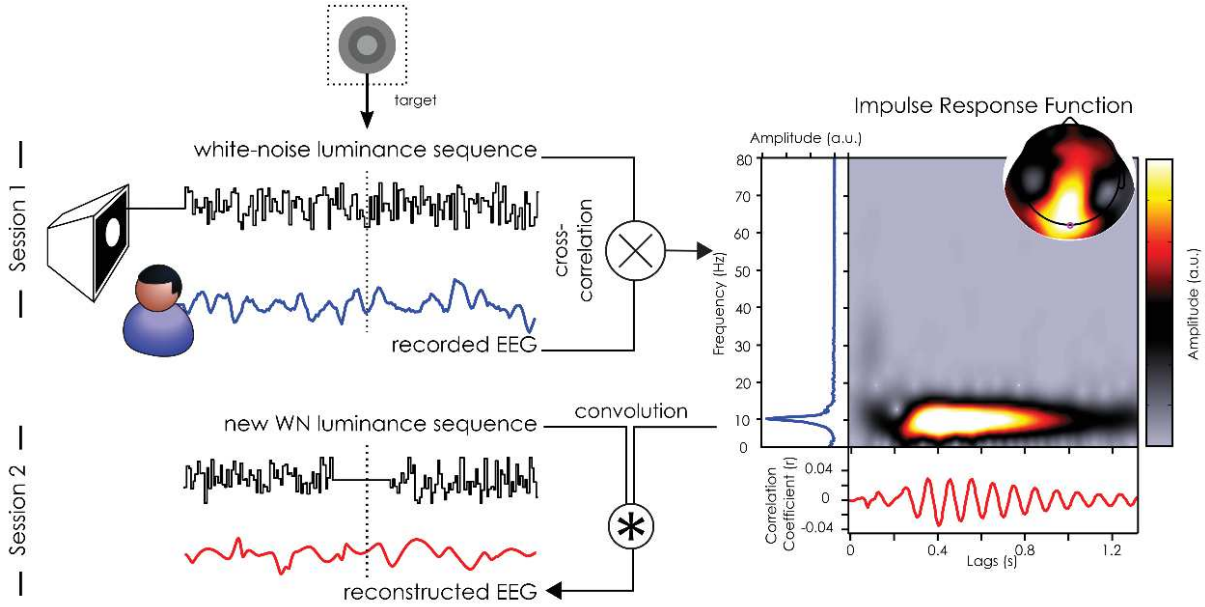

Figure 2-1. Illustration of the White Noise paradigm. The White Noise paradigm is composed of two sessions. During the first session, the impulse response function (IRF) are extracted by cross-correlating the WN sequences with the recorded EEG. Here, an example IRF is shown from one subject on electrode POz, with its time-frequency transform (wavelet), and the frequency spectrum. During the second session, a behavioural task is performed while the WN sequences are presented to constrain the brain response. The IRF (from session 1) can, in turn, be used to reconstruct the brain activity (reconstructed EEG) to any new white-noise sequence by convolution. This allows the direct examination of the causal nature of visual perception.

The strength of the white noise paradigm lies in using the IRF as a model of the brain activity. Reconstructing, rather than recording, the EEG has the advantage of extracting the time course of the phase locked oscillations independently from fluctuations in any of the other brain signals (see description in the introduction, section 1.7.2, table 1). That is, we will be able to test whether the background oscillations have an impact on perception, independently from the impact of all the other signals. Furthermore, it allows us to make direct conclusions about the causal role of oscillations in visual perception. Any modulation in the perception of the target we find as a function of fluctuations in the spectral content of the reconstructed EEG has to be causal in nature. The state of the WN sequences constrain the state of the oscillation, which causal perception to change.

Thus, in this thesis, we will perform all analysis using the reconstructed EEG. In this way, we have access to the instantaneous state of the background oscillations, and can evaluate their relative influence on perception independently of fluctuations in any of the other signals.

2.2 VALIDATION

Using this paradigm, we make a few assumptions that need to be checked before any other analysis is done. Here are the following key points that need to be addressed:

1) How well do the IRF model the brain activity?
2) Are the IRF stable in time?
3) Is there a direct linear relationship between the luminance values of the WN sequences and target perception?
4) Do the white noise sequences drive the detection of targets?

Each of these questions will be addressed in turn in the following sections. First, we will describe the general method used. It applies to all the analyses described in this chapter, as well as analyses described in chapter 6 and 7 (unless stated otherwise).

2.3 GENERAL METHOD

2.3.1 Participants & Procedure

For question 1 and 4, we analysed the behavioural data from the 20 subjects (10 women) who took part in the main EEG experiment of this thesis. The mean age was 28 years old (SD: 3.97 years, 23 to 39 years old) and 15 of them were right handed. The participants underwent two testing sessions: during the first session, we recorded the EEG response to the white noise sequences while in the second session we only recorded the behavioural responses. 11 subjects took part in a session 3 where only behavioural responses were recorded.

For question 2, 17 subjects (7 women) with at least 2 EEG recordings of the IRF were included in this analysis. In total, we had 44 IRF recorded: there were 3 subjects with 4 recording sessions, 4 subjects with 3 recording sessions and 10 subjects with 2 sessions. Note that we only included sessions where the exact same stimuli parameters had been used i.e. the same disk eccentricity and size, as this was found to influence the shape of the IRF (see chapter 3). However, different (new) WN sequences were presented on all sessions. A session from one of the subject had to be excluded because a medical condition arose between the original and the given session that might have changed their oscillatory signature.

To answer question 3, we used the data from a pilot experiment and a control experiment. Twelve subjects (6 women, mean age was 25.7 years, SD: 1.6 years) were

included in the pilot experiment. They underwent two testing session. For the control experiment, 6 participants (3 women) who had not taken part in the previous experiment were included.

All subjects had normal or corrected-to-normal vision, and reported no history of epilepsy. They all gave written informed consent before the experiment, which was carried out in accordance with the local ethics committee. They were compensated (with vouchers) for their participation in the experiment.

Each session (whether EEG or psychophysics) lasted about 1 hour, and was composed of 8 blocks of 48 white noise sequences. The experiment was self-paced: participants started the blocs and trials by pressing the space bar.

2.3.2 Stimuli

For all experiments, 6.25 seconds long white noise (i.e. random luminance) sequences were used as stimuli. The sequences were presented on a CRT monitor with a resolution of 640 by 480 pixels, on a black background. The monitor refreshrate was at 160 Hz, consequently, the WN sequences had a flat power spectrum between 0 and 80 Hz. They were presented within an overhead peripheral disk at 7° eccentricity from fixation and with a diameter of 7° of visual angle on a black background. Participants had to covertly attend to the disk while maintaining fixation on a central point (0.1° of visual angle). To insure that participants covertly paid attention to the disk, near-perceptual threshold targets were embedded within the sequences. Participants were told to detect them as fast as possible by pressing a key ("b"). From 2 to 4 targets could appear in each trial. They were presented for 1 frame only and were composed of a darker surrounding annulus and a lighter circle at the centre presented on a grey background (i.e. medium luminance value). The contrast between the outer and inner parts of the targets was adjusted over the first 100 targets (~30 trials) using the *Quest* function (Watson & Pelli, 1983) in the PsychToolBox (Brainard, 1997) to reach a contrast where participants perceived 50% of the targets. The achieved contrast was then kept constant for the remainder of the experiment.

2.3.3 EEG recording, pre-processing and reconstructing

The brain activity was recorded using a BioSemi 64 channels (1024 Hz sampling rate) EEG with 3 external channels recording the horizontal and vertical ocular movements. The

pre-processing steps were the following: 1) Down-sampling of the EEG to 160 Hz to match the presentation rate of stimuli and thus facilitate the cross-correlation of the two signals, 2) Rejection and interpolation (nearest neighbours) of noisy channels if necessary, 3) Notch filtering (between 47 to 53 Hz) to remove any artefacts due to power line. This was an issue only for a few subjects, but was applied to all in order to standardize pre-processing across subjects, 4) Re-referencing to the common average, 5) Removing drifts through high-pass filtering (>1Hz), 6) Epoching from -0.25s to 6.5s around each WN sequence, 7) Baseline correction (removing. the mean activity from -0.25s to 0 before trial onset), 8) Manual artefact rejection (remove whole epochs as needed to get rid of eye blinks and muscular and movement artefacts).

Once the data had been pre-processed, the impulse response functions were extracted by doing a cross-correlation between the normalized (z-scored across time points) white-noise sequences and the concurrently recorded normalized (z-scored across time points) EEG, for time lags -0.25s to 1.5s (see Figure 2-2):

$$cross - correlation(ch, t) = \sum_{T} WN(T) \cdot EEG(ch, T + t)$$

where *WN(T)* and *EEG(T)* are the z-scored stimuli and EEG signal pair for a given channel at time T, and *t* is the lag between stimulus and EEG signal. The cross-correlations for each trials was then averaged together. The cross-correlation was computed using the *xcorr* function from MATLAB and custom made scripts. Note that using the mTRF toolbox (Crosse et al., 2016) gave identical results. To reconstruct the EEG, we performed a convolution between the white noise sequences presented during the second session and the IRF using the *conv* function in MATLAB for each channel.

2.4 DESCRIPTION OF IMPULSE RESPONSE FUNCTIONS

First, we will start with a brief description of the IRF recorded for the 20 subjects of the main EEG experiment (i.e. data analysed in this chapter and in chapter 4 to 7). As described previously (VanRullen & Macdonald, 2012), the perceptual echoes are characterized by a long lasting reverberation in the alpha band, strongest over the central-occipital channels (see time frequency plots and topographies on Figure 2-2). There is also information in the theta band over fronto-central channels (see Figure 2-2).

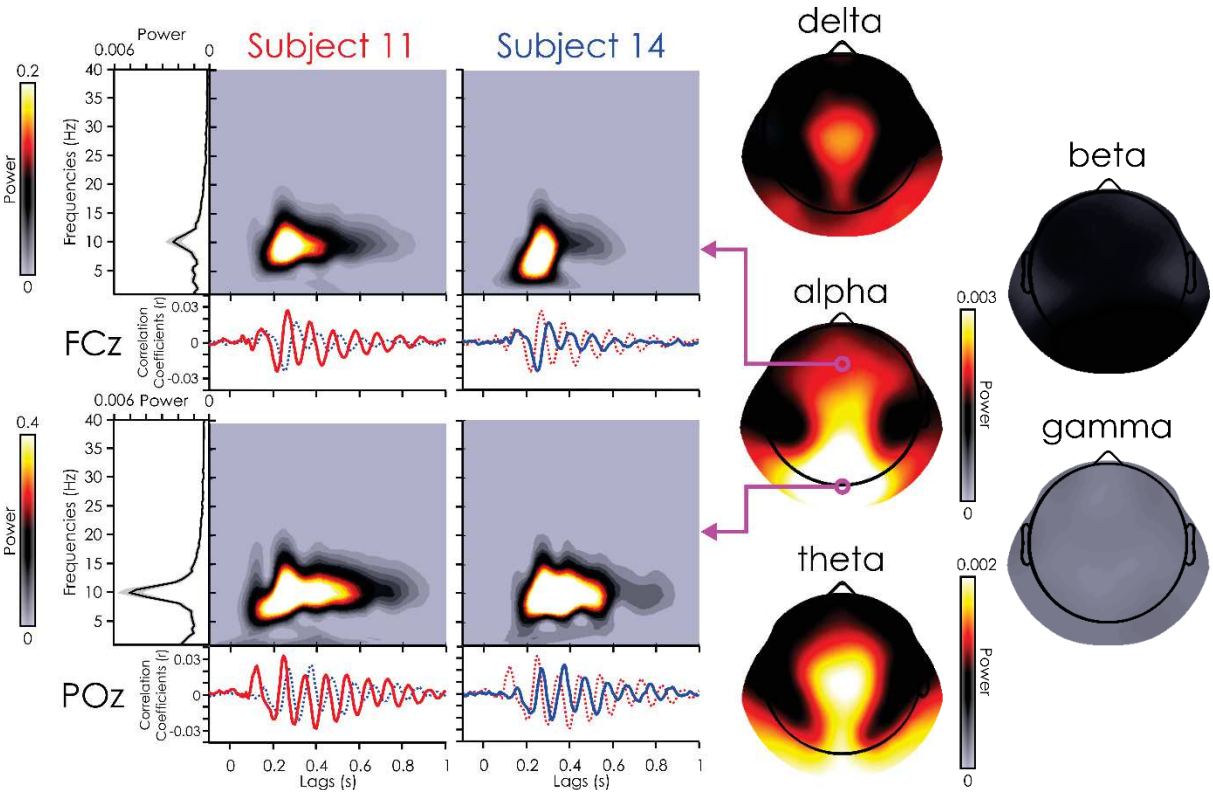

Figure 2-2. Description of echoes. Time frequency transforms for echoes on FCz (top) and POz (bottom) for 2 subjects (subject 11 in red and subject 14 in blue) with the time representation of the echo (below) and the mean spectra across subjects(left; mean in black and standard deviation across subjects in grey). The topographies represent the mean amplitude across subjects for each frequency band from 250ms to 1s. Note the difference in colour bar for the alpha amplitude compared to the other 4 frequency bands.

The echoes represent the correlation between the stimulus intensity and the EEG response at different lags. Moreover, this relationship is subject specific: each subject has a specific resonance frequency and phase relationship to the white noise sequence (see red and blue echoes on Figure 2-2). For example, for electrode POZ, 0.4s after the flash, the EEG of subject 11 is negatively correlated with the flash intensity while the EEG of subject 14 is positively correlated.

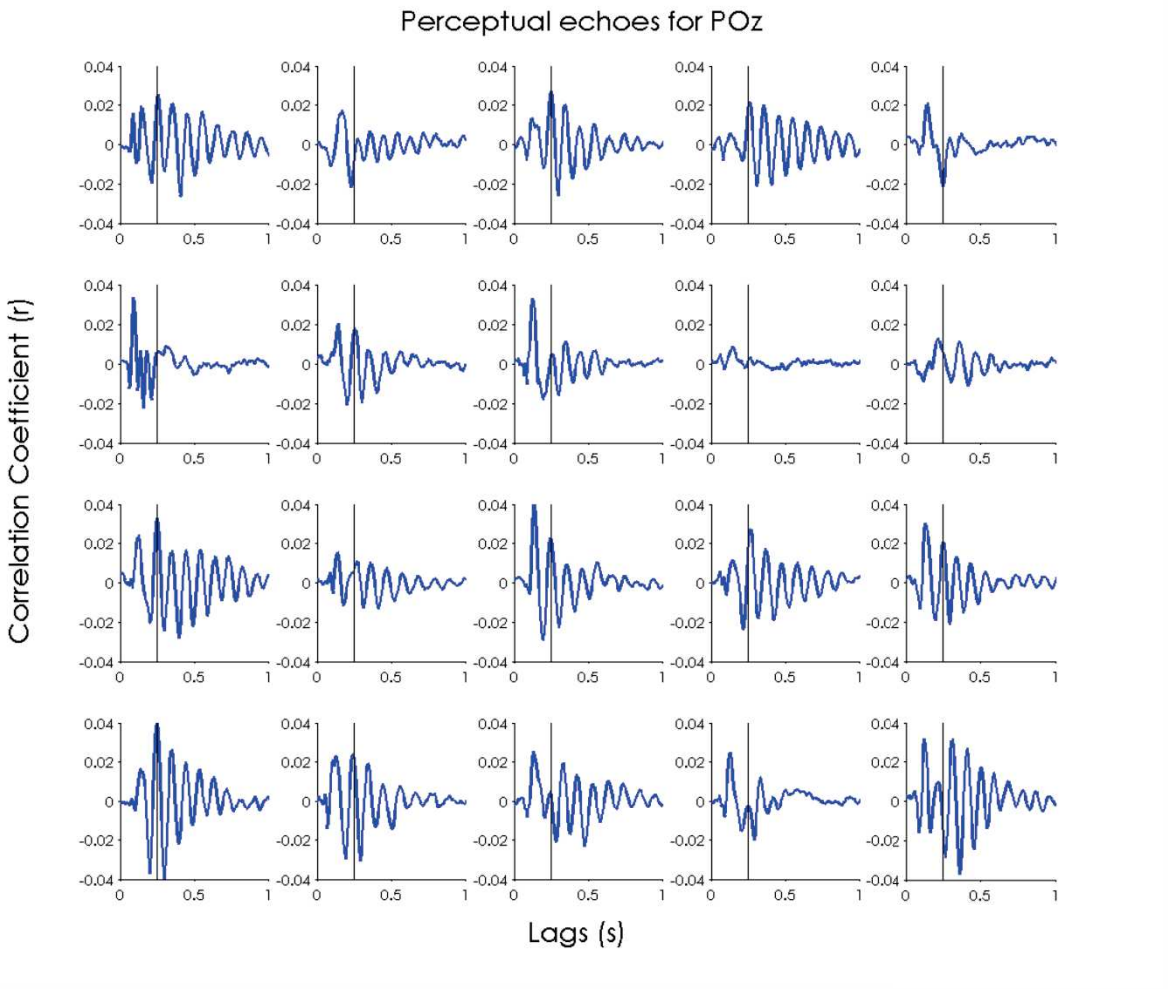

Figure 2-3. Echoes of each subject on the parietal central channel. IRF for each subject on POz. There is a lot of inter-subject variability in the echoes: some subject have very strong and long echoes (e.g. subject 1) and some almost none (subject 6).

These echoes could be divided into two parts (Ilhan & VanRullen, 2012): an "early" component (between 0 and 250ms) corresponding to the VEP (or the VESPA, Lalor et al. 2006) and a "late" component (between 250ms to up to 1s) corresponding to the "echo" (VanRullen & Macdonald, 2012). This distinction between "early" and "late" IRF is relevant to chapter 3 and chapter 7. In chapter 7, we want to use the IRF as encryption keys for subject's perception. It has been shown previously that the exact temporal features of the VEP (corresponding to the early part of the IRF) can be used to differentiate between subjects with an accuracy as high as 86.5% (Power, Lalor, & Reilly, 2006). Does the late part of the IRF (i.e. the perceptual echo) also contain subject specific information? From the spectral analysis, it seems to be the case: each subject has their own resonance frequency. Moreover, we can see from the plots that the echoes vary considerably between subjects in terms to duration: some subjects have very short or inexistent echoes, while others have longer one. Finally, there are

differences in the phase of the echoes: at 400ms for example the EEG of S11 is negatively correlated with the luminance presented before while that of S14 is positively correlated (see Figure 2-3).

2.5 QUESTION 1 – HOW WELL DO THE IRFS MODEL THE BRAIN ACTIVITY?

The extraction of impulse response functions relies on the system identification methods from *linear* system analyses. This makes the assumption that the model under study can, at least in part, be approximated by a linear system with the white noise sequences as inputs and EEG as outputs (Lalor, Pearlmutter, et al., 2006). If the brain were linear, the IRF would be the best possible model of brain activity. However, we know that the brain is highly nonlinear. Therefore, we tested how good of a model the IRF is. In other words, how well does our reconstructed EEG model the recorded EEG? We hypothesize that our ability to reconstruct the activity will be better for specific frequency bands: perceptual echoes have a very strong alpha component, which is likely to lead to the best prediction of the recorded EEG activity.

2.5.1 Method

To test how good a model the IRF was, we did a correlation between the reconstructed and recorded EEG. For this, we used the data from the first session of the experiment for which we recorded the EEG. We used a 10-fold cross validation approach: 90% of trials were used to extract the IRF, which was in turn used to reconstruct the EEG for the remaining 10 % of trials. Note that because subjects had different number of trials available (due to artefact rejection), we sub-sampled the number of trials of each subject to that of the subject with the lowest number of available trials (310 trials). A Pearson correlation was then applied between the reconstructed EEG and the EEG actually recorded on the remaining 10% of (independent) trials. Using this cross-validation strategy ensured that both the recorded and reconstructed EEG were available for the same sets of trials, while avoiding any circularity in the analysis: using the same trials to compute the IRF and the reconstructed EEG could have led to spuriously high correlations between reconstructed and recorded EEG. Moreover, this also reduced the influence of a sampling bias, as repeatedly computing the correlation allowed us to get a better estimate of the true underlying correlation between the two signals. This (cross-

validated) correlation was also evaluated for 5 different frequency bands (FIR filter, delta: 2 – 4 Hz; theta: 4 - 8 Hz; alpha: 7 – 14 Hz; beta: 14 – 28 Hz; gamma: 30 – 60 Hz) to assess whether the prediction would be better for a specific frequency band.

Since the correlation coefficients were not normally distributed (as revealed by a one-sample Kolmogorov-Smirnov test), a Fisher Z transform was applied to the data before averaging. It was computed as follows:

$$Z = 0.5 * ln\left(\frac{1+r}{1-r}\right)$$

To extract the correlation coefficient values for plotting purposes, the inverse Fisher z-transform was applied. Only the correlation coefficients at the strongest channel are presented across subjects. To estimate the statistical validity of the correlation, a one sample t-test against zero was applied on the averaged Z transformed coefficients across all repetitions and subjects for each channel. The p-values were corrected for multiple comparison using a FDR correction across channels. Note that using a non-parametric test (instead of a Fisher Z transform and a t-test) gave equivalent results.

As a comparison for our linear model, we used two other models of perception. First, we tested how much of the variability in the signal could be explained by the event related potential (ERP) to the target embedded in white-noise. Here, we computed the target evoked response (rather than IRF) as a model for the target-evoked activity. Finally, for comparison purposes, we also wanted a measure of how "noisy" EEG data is (i.e., in typical ERP paradigms without ongoing white noise sequences). We tested this using a separate dataset from Busch and colleagues (2009), in which isolated targets were presented on a static background. The target-evoked ERP was again used as a model of evoked EEG activity. In both conditions (ERPs from targets embedded in white noise; ERPs from isolated targets), the ERPs were extracted on 90% of the trials and then convolved with a sequence of target onsets for the remaining 10% of trials, in a 10-fold cross-validation approach. Consequently, the same correlation method between recorded and reconstructed EEG (modelled here by the ERP) was applied as described above for the IRF-based EEG reconstruction model.

2.5.2 Results

We wanted to quantify the goodness of fit of the IRF as a model of brain activity by systematically correlating the recorded and the reconstructed EEG. We also compared the performance of the IRF against a standard model of brain response, the VEP.

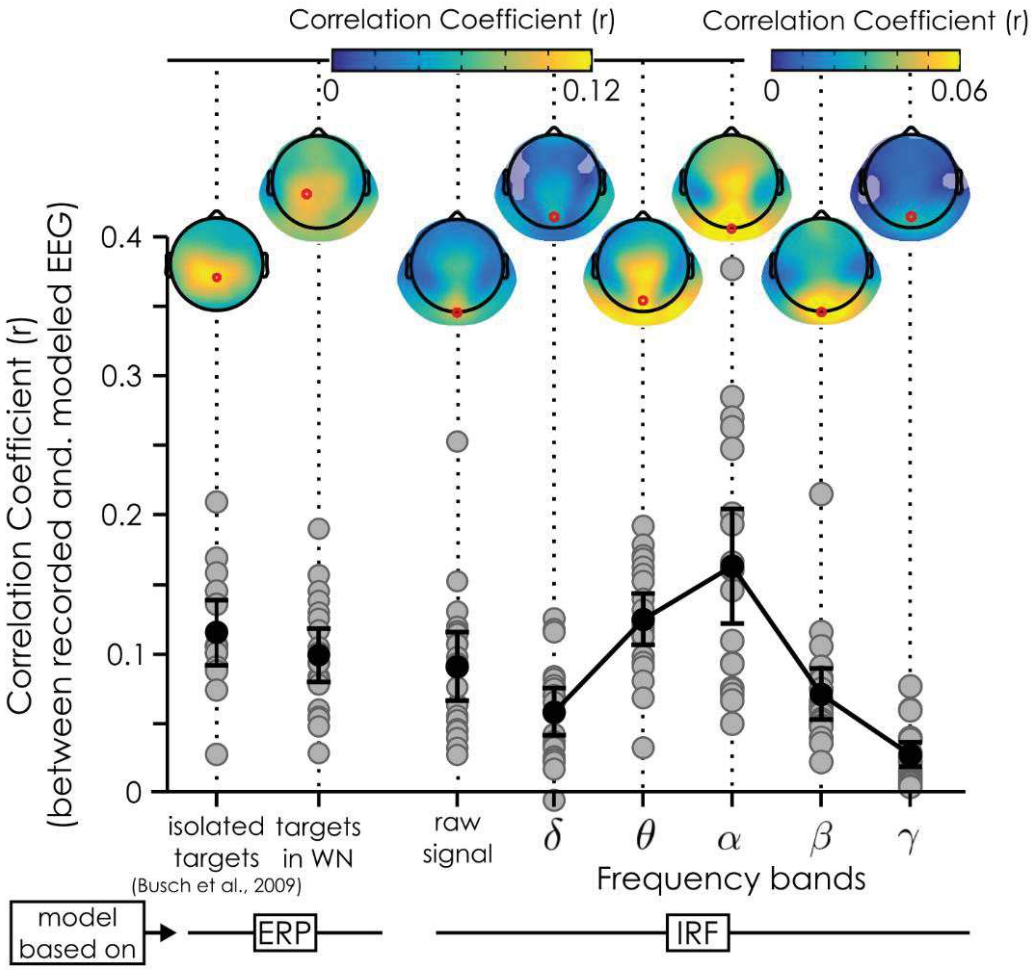

Figure 2-4. Correlation between reconstructed and recorded EEG. Two models of brain EEG activity were tested: one based on the ERP to targets (left) and the other based on the IRF to white noise sequences (right). In both cases, we systematically correlated the prediction of the model with the single trial recorded EEG. For the ERP models, the "isolated targets" coefficients are based on the dataset presented in Busch et al. (2009), while the "targets in white-noise" are from our own dataset, with the ERP extracted relative to the targets embedded in white noise. For the IRF models, we used the IRF to model the brain response to the white noise sequences, and the correlated reconstructed and recorded EEG data using the raw signals, as well as signals filtered in different frequency bands (delta: 2-4Hz; theta: 4-8Hz; alpha: 7-14 Hz; beta: 14-28Hz; gamma: 30-60Hz). The grey dots represent the mean coefficient for each subject (1 dot per subject) across cross-validation runs at the maximum electrode (red dot on the topographies); the error bars represent the 95% confidence interval of the mean (black dot) coefficient across subjects. The topographies represent the mean correlation coefficients across subjects and cross-validation runs. Shaded areas

represent channels not significant after fdr correction. Note the difference of colour scales for the beta and gamma correlation coefficients relative to the other topographies.

Using the IRF as a model and the "raw" reconstructed EEG, we found that the distribution of mean (z-transformed) correlation coefficients was significantly different from zero on all channels. The maximum correlation coefficient (averaged across subjects) was reached on channel Oz with an r of 0.091 ($t(19) = 7.78$, $p=2.51*10^{-7}$, 95% CI for r = 0.066-0.115). We also investigated whether filtered signals would yield a better correlation depending on the frequency band. We found that the correlation strength was strongest for the alpha band on channel Oz (mean r = 0.163, $t(19) = 8.21$, $p=1.14*10^{-7}$, 95% CI for r: 0.121-0.204), closely followed by the theta band on channel POz (mean r = 0.125, $t(19) = 13.69$, $p=2.69*10^{-11}$, 95% CI for r: 0.105 to 0.144). The other frequency bands had lower correlation coefficients, with respectively a mean r of 0.058 (channel POz) for the delta band ($t(19) = 7.28$, $p = 6.60*10^{-7}$, 95% CI for r: 0.041 to 0.075), a mean r of 0.071 (on channel Oz) for the beta band ($t(19) = 7.80$, $p=2.42*10^{-7}$, 95% CI for r: 0.052 to 0.090) and a mean r of 0.027 (on channel POz) for the gamma band ($t(19) = 6.16$, $p=6.37*10^{-6}$, 95% CI for r: 0.018 to 0.037).

As a benchmark measure for our recorded EEG performance, we decided to use two alternative models based on computing the evoked response potentials. First, we used ERPs recorded to targets embedded in white noise sequences (our own paradigm). The correlation coefficient reached its maximum across subjects over the left central parietal channel (CP1) with an r of 0.099 ($t(19) = 10.92$, $p=1.25*10^{-9}$, 95% C.I. for r = 0.080 to 0.118, see "targets in WN", Figure 2-4). Secondly, we also evaluated the amount of signal variability that can be explained by the ERP to isolated targets in a more "typical" visual-evoked potential paradigm (i.e., without the concurrent white noise stimulation), using the data from Busch and colleagues (2009). This was done as a way to estimate the "noise" in a typical EEG setting, for comparison purposes. A one-sample t-test revealed that the distribution of the mean (z-transformed) correlation across subject reached a maximum over channel (Pz) with an r of 0.115 ($t(13) = 9.57$, $p = 2.97*10^{-7}$, 95% C.I. for r: 0.089 to 0.141, see 'isolated targets' in Figure 2-4).

Note that the exact same analysis performed on an independent dataset (i.e. pilot data used to answer question 3) gave very similar results (see Figure 2-5).

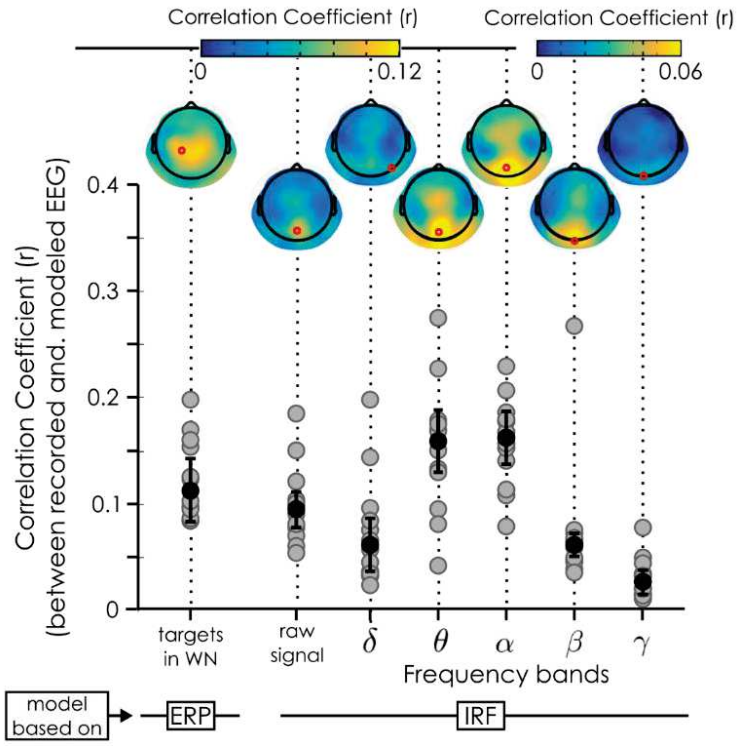

Figure 2-5. Correlation between the reconstructed and recorded EEG for the pilot experiment (N=12 subjects). Same conventions as Figure 2-4. It is important to note that half of the subjects in the pilot experiment were also in the main experiment although the white noise sequences used to record the IRFs were different.

2.5.3 Discussion

We showed that the IRF is a useful model of EEG, though it is far from perfect (i.e. r<0.2). The best reconstruction of brain activity was in the alpha band, closely followed by the theta band. However, the obtained EEG reconstruction was still clearly noisy. This is not surprising: it is well known that EEG is a noisy recording method (Picton et al., 2000). Despite this, the reconstruction error was on par with that observed using more traditional models of brain activity, i.e. ERP-based models of target-evoked activity. We used the ERPs extracted from two different paradigms, either to targets presented in WN (in our experimental paradigm), or from previously published data as a comparison. Both models could explain the same amount of variability in the recorded EEG signal as the IRF-based reconstructed EEG. These results are in line with what has been found before: using the VESPA as a model for brain activity, a mean correlation (on channel Oz) of 0.084 between the reconstructed and recorded EEG was found (Lalor, 2009). Interestingly, using a quadratic model of the IRF yielded only marginally better predictions: it reached on average across subjects a r of 0.097 (Lalor, 2009).

2.6 QUESTION 2 – ARE THE IRF STABLE IN TIME?

Next we asked whether the IRF were stable in time. If we want to use the IRF to reconstruct the brain activity at various time points, we need to know how long we can re-use the same IRF. In their original paper, VanRullen & Macdonald (2012) tested the stability of echoes at 6 months interval and found a high positive correlation within subjects. However, the number of subjects was relatively small (N = 3) and the stability of the echoes has not been tested at longer time delays. For the forward models presented in chapters 6 and 7 to work, we need the prediction to be stable at longer time intervals, or at least to know how often the IRF need to be recorded. Thus, a longitudinal analysis of the echoes recorded at longer delays (from various experimental paradigms) was performed.

2.6.1 Method

We correlated the IRF of the same subjects recorded on two different sessions using a Pearson correlation and kept the delay between sessions as the independent variable. This was done separately for each channel and subject, systematically correlating all possible pairs of sessions, yielding 44 "within subject" comparisons for each of the 64 channels. We also computed the correlation between echoes from different subjects to evaluate what the "chance correlation" between the echoes could be based on similarity in echo peak frequency. We thus systematically correlated the echoes of all possible different pairs of (different) subjects. This yielded a total of 906 "between subjects" comparisons for each of the 64 channels.

We restricted our analysis to the electrode showing the strongest alpha amplitude across subjects POz (see Figure 2-2). To evaluate the significance of the effect, we created linear regression models (using the function *regress* in MATLAB) which attempted to explain, independently for the within and between subjects comparisons, the correlation coefficients as a function of the delays between sessions. Using these linear regression models, we also wanted to know the number of years it would take for the within subject regression line to cross the between subject intercept and how long it would take for the two predicted lines to cross.

2.6.2 Results

Over the 17 subjects and the 44 sessions recorded, we found that, regardless of the delay between recordings, there is a very strong correlation between two echoes recorded within subjects for channel POz (r =0.82, 95% confidence interval for r: 0.79 to 0.86), while the between subject one was much weaker (r=0.34,95% confidence interval for r: 0.32 to 0.37). We can see from the topography (see Figure 2-6), that this stronger correlation within subjects is found mainly over parietal and occipital channels.

To see whether the delay had an impact on the correlation between echoes recorded on different sessions, a linear regression was used, assessing how correlation coefficients were affected by the delay between sessions for the within and between subjects comparisons independently. For the within subject correlations, we found a significant effect ($F = 7.3$, $p = 0.0101$, $R^2 = 0.16$) with an intercept at 0.88 (95% CI = 0.79 – 0.97) and a slope of -0.0002 (95% CI = -0.00035 to $-5*10^{-5}$). This means that the correlation coefficient between the echoes recorded on two different sessions for the same subjects is predicted to be 0.88 – 0.0002 per 1 day's delay (e.g., it would be predicted to decrease to 0.86 after 100 days). A significant effect was also found for the between subject correlations ($F = 6.1$, $p = 0.013456$, $R^2 = 0.0064$), with an intercept of 0.34 (95% CI = 0.31 – 0.38) and a slope of $-8.1*10^{-5}$ (95% CI = -0.00014 to $-1.7*10^{-5}$). This means that the correlation coefficient between the echoes recorded on two different sessions for two different subjects is predicted to be 0.34 – 0.000081 per 1 day's delay (e.g. it would decrease to about 0.332 after 100 days).

To better understand the stability of the echoes in time, we computed two metrics. First, we computed the number of years it would take for the within subject regression line to cross the intercept of the between subject regression. We found that this would take 7.35 years. Second, we also computed the time it would take for the within subject regression line to cross the between subject regression line and found that it would be 12.316 years.

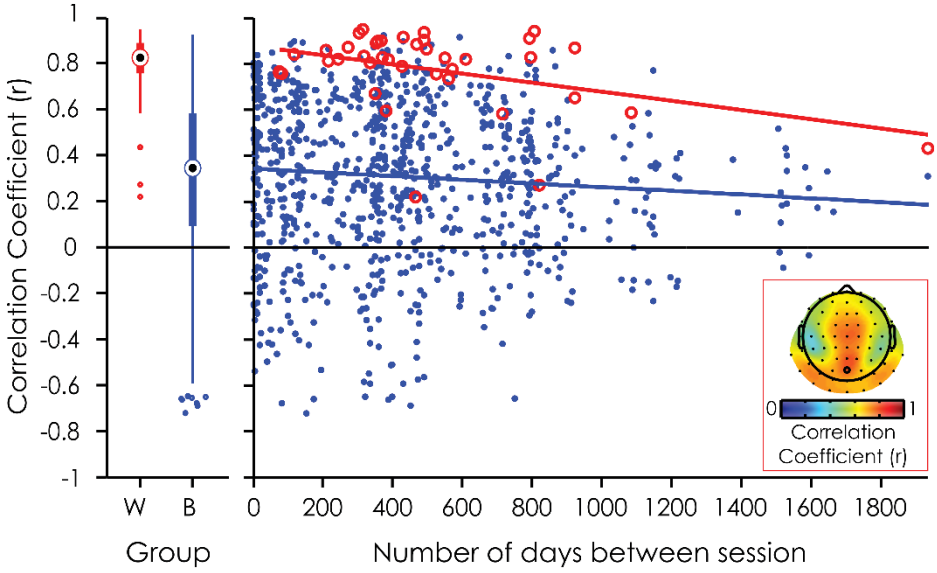

Figure 2-6. Scatter plot and linear regression lines of the correlation between echoes as a function of the delay between sessions for channel POz. The red circles represent the correlation coefficient within subjects, while the blue dots represent the correlation coefficients between subjects. The boxplots represent the median and distribution of the correlation coefficients within (red) and between (blue) subjects on POz, regardless of the delay. The topography represents the within subject coefficients averaged across subjects and delays.

Despite the very small size of the effect (0.000081 with an R^2 of 0.0064), the fact that the linear regression is significant for the between subject correlation coefficients is unexpected. Since the recordings compared are obtained from distinct subjects, and the order in which we ran the subjects or assigned them to our different experiments was arbitrary, we would have expected this correlation to remain relatively stable. There should be no reason why the two echoes recorded on different subjects are more alike if they are recorded around the same time than a few months or years apart. This effect could be explained, however, by variations of many factors between the sessions: the data was collected from different experimental paradigm (i.e. in some experiments other tasks were carried out by the subjects), different stimulation programs, different numbers of trials, different experimenters or even EEG sets and recording amplifiers.

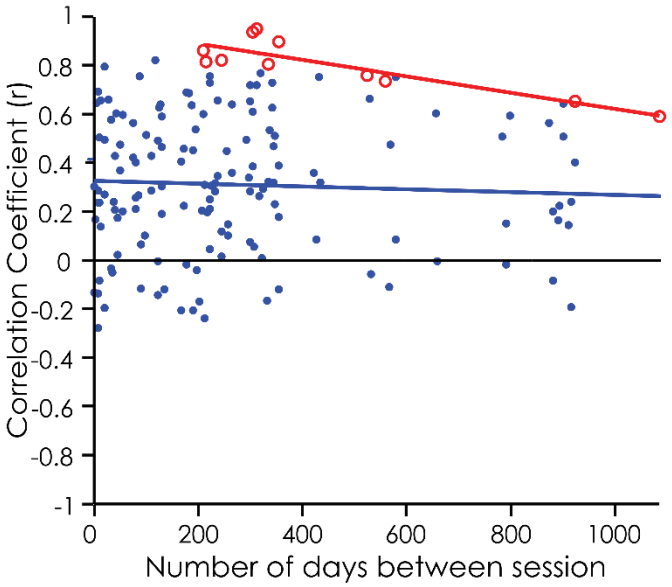

Figure 2-7. Scatter plot and linear regression lines only including subjects for whom both sessions were run by the same experimenter. Here again, within subject coefficients and regression line are in red while the between subjects coefficients and line are in blue. In this case, the between-subject slope is not significantly different from zero.

To be sure that this significant effect might just reflect these secondary factors, we ran a control analysis. The linear regression analysis was re-run, but only including the participants run by the same experimenter and on the same setting. Whereas less data points were available in this analysis, and with shorter delays between recordings, the results (see Figure 2-7) reveal that while the within subject correlation coefficients are still significantly influenced (predicted r = 0.95 (95% CI = 0.88 to 1) – 0.00033 (95% CI = -0.00048 to -0.00019) * days) by the delays ($F = 27$, $p<0.0006$, $R^2=0.75$), this effect has now disappeared for the between subject coefficients ($F=0.85$, $p=0.3582$, $R^2=0.0053$).

2.6.3 Discussion

Comparing the correlation between echoes recorded within or between subjects, we find that, while the IRF recorded within subjects are highly strongly correlated, the correlation between subjects is much smaller. This variability between subjects could potentially be explained by variability in the "early" or the "late" part of the IRF. It has been shown that the VEP, corresponding to the "early" IRF is highly variable between subjects. Steven Luck for example remarked that in ERP shape he generally finds more variability between subjects than within subjects (Luck, 2005b, pp. 18–19). Part of this between subject variability he ascribes to the idiosyncratic folding of the visual cortex, which is thought to be the generator

of the ERP (Di Russo, Martínez, Sereno, Pitzalis, & Hillyard, 2002). As the early visual cortex has also been shown to be the potential source of the IRF (Lalor et al., 2012), this could explain part of the poor correlation of between perceptual echoes of different subjects. Furthermore, the variability between subjects could also be determined by the "late" part of the IRF. In fact, different subjects have different peak frequency (VanRullen & Macdonald, 2012).

Moreover, the IRF recorded within subjects are very stable in time: the correlation between two given echoes decreased by about 0.07 point per year. The correlation coefficient for the between subject comparison was much smaller: the linear model predicted that it would take 7.35 years before the within subject coefficients would become as bad as the initial between subjects correlation. Furthermore, the linear model also predicted that it would take 12.316 years before we are completely unable to differentiate the two groups. Therefore, we conclude that IRFs can be used to reconstruct brain responses in the absence of recorded EEG in a reliable manner. This is especially important for chapters 6 and 7, where we develop forward models of perception. It is crucial that we know how long the echoes are stable for, and thus for how long the prediction would hold before we have to record the IRF again.

2.7 Question 3 – Classification Image: can luminance directly modulate target perception?

Studies have shown that the noise content around the presentation of a target could directly influence its detection (Neri & Heeger, 2002). Finding a direct relationship between stimulus and perception would render any connection between prediction EEG and perception trivial (as potentially driven by the link between stimuli and perception on the one hand and between stimuli and brain response on the other hand).

2.7.1 Method

First, we used the noise-image classification (Ahumada, 2002) to extract the mean stimulus properties around each target. The stimuli were epochs in 1.6 seconds long epochs around each presented targets. Practically speaking, the mean luminance values were

computed separately for the detected versus missed targets trials for each subject. An independent samples t-test was applied to compare the distribution of luminance values for hits and missed targets epochs for each subject independently. A FDR correction across time and subject was applied to correct for multiple comparisons in each session. Session 1 data revealed that immediately surrounding luminance values influenced target visibility. This influence could have masked the (potential) effect of oscillatory phase. Therefore, in session 2 we decided to remove any luminance fluctuations around targets, setting 14 frames before (i.e. 87.5ms) and 11 frames after (i.e. 68.75ms) the target to the same medium grey value.

Using a similar design and protocol, we ran a control experiment to assess whether the fluctuation-free periods could be used by subjects (N=6) to detect the targets. This time however, we created catch trials by not presenting the targets in 1/3rd of the flattened luminance time windows. As with the main experiment, any button press after a target (from 150ms to 800ms) was counted as a hit. The same response window was used for catch trials: any button press within this 650ms time window after the moment where a target would have been was counted as a false detection. The percentage of catch trials where a subject pressed the button was computed, similarly to the percentage of detected targets. Separately, we also assessed the false alarm rates of subjects, i.e. button presses falling outside of response time windows for either presented targets or "catch trials". To get an estimate of the false alarm rate that is comparable with the other detection rates, we counted the number of 650ms long time windows outside of target times and divided the number of false alarms by this number.

2.7.2 Results

The systematic relationship between luminance values in the WN sequences and visual perception was probed using classification images (Ahumada, 2002) to test whether the luminance values had a direct influence on visual detection. We wanted to ensure that there were no direct linear effects (i.e. masking or unmasking) of the luminance on the perception of targets, as these might hide any phasic influence of the constrained brain oscillations on behaviour.

In a pilot experiment, we presented the WN sequences and used classification images to separate the mean luminance values between seen and missed target trials. There seems to be a large "masking" effect of the target by the luminance values presented around the target (at time t = 0ms, see Figure 2-8). If higher luminance values are presented 2 frames before and/or 1 frame after the target, and/or a weaker luminance value 3 frames before, then the

target will be less visible. The exact opposite pattern leads to better-than-average perception. Thus, there is a direct linear relationship between the properties of the luminance sequence and the visibility of the target (see Figure 2-8).

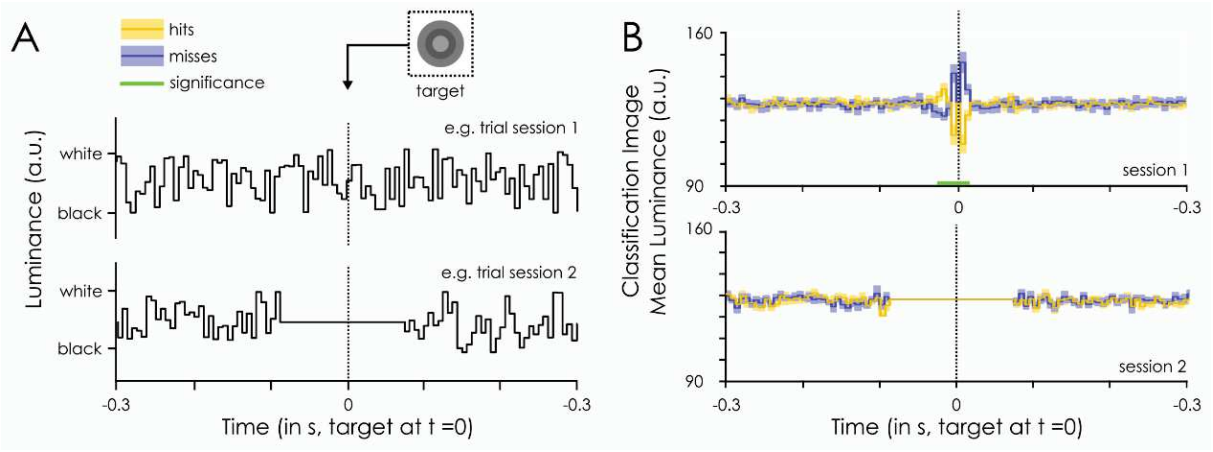

Figure 2-8. Classification image and example WN sequences for session 1 and 2. A. Example of white noise sequences used in session 1 (top) and in session 2 (bottom). The target was presented at time 0. B. Classification image for the hits (yellow) and the misses (blue) for session 1 (top) and session 2 (bottom). Darker areas represent the mean across subjects, the shaded areas represent SEM across subjects. The only significant effect between hits and misses was found in session 1 (green line) just around target presentation (-18.75,-6.25ms, 6.25ms and 12.5ms).

Given this direct linear relationship, there is no need to look at brain signals: we need only investigate the effect of luminance masking on perception. This doesn't tell us anything about how and why perception is linked with oscillations. Thus, we decided to remove the linear relationship by removing the random luminance fluctuations around the target. After further piloting, we decided to suppress (i.e. present grey frames) 14 frames before and 11 frames after the target. Consequently, there were no random luminance fluctuations for 162.5ms around the target (see Figure 2-8A). This effectively removed the masking effects (see Figure 2-8B), and did not influence the power spectrum of the stimuli (i.e. it was still flat). In the main EEG experiment, we thus decided to keep the intact sequences for the EEG recording (i.e. in session 1) and present the altered sequences to measure the behavioural response (in session 2).

Once this problem was overcome, we still wondered whether the subjects could use these moments of suppressed luminance fluctuations around the targets to detect them. In

order to verify that this was not the case, we ran a control experiment. We used the same paradigm as before, only for 1/3rd of the fluctuation stops, we did not include a target, creating "catch trials" of sorts. We compared the detection of catch trials to targets and false alarm rates using Student's T-test. On average, people did not significantly detect the catch trials more than they did false alarms (mean % detection of 8.94% versus 6.30% respectively, $p = 0.646$). The hit rate (mean = 52.80%) was significantly greater than both the catch trials detection rate and the false alarms ($p < 0.002$).

2.7.3 Discussion

We wanted to test whether the detection of targets embedded within the white noise sequences could be directly, linearly linked to the luminance values in the WN sequences. Studies have shown that the random noise around a target could influence perception (Neri & Heeger, 2002). Classification image techniques revealed that the luminance values flashed around the target's presentation could indeed mask or enhance its detection. Since we want to better understand how the constrained background oscillations influence performance, we decided to remove these effects by suppressing the luminance fluctuations around targets so as not to hide any potential impact of the oscillations' features (i.e. amplitude or phase) on perception. In the second session of the experiment, we suppressed the random luminance fluctuations for 162.5ms around the presentation of the target, which effectively removed any direct linear influence of the luminance on visual perception.

In a separate control experiment, we found that the periods of suppressed luminance fluctuation were not perceived by the subjects and could not be used to detect the target. Consequently, we decided to present the "unaltered" WN sequences during session 1, to extract the IRF and the "flattened" sequences during session 2 to investigate their influence on behaviour. Once the direct linear influence of luminance on perception had been removed, we asked whether we could still reliably modulate the behaviour of subjects using white noise sequences.

2.8 Question 4 – Do the white noise sequences drive the detection of targets?

Finally, we wanted to verify that the detection of targets could be influenced by the white noise sequences. If we are to uncover a link between the constrained oscillations and perception, perception needs to be influenced in the first place. This is especially relevant for chapters 6 and 7, where we are trying to build a classifier or "predictor" of the behaviour of subjects based on the spectral features of the reconstructed EEG. For example, in chapter 6, we are trying to predict, for a given target, whether it will be seen or not: this is a subject independent (or universal) approach. For this, we build a "universal" classifier, extracting the subject independent features that link perception to detection. In chapter 7, we build a predictor based on subject dependent features that will be able to optimize the presentation of targets so that one subject might perceive it, but not the other. For these to work, the perception of subjects has to be systematically, predictably influenced by the WN sequences. If there is absolutely no agreement rate between or within subject, seeing the target once cannot be used to tell whether the subject will see the target again (i.e. we cannot build models to predict perception, see Figure 2-9). If a WN sequence reliably influences perception, that subject would be the best predictor of him-self: presented with the same targets, embedded within the same WN sequences his perception should "agree" (i.e. either both seen or both missed, see Figure 2-9). Thus, in this case, we can wholly predict whether a target will be seen again (i.e. "good" model) or (i.e. "bad" model on Figure 2-9). This nomenclature of "good" and "bad" refers to the detection of the target. The difference between the "good" and "bad" prediction (see Figure 2-9) is what we call the "expected performance modulation". This reflects how much modulation in performance we can expect to see in future chapters, this is the upper limit for the results presented in chapter 6 and 7.

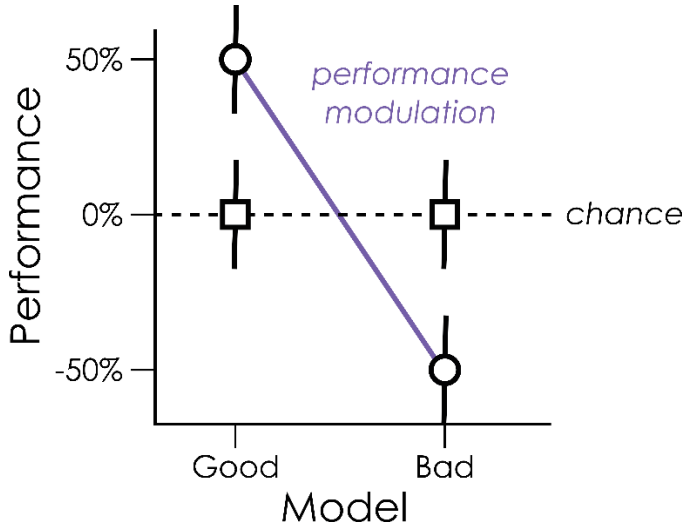

Figure 2-9. Illustration of agreement rate. The perception of new targets is predicted on the basis of previous detection (i.e. model). If the previous target was seen the model tells us perception will be "good", while if it wasn't it predicts "bad" perception. Without agreement within or between subjects, seeing the target once cannot be used to tell whether the subject will see the target again (i.e. squares on the figure). Our prediction ability is at chance. However, if the agreement is perfect, then, based on the perception of the target the first time (i.e. model), we can wholly predict the detection of the target at the second session (or of the other subject). In other words we have perfect classification. The performance modulation is the thus the difference in performance between the good and bad prediction.

To test whether the white noise sequences could still significantly modulate perception after the linear/direct influence of luminance sequences had been removed, we looked at how systematic the detection of a specific target was between or within subjects. In other words, we asked whether the same targets embedded within the same WN sequences were reliably seen or missed by a subject seeing the sequence twice. This is the *within subject* or *subject dependent* approach. We also asked whether part of the target's visibility could be explained by a *subject independent* modulation (i.e. do two subjects reliably miss and see the same targets?). Ultimately, the agreement rate within subject will be (at least in part) explained by the between subject agreement rate. Therefore, we will start by examining the between subject agreement rate.

2.8.1 Computing between subject agreement rate

The between subject agreement rate is represented on Figure 2-10. It was computed, for each target as the proportion of subjects (N=20) who agreed in its detection (i.e. either miss or see). The distribution across all targets presented during the second session (N=821 targets) was then computed (see Figure 2-10, red line). Since there seems to be a large number

of targets which are perceived by about 50% of subjects, we ask whether this distribution could be expected by chance. If our data could be explained by a random process, the distribution would be well approximated by a binomial distribution with a mean hit rate of 0.50. To test this, we fitted our observed distribution with simulated binomial distributions. The 100000 surrogates were computed to have the same number of targets (N=821) and mean hit rate (0.48 hit rate) as our data (see Figure 2-10A, black line). To test the statistical significance of this effect, we applied a chi squared test for the goodness of fit between the observed distribution and the mean of all surrogates (i.e. our expected distribution). We found that it is highly unlikely that a random process explains all the information in our dataset (X^2 (11, N = 821) = 218.7455, $p < .01^{-40}$, see Figure 2-10A).

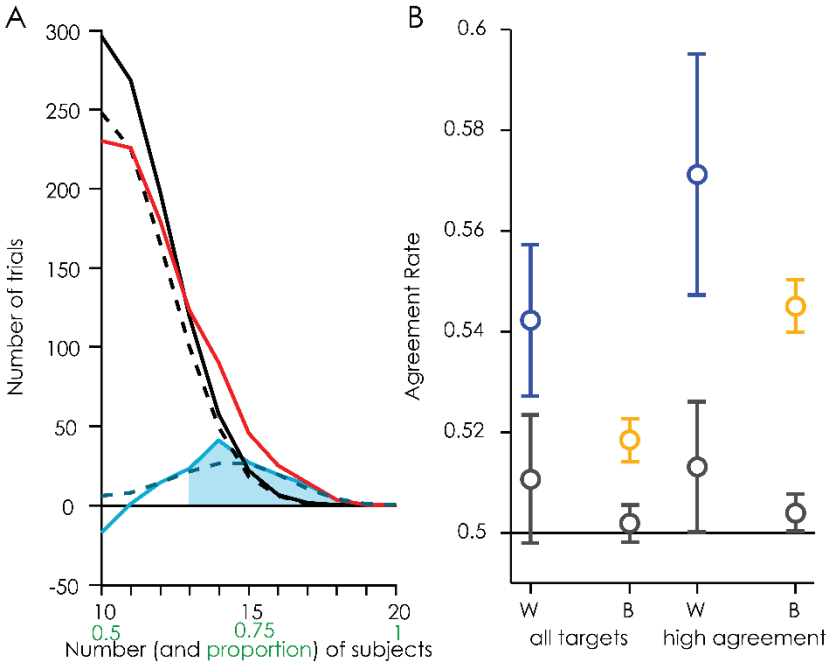

Figure 2-10. **A.** Distribution of targets depending on the number of subjects who saw them. The red line represents the observed distribution. The full black line represents the mean binomial distribution (100000 surrogate distributions) with the same number of targets with p=0.48 and N = 821. The dotted black line represents the fitted "random" process (binomial distribution with 689 targets and mean 0.48). The light blue line represents the sequence induced targets (132 targets i.e. difference between the random process and the observed data). The shaded area represents the high agreement targets (i.e. 13 subjects in agreement). The dotted blue line represents the best fit for the sequence induced process. Note that for representation purposes, the data was collapsed across agree seen (i.e. all subjects saw the target) and agree miss (i.e. all subjects missed the target), but the data was analysed as a full distribution. **B.** Agreement rate computed as the proportion of targets seen in the same way within subjects (in blue) and between subjects (in yellow) with the corresponding null hypothesis (in black). The circles represent the mean across subjects and the error bars

represent the standard error of the mean. The agreement rate was computed either on "all targets", or on a subset of targets (seen on shaded area in panel A)

Since a random process could not explain the agreement rate entirely, we hypothesised that our distribution could be modelled by two underlying processes. The first distribution representing "random targets" on which the constrained oscillations have no influence and from which we can never hope to extract any meaningful information about visual perception (i.e. the hit rate for these targets will be at 50%). The second distribution represents the "sequence driven targets" (i.e. targets whose detection is determined by the WN sequences to some extent). Any deviation in our observed distribution from the random distribution will likely reflect the subjects' propensity to detect the target in the same way. The first distribution of "random targets" was extracted by refitting the observed distribution with a binomial distribution to find the best fitting curve (peak to peak). We iteratively reduced the number of targets in the binomial distribution (mean hit rate of 0.48) and tested the fit of the curve using a least square fit. We found that the best fitting curve of the "random targets" distribution was modelled by a binomial with 689 targets (see Figure 2-10C, dotted black line). In other words, there is at most 84% of the targets we presented which are due to random fluctuations. This might be due to fluctuations in the attention of subjects or concentration. In any case, we cannot hope to make any predictions in terms of perception for these targets. Nevertheless, this also means that there are 132 out of 821 targets whose detection was determined by the entraining sequences (Figure 2-10A, green line).

Next we extracted the distribution of the sequence driven targets from the observed distribution by subtracting the random targets distribution. We can clearly see only figure 2-9 that this distribution seems to have a peak between 13 to 15 subjects (i.e. a proportion of 0.65 to 0.75 of subject in agreement). We wanted to quantify exactly what the mean hit rate of this distribution was (i.e. only considering the sequence driven targets). We fitted the sequence driven targets distribution by systematically increasing the mean hit rate of a (binomial) distribution from 0.48 from +/-0 to +/-0.3 in steps of 0.01. Here again, the goodness of fit was tested using the least square difference method. We found that the best fitting curve had a deviation 0.2066 (see Figure 2-10, dashes blue line). For these 132 targets, when the target is seen, the agreement rate across subjects will be of 0.48+0.2066 = 0.6866. On the other hand, when the target is not seen, the agreement rate across subjects will be of 0.48-0.2066 = 0.2734. In total, this makes a performance modulation of 0.6866 – 0.2734 = 0.2066 * 2 =

0.4132. In the best case scenario, if we were able to completely isolate these 132 sequence induced targets, we could expect a modulation of 41.32%. However, we don't know which targets these actually are, they cannot be easily isolated from the rest of the targets. Therefore, we start by computing the expected performance modulation when all targets are taken into account and assuming that the observed distribution is composed of 132 targets, which have a modulation of 41.32% and 689 targets which have a modulation 0%. Taken together, this gives us an expected performance modulation of 132*41.32 + 689*0 = 6.6434% across all targets.

As mentioned above, ultimately we want to construct classifiers which will allow us to predict target detection, and in chapter 6, we want to do so in a subject independent way: for any target presented, we want to know what the likelihood of perceiving it is, for anyone. Thus, we decided to make predictions only about targets which have a high agreement rate. Here, we simply subsample the targets to extract the performance modulation only for targets where at least 13 subjects agree in their percept (see Figure 2-10A, shaded area). There are 300 targets which have been seen by 65% of subjects, of which 132 targets are predicted to fall within the target induced sequences. On these targets (132 at 41.32% and 168 at 0%), we can expect to have a maximum performance modulation of 18.18%. Note that in chapters 6 and 7 the between subjects agreement rate will be computed on an independent group of subjects to avoid circularity.

2.8.2 Computing within subject agreement rate

We now turn to the within subject agreement rate. In chapter 7, we want to present a target within the white noise sequences so that one subject will see it while another subject will not. For this to work, we need the targets to be reliably detected within subjects (i.e. between sessions). For this, we tested the within subject agreement rate on 11 subjects who agreed to come back for a second psychophysical session. They saw the same WN sequences in a randomized order. We computed the within subject agreement rate as the proportion of targets perceived in the same way (either missed or seen) on the 2 different sessions. We compared these agreement rates to a "null hypothesis agreement rate" (or Anh), which reflects what we would expect based only on the similarity of hit rates between sessions (i.e. the sum of the product of hit rates and the product of miss rates). Ultimately the within subject agreement rate will, at least in part, be explained by the between subject agreement rate.

Therefore, we also computed the between subjects agreement rate using the second session's behavioural data, where all subjects had seen the same WN sequences (albeit in a pseudo-randomized order). We then systematically computed the between subject agreement rate as the agreement rate between each possible pair of (different) subjects. Significance was tested by applying a paired t-test between the agreement rate and the corresponding null hypothesis across subjects. Let us first start with the between subject agreement rate.

The between subjects agreement rate (see Figure 2-10B, blue) is 51.84% and the corresponding null-hypothesis agreement rate is 50.19%, yielding a difference of 1.65% (paired t-test of the difference with null hypothesis, $t(208) = 14.4847$, $p = 2.4*10^{-33}$, ci = 1.427 to 1.8767). This means that if we took the "seen" targets of session, we would expect the hit rate to be at 50-1.65 = 48.35% and vice versa for the "missed" targets, we would expect a hit rate of 51.65%. Therefore, this would give us an expected performance modulation (see Figure 2-9) of 3.3%.

Here again, we decided to subsample the targets to compute the agreement rate only for high agreement targets (i.e. shaded area on Figure 2-10A and high agreement targets in panel B), we expect the agreement rate between subject to increase. The agreement rate between subjects (on average) goes up to 54.46% with a null hypothesis is 50.3857%, yielding a difference between the two of 4.07% ($t(208) = 21.7339$, $p = 1.96*10^{55}$, 95% confidence interval of the difference is 3.71 to 4.45). Consequently, following the same logic as above, the performance modulation is of 8.14%. This increase in performance modulation for high agreement targets is expected as the two measures of similarity in perception presented above are not independent.

We now turn to the computation of the within subject agreement rate (see Figure 2-10B, yellow). In the same way as we did for the between subjects analysis, we computed the within subjects agreement rate on all targets. The agreement rate within subjects is 54.22% (null hypothesis is 51.07%), yielding a difference of 3.15% ($t(10)=6.811$, $p=4.6*10^{-5}$, 95% CI of the difference is from 2.12% to 4.18%). The expected performance modulation on the within subjects analysis of 6.3%. If we subsample the targets as we did for the between subjects analysis, our prediction ability (i.e. performance modulation) goes up to 57.06% compared to a null hypothesis of 51.29%, yielding a difference of 5.77 % ($t(10) = 7.317$, $p=2.54 *10^{-5}$, 95% CI of the difference is from 4.01% to 7.52%). In total, we expect a modulation of 11.54%. This increase is non -trivial because the targets were defined independently of the within subject agreement rate.

2.8.3 Discussion

We found the WN sequences could reliably influence the detection of the embedded targets, with two caveats. First, there only seem to be a small proportion of targets whose perception appears to be modulated by the sequence. This might be due to the difficulty of the task or to fluctuations in the attention or alertness of subjects. Moreover, a large amount of the effect of WN sequences on performance was mediated by a between subject agreement rate. In other words, the subjects were likely to miss or see the same targets. There was still, however, a significant part of behaviour that could be explained by a within subject agreement rate, beyond the shared effects across subjects. These results confirm that attempting to build forward models of perception on subject-independent and subject-specific features of the reconstructed EEG is a feasible enterprise. However, it does put an upper boundary on what amount of performance modulation we can expect from using this task, we cannot hope to modulate perception completely, but it will serve as a proof of concept (this will be discussed further in chapters 6 and 7).

2.9 Conclusions

In a nutshell, we have found that the IRF is a good (thought far from perfect) model of the brain activity. Morevoer, it is subject specific and stable in time. Further, behaviour is modulated by the WN sequences (in a non-trivial, non-linear manner). Now that the framework has been validated, we can answer the important questions: where do these background oscillations come from in the brain? Are they related to perception? And if so, how much of the variability in performance can they explain compared to the recorded, ongoing oscillations? Can they enlighten us about the link between oscillations and visual perception, so that we can control the perception of subjects?

Chapter 3. NEURONAL BASIS OF PERCEPTUAL ECHOES

In the previous chapter, we validated the White Noise Paradigm as tool to study the link between ongoing oscillations and perception. In this chapter, we will investigate what the neuronal basis of the perceptual echoes might be. This project was done in collaboration with Leila Reddy and Isabelle Berry. Before evaluating how the constrained oscillations might influence the detection of near-perceptual targets, we first wanted to find what brain areas gave rise to the perceptual echoes to better understand the nature of the reconstructed EEG.

3.1 ARTICLE 1: NEURONAL BASIS OF PERCEPTUAL ECHOES (IN PREPARATION)

List of all Author Names and Affiliations

Sasskia Brüers[1,2], Leila Reddy[1,2], Isabelle Berry[1,3], and Rufin VanRullen[1,2]
[1]Faculté de Médecine, Université de Toulouse Paul Sabatier, Toulouse, France
[2]Centre de Recherche Cerveau et Cognition, CNRS, UMR 5549, Toulouse, France
3Toulouse NeuroImaging Center, INSERM, U825, Toulouse, France

Author Contributions

RV and SB Designed research, IB and SB performed research, LR and SB analysed data; SB wrote the first draft

Correspondence should be addressed to (include email address)

Rufin VanRullen, Centre de Recherche Cerveau et Cognition (CerCo), CNRS, UMR 5549, Pavillon Baudot, CHU Purpan, BP 25202, 31052 Toulouse Cedex 03, (rufin.vanrullen@cnrs.fr)

Acknowledgements
The authors would like to thank the technical team at the neuroimaging centre for their help and support.

Conflict of Interest
Authors report no conflict of interest

Funding sources
This research was supported by the ERC grant P-CYCLES (N°614244) to RV.

3.1.1 Introduction

Although a lot is known about the temporal profile of the impulse responses to white noise random luminance sequences, little is known about what their neuronal basis might be.

Many studies have evaluated the exact temporal profile of the EEG impulse response using random luminance sequences: it is characterized by an alternation of positive and negative potentials every 100ms or so (Lalor, Pearlmutter, et al., 2006; VanRullen & Macdonald, 2012). Moreover, this response can be separated into two parts. The "VESPA" (visual evoked spread spectrum potential activity, Lalor, Reilly, Member, Pearlmutter, & Foxe, 2006) corresponds to the "early" part of the IRF. In terms of temporality, the VESPA are correlated to the visual evoked potential (VEP), more specifically, there is a strong correlation between the topographies of the first early component of the VESPA (C1) and that of the VEP (Murphy et al., 2012). The "late" part of the IRF is composed of a long lasting, subject-specific, ringing in the alpha frequency band (VanRullen & Macdonald, 2012), as if the luminance information presented was echoed in the brain. Both responses have their peak amplitude over the occipital channels. However, little is known about the underlying neuronal sources of these different components.

Some studies have evaluated the brain regions which might generate the VESPA. Because there is a strong correlation between the amplitude of the C1 and P1 components of the VESPA with the VEP's C1, they could be the result of shared widely spread generators in the early visual cortex (Murphy et al., 2012). In fact, this was supported by evidence from a retinotopic mapping of the VESPA response using high density EEG: applying source localization algorithms revealed generators of the C1 and P1 components of the VESPA in the calcarine cortex (Lalor et al., 2012). More generally, the brain basis of the VEP are thought to arise from early visual areas: the activity underlying the C1 component arises in area V1 while a later component such as P1 has its basis in V3 and middle occipital gyrus (Di Russo et al., 2005, 2002).

In terms of frequency, the "late" part of the IRF (corresponding to the perceptual echoes) is characterized by a strong oscillatory component within the alpha band (~10 Hz), which is correlated in frequency with that of the subject's resting state alpha (VanRullen & Macdonald, 2012). The source of alpha oscillations has been a matter of debate. Alpha oscillations in the behaving monkey have been shown to have a cortical generator in both V2, V4 and IT (Bollimunta et al., 2008) although thalamo-cortical networks have also been

implicated (Lopes da Silva, 1991; Steriade et al., 1990). It is likely to be the product of both thalamo-cortical and cortico-cortical networks. A few studies looked into the relationship between the trial by trial correlation between the alpha power and the blood-oxygen level dependent (BOLD) activity, early studies focusing on the alpha power at rest. There is a negative correlation between the frontal and parietal activity and the alpha power, although there were a few voxels which were found to be positively correlated with alpha power in the rest condition over the occipital lobe (Laufs et al., 2003; Moosmann et al., 2003). Later studies also investigated the correlation of alpha activity as a function of visual tasks. Overall, they found that the ongoing alpha amplitude during visual tasks was negatively correlated to BOLD signal over the visual cortex (Moosmann et al., 2003; Scheeringa et al., 2009; Scheeringa, Koopmans, van Mourik, Jensen, & Norris, 2016; Zumer, Scheeringa, Schoffelen, Norris, & Jensen, 2014). These studies support the inhibitory role of alpha activity: it serves as a mechanism through which information is routed in the brain and reflects the top-down influence of attentional control (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007; Mathewson et al., 2011). Could it be that the same mechanism underlies the generation of perceptual echoes? Although in the same frequency band and largely correlated with the alpha peak frequency, the perceptual echoes show some functional differences. For example, attention to a white noise sequence presented to one hemi- field has been shown to increase alpha amplitude in the perceptual echo recorded over the contralateral channels (VanRullen & Macdonald, 2012). This effect is however inverse for the ongoing alpha oscillations, which decrease over contralateral channels as a function of attention allocation (Rihs, Michel, & Thut, 2009; Romei et al., 2010; Sauseng et al., 2005; Thut et al., 2006).

In this study, we investigated the neuronal basis of the perceptual echoes. Using the WN paradigm (over recording the EEG-fMRI concurrently), was advantageous for two reasons. First, while recording EEG-fMRI at the same time takes advantage of both the good temporal resolution of EEG and the good spatial resolution of fMRI, it has the drawback of creating artefacts in both the EEG and the BOLD activity (Huster, Debener, Eichele, & Herrmann, 2012). Furthermore, it has been shown that spurious correlations between EEG and BOLD activity could be created from movement in the scanner (Fellner et al., 2016). Thus, having the opportunity to record the EEG outside of scanner and reconstruct the EEG signal rather than record it is a strength of the White Noise paradigm. Secondly, using the white noise paradigm allows us to isolate the brain activity specifically constrained by the

sequences for that particular subject. Since the IRF's exact frequency is subject specific, the brain activity elicited by the stimulation will also be subject specific. Thus reconstructing, rather than recording the EEG activity can extract the brain activity for a given subject which reflects the time course of the perceptual echo (i.e. the phase locked activity) as the stimulus unfolds. However, if we recorded the EEG, we would also extract information about the time course of the induced signal or spontaneous oscillations potentially present in the EEG, which might also correlate with the BOLD signal.

First, the impulse response functions (IRFs) were recorded outside of the scanner. They were then used to reconstruct the brain activity to the white noise sequences presented inside the MRI. In particular, we were interested to find which brain regions had BOLD activity which co-varied with the brain oscillations elicited by the WN sequences. Since we know that the perceptual echoes are composed of a strong oscillatory component in the alpha band, we chose to focus on the alpha activity specifically: for each subject, the reconstructed EEG alpha envelope of each trial was extracted and used as a regressor in the GLM for the BOLD activity.

3.1.2 Materials and Methods

3.1.2.1 Participants

In total, 22 participants were included in the study after a medical interview, and giving written informed consent. In total 18 subjects completed the experiment. Two subjects had to be excluded after the first recording session because their alpha level in the echo did not reach the pre-determined threshold (see criteria below in EEG recording). This study was approved by the ethics "Comité de Protection des Personnes Sud-Méditerranée I" (N°2016-A01937-44).

3.1.2.2 Experimental Protocol

The experimental protocol consisted of 2 sessions: the recording of impulse responses using electro-encephalography (EEG) during a first session and the acquisition of anatomical and functional activity using magnetic resonance imaging (MRI) during the second session. Each session lasted about 1h45, including rests and the instruction period, and approximately 1h15 minutes recording per session.

The EEG session was composed of 8 runs of 48 WN sequences each lasting 6.25s. The experiment was self-paced to allow participants to blink as needed: they pressed a button to start blocks and trials. The fMRI session was composed of 12 functional runs. Each experimental run consisted of 7 WN sequence stimulation lasting 30 seconds each with an ITI of 12 seconds. An anatomical scan was recorded at the beginning of the functional mapping.

3.1.2.3 Stimuli

In both sessions, white noise (random luminance) sequences were presented in a disk (diameter of 2° of visual angle) in the upper left quadrant. The centre of the disk was at 5° of eccentricity from the fixation point (0.1° of visual angle). In both sessions, the task was to detect, as quickly as possible, near perceptual threshold targets embedded within the white noise sequences. The targets were composed of a lighter disk surrounded by a darker annulus presented on a medium grey background. They had a mean luminance of "medium grey", identical to the mean "medium grey" level of the white-noise sequence. The targets were presented for 1 frame only. Using a staircase procedure on the first 100 targets presented (i.e. about 30 trials), we manipulated the visibility of targets by changing the contrast between the outer (darker annulus) and inner (lighter disk) parts to reach the contrast at which participants perceived about 50% keeping the resulting contrast constant for the reminder of the session. The perceptual threshold was computed for each session independently using the *quest* function (Watson & Pelli, 1983).

Because of differences in setup, the stimuli presented in the EEG and fMRI recording session were not identical in terms of temporal frequency. In session 1, the CRT monitor on which the WN sequences were presented had a resolution of 640 by 480 pixels and a refreshrate of 160 Hz, giving the white noise sequences a flat power spectrum between 0 and 80Hz. During session 2, the stimuli was presented on a 24 inches fMRI compatible LED screen from Cambridge Research Systems, with a maximum presentation rate of 60 Hz. This means that the white noise sequences only had a flat power spectrum between 0 and 30 Hz, yet well above the temporal resolution of the BOLD activity.

3.1.2.4 EEG recording, pre-processing, extraction of IRF and regressors

During the first session, the brain activity was recorded using a 64 channels BioSemi EEG (1024 Hz sampling rate), with 4 external ocular electrodes recording the horizontal and vertical oculogram. Offline, the data were pre-processed using the EEGlab plugin (Delorme & Makeig, 2004) and custom written code in MATLAB. The following pre-processing steps were applied to the EEG data: 1) rejection and (linear) interpolation of noisy and/or artifacted channels (as needed) , 2) down-sampling to 160 Hz to facilitate the cross-correlation with the stimuli, 3) notch filter (47 to 53 Hz) to remove artefacts from the power line, 4) average re-reference, 5) creating data epochs (384) around the white noise sequences, from -250ms to 6.5s around the sequences, 6) baseline correction (from – 250ms to 0ms), 7) rejection of whole epochs (as necessary) containing ocular and movement related artifacts. The IRF were extracted by doing a cross-correlation between the pre-processed EEG data and the corresponding WN sequences (VanRullen & Macdonald, 2012).

From previous studies, we saw that the alpha power in the echo was correlated with the accuracy of the reconstructed EEG as a model of the recorded EEG. Therefore, based on these 20 subjects, we set a minimum alpha power threshold that the echoes of the subjects had to achieve to be included in session 2. At this stage, 2 subjects were removed because their echoes did not have enough alpha power.

The IRF were then used to reconstruct the brain activity to the second session's WN sequences. First, the IRF were down-sampled to 60 Hz (to match the presentation rate of the WN sequences in the scanner). The reconstructed EEG was extracted by convolution of the IRF and WN sequences. Because a lot of information was shared between channels in terms of envelope, we decided to only use one electrode for each subject. We chose the electrode which had the most power in the alpha band in the "late" part of the echo (from 250ms to 1250ms).

This reconstructed EEG signal was then used as a regressor of the BOLD activity. Because the perceptual echoes are characterized by a frequency specific resonance in the alpha band, the reconstructed EEG was filtered (using the *filtfilt* function and a FIR filter) in 4 frequency bands (delta: 2-4 Hz, theta: 4-8 Hz, alpha: 7-14 Hz, beta: 14-28Hz). The absolute value of the Hilbert transform data was taken as the envelope. Finally, the regressors were 'clipped' by removing 4s at the beginning and 2s at the end of each block to remove the

systematic filtering artefacts present (across all trials, subjects and frequency), which could have led to spurious correlations.

3.1.2.5 fMRI recording and pre-processing

The MRI session data were acquired using a 3T Philips (Amsterdam, The Netherlands) ACHIEVA scanner with a 32-channel head coil. First we recorded anatomical images from each participant with a high resolution: 170 sagittal slices were acquired with a voxel size of $1mm^3$, a repetition time (TR) of 8.13ms, and a time to echo (TE) of 3.74ms. The functional images were acquired using a gradient-echo pulse sequence with a TR of 2 seconds and a TE of 35ms of 39 slices positioned to cover the whole brain. The voxel size was of $3mm^3$. The analysis of the functional images was as followed: 1) all images were motion-corrected, 2) slice time-corrected, 3) intensity-normalized, and 4) smoothed with a 3mm Gaussian kernel. Then, we extracted regions of interest (ROI) based on the regions which were responsive to the stimulation. To extract them, we contrasted the maps in the right hemisphere for the stimulus "on" versus stimulus "off" time points across all runs and 16 subjects using a generalized linear model (GLM) in the Freesurfer average brain space. A group analysis was then performed to define significant clusters which reached a voxel-wise $p<10^{-4}$ and cluster wise p 10^{-1}. The activity was then overlaid on an anatomically defined (on the average brain) V1 and V2. A GLM was applied in each subject to test the correlation between the BOLD activity and the reconstructed EEG alpha power in the ROIs and in anatomically defined brain regions V1 and V2. We also tested whether the correlation between alpha activity and BOLD was specific to the alpha frequency by including the delta, theta and beta frequency bands as regressors in a separate GLM. The statistical significance of the effect at the group level was tested using a t-test of the coefficient for the single subjects against 0. An fdr correction was applied to correct for multiple comparisons.

3.1.3 Results

Using the white noise paradigm, we investigated the neuronal basis of the perceptual echoes (VanRullen & Macdonald, 2012). We took advantage of the EEG reconstruction method to extract the subject specific time course of this alpha reverberation in response to the WN sequences presented inside the scanner. This reconstructed EEG envelope was then used as a regressor of the fMRI BOLD activity to extract the brain regions which co-varied

with the fluctuations in alpha power in response to WN sequences. Using the reconstructed rather than the recorded EEG provided a direct access to the time course of the activity of the "perceptual echo" rather than all the other activities typically recorded in the EEG that co-vary with it. Because the perceptual echoes are characterized by a peak in the alpha band, we first looked at whether this alpha activity was driving the BOLD in any of the regions of interest defined from the retinotopic mapping.

First, we extracted the retinotopic brains regions activated by the stimulus in the upper left quadrant. For this, we did a contrast between times were the stimulus was "on" versus "off" (i.e. inter trial) intervals. An automatic cluster extraction revealed the presence of two clusters of activity in the right hemisphere: cluster one (C1, see Figure 3-1, left) was located in the lingual gyrus below the calcarine fissure ($p<0.05$ corrected, MNI coordinates: x = 21.9, y = -79, z = -8.3). The second cluster (C2, see Figure 3-1, right) of activity was found in the lateral occipital gyrus ($p<0.05$ corrected, MNI coordinates: x = 41.3, y = -72.6, z = -0.5). There was no corresponding cluster in the left hemisphere (since the stimulus was lateralized to the left).

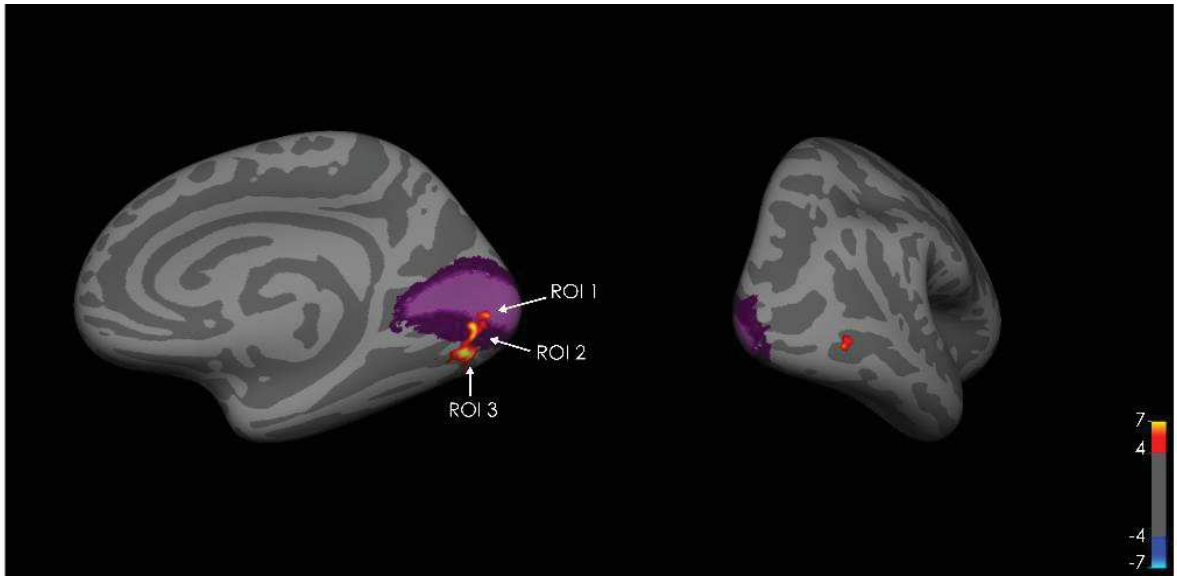

Figure 3-1. Retinotopic activation and early visual areas. Representation of the regions retinotopically activated by the presentation of the target is in red/yellow. These were extracted from the contrast between the stimuli "on" versus "off" intervals, and overlaid on a flattened cortical surface. The light purple represents the anatomical location (extracted from the average brain) of the primary visual area (V1) and the darker purple represents the location of the secondary visual area V2. Statistical analysis revealed 2 clusters of activity: cluster one (C1) on the left is located below the calcarine sulcus (MNI coordinates: x = 21.9, Y=-79, Z=-8.3) and cluster 2 can be found in the second representation, possibly corresponding to lateral occipital cortex (LOC).

From the two clusters of activity extracted, we created regions of interest (ROIs): the first two clusters already correspond to two different ROIs: C1 and C2 respectively. We also wanted to see if there was a specific involvement of V1, V2 and thus we created 3 further ROIs within the first cluster. R1 was created as the intersection of cluster 1 with V1, R2 as the intersection of cluster 1 with V2 and R3 was the rest of the activity in cluster 1 that neither overlapped with V1 nor V2. C1 was thus the union of R1, R2 and R3. A GLM was then applied to these 5 regions of interest: C1, C2, R1, R2, R3 (see Figure 3-2).

First, in the ROIs extracted from the brain regions responsive to the presence of the stimulus, we found that the BOLD activity was not significantly modulated by the envelope of the alpha band reconstructed EEG (see Figure 3-2A). In ROI 1 to 3 or in cluster C1, the mean coefficient across the group was close to 0 and characterized by a large inter-subject variability (ROI 1 $p = 0.93$, ROI 2 $p = 0.88$, ROI 3 $p = 0.69$, C1 $p = 0.78$). This is somewhat reduced in C2, but the mean coefficient is still not significantly different from 0 (C2 $p = 0.13$).

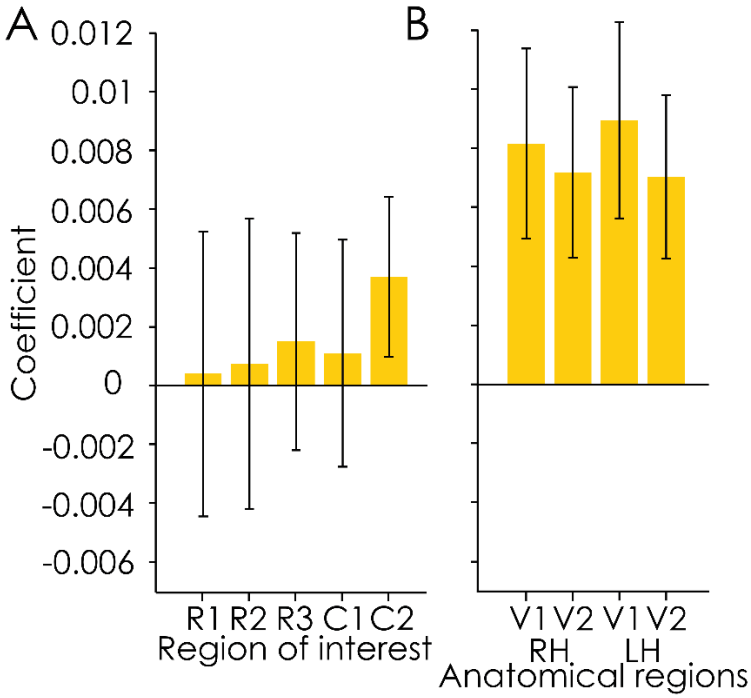

Figure 3-2. Coefficients of the GLM looking at the alpha frequency band envelope as a regressor of the BOLD activity (A) in the five regions of interested extracted from the retinotopic mapping of the stimulus responsive brain regions (see Figure 3-1) and (B) to anatomically defined early visual areas, namely left and right V1 and V2. The bars represent the means, the error bars represent the standard error of the mean (SEM)

Because we found no effect in the retinotopic regions, we decided to look at the whole brain activity. Could we see some other brain regions in the visual cortex whose activity was correlated with the time course of the reconstructed alpha activity? In fact, this analysis reveals that there is BOLD activity that co-varies with the alpha power: we see widespread activations in both V1 and V2 on both sides (see Figure 3-3). This activity is not restricted at all to the ROI, but extends well outside of these regions. In fact, there seem to be absolutely no activity in the ROIs (as confirmed by analysis of the mean coefficients across subjects, see Figure 3-2).

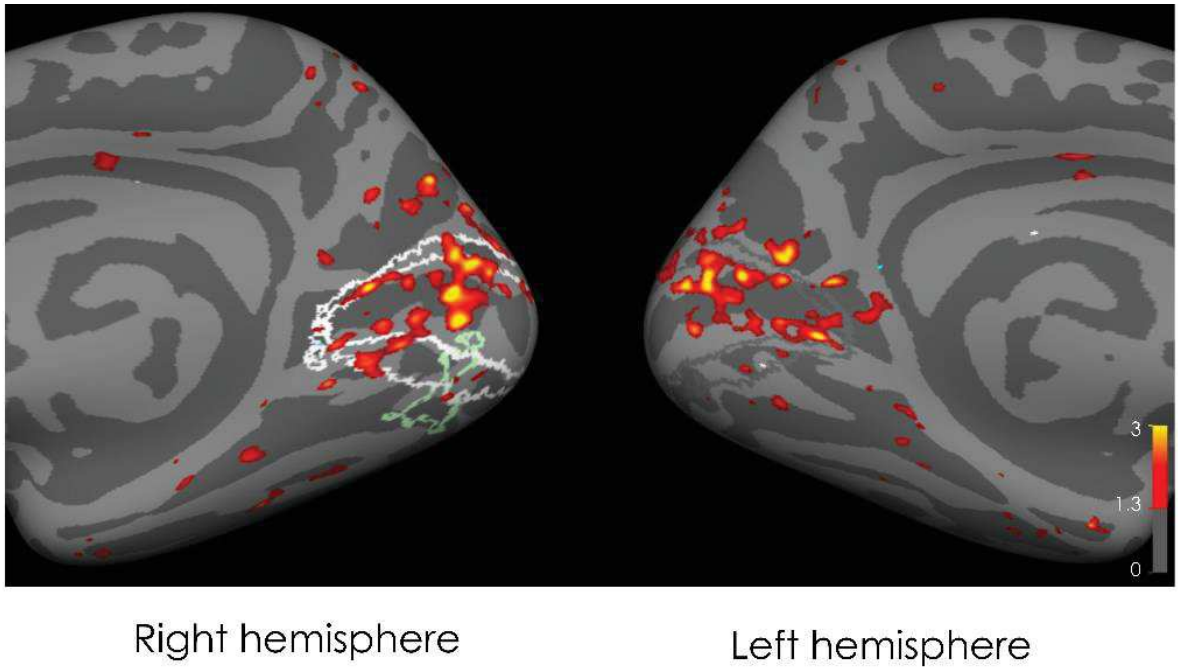

Figure 3-3. Whole brain analysis of regions which are correlated to fluctuations in the alpha power envelop in the right and the left hemisphere. The activations extend well beyond the regions of interests based on the stimulus: there is wide spread activity in both V1 and V2 bilaterally.

We also looked whether there were any statistically significant correlation between the BOLD activity and the reconstructed EEG alpha power in anatomically defined early visual areas V1 and V2 (see Figure 3-2B). We found that there was a significant correlation in the right hemisphere in V1 ($t(15) = 2.5241$, $p = 0.0233$, mean = 0.008, 95% CI = 0.001 – 0.01), and V2 ($t(15 = 2.4849$, $p = 0.024$, mean = 0.007, 95% CI = 0.001-0.013). Surprisingly the

effects were also significant in the left hemisphere in V1 (t(15) = 2.6979, p = 0.0165, mean = 0.008, 95% CI = 0.001 – 0.016), and V2 (t(15 = 2.4849, p = 0.024, mean = 0.007, 95% CI = 0.001-0.012).

Finally, we asked whether the information in the alpha band was specifically modulating the activity in these early brain regions, or whether the activity in other frequency bands could also drive the BOLD activity (see Figure 3-4).

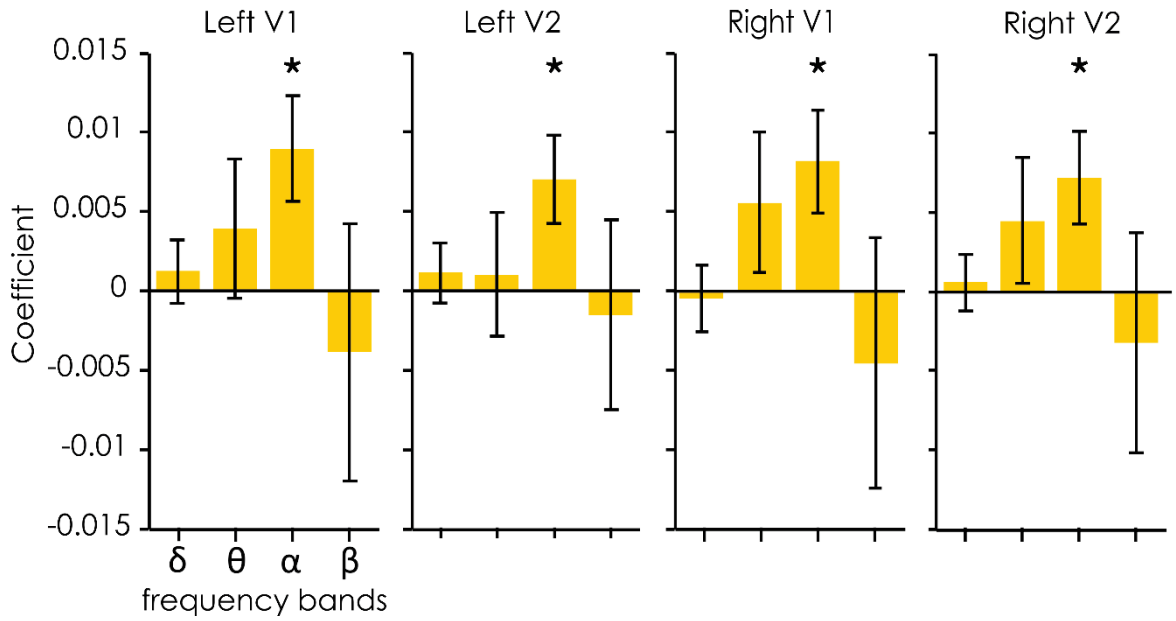

Figure 3-4. Coefficient of GLM for the regressors of the BOLD activity based on the envelope of the filtered reconstructed EEG in the delta (δ), theta (θ), alpha (α) and beta (β) frequency bands for the left and right V1 and V2. The bars represent the mean and the error bars represent the standard error of the mean.

A GLM was constructed with the envelopes of the reconstructed EEG in the delta, theta, alpha and beta band as regressors for the time course of the BOLD activity. We found that the alpha regressors were specifically modulating the BOLD activity, but not the other frequency bands.

3.1.4 Discussion

In summary, we find that the alpha band activity representing the resonance of the brain in response to WN sequences is positively correlated with the activation of both the primary and secondary visual cortex areas in both hemispheres. As the alpha power increases, so does the BOLD activity in these brain regions. The effects however, are not driven by any activity in the brain regions which are otherwise responsive to the presence of the stimulation.

Moreover, the effects are specific to the alpha band envelope, the BOLD activity in V1 and V2 did not correlate with the time course of the delta, theta and beta activity.

First, the fact that we do not find any correlation between the BOLD activity and the alpha envelope in retinotopically specific locations, while there is a correlation with the wider V1 and V2 brain regions is puzzling. This effect might be due to a saturation of the BOLD activity in those retinotopic regions due to the high luminance fluctuation presented. Once at its peak due to the strong sensory stimulation, the BOLD activity might not be sensitive at all to the relatively small changes induced by oscillatory fluctuations during the session, which would explain the absence of response in these regions. The rest of V1 and V2, on the other hand, does not receive direct retinotopic visual stimulation, and remains far from a potential saturation point.

Second, we found that the BOLD activity in the whole of V1 and V2 in both hemisphere co-varies with the alpha envelope. We cannot completely discard the hypothesis that part of this activation (especially in the left hemisphere) could be driven by an indirect stimulation of the right visual field by the stimulus. In fact, the bore of the scanner is white and the luminance from the stimulus could have reflected off the sides. This seems an unlikely explanation as the size of the stimulus was small (2° of visual angle). Another alternative hypothesis is that these wide spread activation could be the result of the activation of a travelling alpha wave by the long-lasting stimulation. Research in our lab has shown that the perceptual echoes created travelling alpha waves in the visual cortex (Lozano-Soldevilla & Vanrullen, 2016). Furthermore, studies have shown that the ERP components (such as the P1) were actually the result of underlying alpha travelling waves going from the occipital to the parietal cortex (Klimesch, Hanslmayr, Sauseng, Gruber, & Doppelmayr, 2007). A travelling alpha wave generated through the stimulation sequence could explain why we find large-scale activations in both V1 and V2, in both hemispheres.

This positive correlation between reconstructed EEG alpha power and BOLD activity is different than what was expected from the literature. As mentioned in the introduction of the present chapter, previous studies have reported a negative correlation between the BOLD signal and alpha power over the visual cortex (Laufs et al., 2003; Moosmann et al., 2003; Scheeringa et al., 2009, 2016; Zumer et al., 2014), although there were a few voxels which were found to be positively correlated with alpha power in the rest condition over the occipital lobe (Laufs et al., 2003). This further supports the idea that the perceptual echoes represent a mechanism which has a different functional role from the alpha oscillations acting as a top-

down inhibitory mechanism (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007; Mathewson et al., 2011).

In conclusion, we find that the BOLD activity evoked by the white noise sequences is significantly correlated to the reconstructed alpha power. This alpha power is likely to reflect activity in the "late" part of the impulse response functions (i.e. the perceptual echoes) which are characterized by a long-lasting ringing at about 10 Hz. Thus, our result suggest that the perceptual echoes originate in early visual cortex. Specifically, it is driven by a wide spread activity in V1 and V2, which is not retinotopically restricted, but possibly reflects the activity of a travelling alpha wave.

3.2 Conclusions

In this chapter, we have shown that the alpha power was positively correlated with the activity in V1 and V2 area. Although not conclusive (correlation does not show causation) this suggests that the perceptual echoes, just as the VESPA, have their source in the early brain areas.

Furthermore, this positive relation between alpha power and BOLD activity supports the notion that the perceptual echoes, although in the general alpha frequency band, reflect a different mechanism than the traditional endogenous alpha oscillations. Under the taxonomy introduced in the chapter 1, the endogenous alpha oscillations would (when driven by an endogenous factor such as attention) fall within the "relevant spontaneous oscillations" (signal S2). Our results thus seem to validate the functional distinction we made in chapter one between these S2 signals and the "phase locked oscillations" (signals 3 and 4, see taxonomy of rhythms proposed in chapter 1) as their power time course has an opposite relation with the BOLD activity.

That being said, there are limitations to this study. One potential limitation is that we did not record the echoes directly in the scanner. Although the experimental setup of the two sessions were made to be as similar as possible, there were still differences. For example, the ambient luminance in the scanner was higher than in the EEG room. The screen used for the fMRI experiment had a strong luminance, which reflected off the walls of the fMRI. We did decrease the luminance of the stimulus as much as possible, however. Moreover the sampling rate of the stimulation was different in both cases. Thus, it could be that despite our best efforts, the IRF recorded inside the scanner would have looked differently than those acquired

outside the scanner. One potentially easy way to circumvent this is by recording the IRF inside of the scanner (scanner off), to confirm that the echoes would be the same. By recording the brain activity inside the scanner, we could also extract information about the time course of the endogenous/spontaneous alpha activity during the white noise sequences. By using both the power of the recorded EEG as well as the reconstructed EEG, we could then directly compare the time course of these two activities. This could highlight differences between endogenous driven alpha oscillations and the sequence drive alpha oscillations. Finally, the specificity of the EEG regressor might be further refined. Recently, a paper in our lab showed that the perceptual echoes could be modelled by two underlying sources of activity: a parietal and an occipital alpha source, which showed slight within subject differences in their alpha peak frequencies (Gulbinaite et al., 2017). It might be interesting to see whether these effects could be replicated at the level of the reconstructed EEG inside the scanner.

Chapter 4. BACKGROUND OSCILLATORY PHASE AND VISUAL DETECTION

4.1 INTRODUCTION

In this chapter, we aim to answer three questions. First, we wanted to explain the pre-stimulus phase effects reported in the literature. If the phase of ongoing oscillations causally determines whether a target will be seen or not, then phase effects should be found in the post-stimulus time window too, while the stimuli are being processed in the brain. However, most papers in the literature report effects exclusively in the pre-stimulus time window, peaking up to 100ms before the target is even presented. We hypothesize that this effect is due to presence of large amplitude evoked responses (i.e. signal S2 in the general framework, chapter 2).

In parallel this chapter gives the ground work for further chapters (6 and 7). The analysis of the reconstructed EEG also allow us to test whether the phase of the reconstructed EEG is related to the perception of the targets embedded in the sequences, which is crucial for analyses presented in these chapters.

Finally, this also gives us the opportunity to disentangle the relative contribution of signals S2 and S4 to the phase effects in the EEG. There might be a possibility that only the phase of relevant spontaneous signal (S2) influences perception: this would be the case if the effects was wholly driven by a top-down attentional effect for example. In this case, we should not see any phase effect using the white noise paradigm, since the reconstructed EEG only models the time course of signal S4.

4.2 ARTICLE 2: AT WHAT LATENCY DOES THE PHASE OF BRAIN OSCILLATIONS INFLUENCE PERCEPTION? (*PUBLISHED IN ENEURO*)

eNeuro

New Research

Cognition and Behavior

At What Latency Does the Phase of Brain Oscillations Influence Perception?

Sasskia Brüers,[1,2] and Rufin VanRullen[1,2]

DOI:http://dx.doi.org/10.1523/ENEURO.0078-17.2017

[1]Université de Toulouse Paul Sabatier, 31062 Toulouse cedex 9, France and [2]Centre de Recherche Cerveau et Cognition, CNRS, UMR 5549, BP 25202, 31052 Toulouse Cedex Toulouse, France

Abstract

Recent evidence has shown a rhythmic modulation of perception: prestimulus ongoing electroencephalography (EEG) phase in the θ (4–8 Hz) and α (8–13 Hz) bands has been directly linked with fluctuations in target detection. In fact, the ongoing EEG phase directly reflects cortical excitability: it acts as a gating mechanism for information flow at the neuronal level. Consequently, the key phase modulating perception should be the one present in the brain when the stimulus is actually being processed. Most previous studies, however, reported phase modulation peaking 100 ms or more before target onset. To explain this discrepancy, we first use simulations showing that contamination of spontaneous oscillatory signals by target-evoked ERP and signal filtering (e.g., wavelet) can result in an apparent shift of the peak phase modulation towards earlier latencies, potentially reaching the prestimulus period. We then present a paradigm based on linear systems analysis which can uncover the true latency at which ongoing EEG phase influences perception. After measuring the impulse response function, we use it to reconstruct (rather than record) the brain activity of human observers during white noise sequences. We can then present targets in those sequences, and reliably estimate EEG phase around these targets without any influence of the target-evoked response. We find that in these reconstructed signals, the important phase for perception is that of fronto-occipital ~6 Hz background oscillations at about 75 ms after target onset. These results confirm the causal influence of phase on perception at the time the stimulus is effectively processed in the brain.

Key words: EEG; oscillation; phase modulation

Significance Statement

When investigating the relationship between ongoing electroencephalography (EEG) oscillations and perception in humans, most studies report a peak influence of the EEG phase before stimulus onset. However, we should also be able to measure these effects poststimulus, when the target is actually processed by the brain. First, we use simulations to show that a combined influence of the target-evoked potential and filtering can explain the lack of poststimulus phase modulation in typical studies. Crucially, we then introduce a paradigm to uncover the true latency at which phase influences perception. The white noise paradigm allows us to model background oscillations without target-evoked potentials. For the first time, we show that a θ-band ongoing oscillation influences perception ~75 ms after target onset.

Introduction

Instead of continuously processing the incoming visual information, the "perceptual cycles" hypothesis suggests that our brain relies on a rhythmic sampling of the input. Ongoing oscillations have been proposed as a mechanism through which our environment is sampled from

Received March 6, 2017; accepted May 15, 2017; First published May 19, 2017.
The authors declare no competing financial interests.

Author contributions: R.V. designed research; S.B. performed research; R.V. and S.B. analyzed data; R.V. and S.B. wrote the paper.
This work was supported by the ERC Grant P-CYCLES 614244 (to R.V.).

5-15 times a second (VanRullen, 2016b; VanRullen and Koch, 2003). At the neuronal level, this is realized through a biasing of the neuronal firing by local field potential (LFP) phase in various frequency bands (Fries et al., 2002; Jacobs et al., 2007). This results in periodic fluctuations of the excitability of the cortex (Fries, 2005; Lakatos et al., 2008): spikes in sensory cortices are more likely to occur at a given phase of LFP α oscillations than at the opposite phase (Haegens et al., 2011; Haegens et al., 2015).

At the global level, this relationship between excitability and phase has been studied using transcranial magnetic stimulation (TMS) and electroencephalography (EEG): an illusory percept (a "phosphene") was most likely to occur when the TMS pulse was applied at a certain phase of the ongoing EEG oscillation than at the opposite phase (Dugué et al., 2011). Moreover, ongoing EEG oscillations have also been linked with visual perception: for example, the ~7 Hz ongoing oscillation phase over fronto-central channels could account for 16% of the variability in the detection of near-perceptual threshold peripheral targets (Busch et al., 2009). Various studies have related ongoing EEG phase to behavioral and perceptual outcome using near-perceptual threshold target detection with attentional manipulation (Busch and VanRullen, 2010) or without it (Nunn and Osselton, 1974), using contour integration tasks at perceptual threshold (Hanslmayr et al., 2013), using suprathreshold stimuli detection (Callaway and Yeager, 1960; Mathewson et al., 2009), using eye-movement initiation (Drewes and VanRullen, 2011) and mislocalization (McLelland et al., 2016) or using a temporal illusion (Chakravarthi and VanRullen, 2012). Crucially, most of these experiments find the effects of the ongoing θ and/or α-phase during the prestimulus period, sometimes peaking even up to 200 ms before stimulus onset (for a review, see VanRullen, 2016b).

Superficially, this prestimulus effect might seem counter intuitive: the critical phase for perception should be the one present in the cortex during stimulus processing. If oscillatory signals are consistent over time, this phase influence may of course be visible several hundred of milliseconds before, but why does it vanish as stimulus onset approaches? The event-related potential (ERP) evoked after target presentation could be causing this seemingly contradictory finding: this relatively high-amplitude signal with similar phase values on every trial is likely to obscure any difference in background oscillatory phase of perceived and unperceived trials that would be present after stimulus onset, and thus apparently "push back" in time our ability to detect any such phase difference. This is especially true when considering window-based time-frequency analysis methods (a necessary step in oscillatory phase analysis), which will smear the effect in time (Lakatos et al., 2005; VanRullen, 2011, 2016a; Hanslmayr et al., 2013). Can we overcome these biases, and uncover the exact latency at which the phase of ongoing EEG oscillations modulates perception?

First, we use simulations to control the exact timing and oscillatory frequency at which a phase modulation of perceptual outcome is inserted into an artificial EEG dataset. We then assess the latency at which a significant phase difference between two conditions can be detected, and verify that this latency can be vastly underestimated.

Secondly, we introduce the white noise (WN) paradigm, based on linear-systems analysis (Marmarelis and Marmarelis, 1978) and reverse correlation methods (Ringach and Shapley, 2004), which are used to characterize the systematic relationship between visual stimulation and brain response, i.e., the impulse response function (IRF). This approach has been applied to EEG to measure evoked potentials (VESPA; Lalor et al., 2006) or to reveal "perceptual echoes" (VanRullen and Macdonald, 2012). Once extracted, a simple convolution with these IRFs can be used to model the brain's EEG response to any new WN sequence presented (Ringach and Shapley, 2004). Within these new WN sequences, we also embedded near-perceptual threshold targets, which had a medium gray luminance level with the same properties as any other frame in the sequence. Accordingly, they did not affect the convolution result. Thus, we effectively removed the target-evoked response from our signal, and only modeled the background oscillations. Therefore, we could measure the real latency at which this background oscillatory phase impacts visual perception.

Materials and Methods

Measuring phase differences

Both the simulation and the experimental paradigm relied on quantifying the phase difference between two conditions, whether simulated or based on the actual behavioral outcome of human observers. To this end, we used the phase opposition sum (POS) measure (VanRullen, 2016a), which relies on the intertrial phase clustering (ITPC, also phase locking value or factor; Tallon-Baudry et al., 1996; Lachaux et al., 1999) and is computed as follows:

$$ITPC_{(t,f)} = \left| \frac{1}{n} \sum_{k=1}^{n} e^{(i\varphi_k(t,f))} \right|$$

The ITPC is, for a given time point t, and frequency of interest f, the norm of the complex average across n trials of the vector with unit length and phase φ. Consequently, the POS (VanRullen, 2016a) was computed as the sum of the ITPC for each condition corrected by subtracting the ITPC for the combined conditions, as follows:

$$POS = ITPC_A + ITPC_B - 2ITPC_{all}$$

Theoretically, the POS takes values between 0 and 2. A value of 0 arises when phase distributions for each con-

Acknowledgements: We thank T. Baudry for his help with the control experiment's data acquisition.

Correspondence should be addressed to Rufin VanRullen, Centre de Recherche Cerveau et Cognition, CNRS, UMR 5549, Pavillon Baudot, CHU Purpan, BP 25202, 31052 Toulouse Cedex 03, France, E-mail: rufin.vanrullen@cnrs.fr.

DOI:http://dx.doi.org/10.1523/ENEURO.0078-17.2017

dition are fully random (uniform) or when both conditions have their phases locked to the same angle. A value of 2 represents a perfect phase opposition between the two conditions, that is, all trials with behavioral outcome A have the same phase angle, and all trials with outcome B have the opposite angle (with no ITPC across the entire set of trials). Intermediate POS values reflect a partial phase opposition (a more plausible scenario) in which the phase modulates the probability of outcome A versus B. These theoretical values rely on the assumption that there is a perfect uniform sampling in the phases across all trials. Since this is unlikely to be the case in real experiments, it is better to compare the POS value to a corresponding distribution of surrogates POS values obtained by shuffling the labels between conditions. This will account for any bias in the underlying ITPC across all trials.

The POS was computed separately for each dimension of the dataset (time, frequency, channels, and subjects). The strength of the effect was evaluated at each of these dimensions by creating a surrogate distribution of 1000 permuted POS values through shuffling of the labels between conditions (i.e., creating random partitions). The mean and variance of the surrogate distribution were extracted to compute the *z* score of the observed POS, which was then transformed into a *p* value using the normal cumulative distribution function (for a description of this method and a comparison with other measures, see VanRullen, 2016a).

Simulations

In a first part, we used simulations of artificial datasets to look at how the ERP shape and frequency content, coupled with the time-frequency decomposition, influenced the latency at which a phase difference between two conditions could be detected, depending on the frequency of the phase modulation.

Creating artificial datasets

To evaluate the full extent of the effect, we systematically varied the frequency at which the phase modulation was inserted from 3.99 to 100 Hz in 24 logarithmically spaced steps. For each of the 24 frequency of interest, 100 artificial datasets (corresponding to the "subjects" in traditional EEG experiments) were created using an approach similar to that described in VanRullen (2016a). First, the background electrophysiological signal was simulated by creating 500 WN sequences drawn from a Gaussian distribution with a μ of 0 and a σ of 10 arbitrary units (Fig. 1). These sequences lasted 3 s ([−1.5 to 1.5 s]) and had a sampling rate of 500 Hz.

Once the artificial datasets had been generated, a phase modulation between two experimental conditions (i.e., trial groups) was artificially created using the phase of the frequency of interest at an arbitrarily chosen time point (40 ms after target onset; Fig. 1, green line). This phase was extracted by filtering the datasets at the frequency of interest and then applying a Hilbert transform. It was then used to assign an experimental condition label to each trial. Each of the two conditions was equally likely to occur overall (i.e., mean probability $\overline{p_A}$ of outcome A was equal to the probability $\overline{p_B}$ of outcome B). However, the likelihood of a trial outcome was modulated using a cosine function of the phase angle at the critical time, with a modulation depth (denoted as *mod* in the following equation) fixed at 0.4 (arbitrarily defined parameters). It was computed as follows:

$$p_{A(angle)} = \overline{p_A} + mod \cdot cos(angle)$$

In our case, this means that for trials at phase 0, there was a 70% chance of the trial yielding outcome A, while at the phase π, the trial had 70% chance to yield outcome B. Finally, an ERP was added to each trial. Both conditions had an ERP drawn from the same process, which was composed of a P1 and an N1 waves with (arbitrarily defined) parameters. The exact shape of the ERP differed slightly on each trial, as the parameters were drawn from normal distributions with known mean and σ defined for our purposes as follows (Fig. 1, average representation): the P1 mean amplitude was fixed at 20 units (amplitude σ of 5) and its mean peak latency at 65 ms (peak latency σ of 10 ms) with a mean duration of 50 ms (duration σ of 10 ms). The N1 mean amplitude was fixed at 30 units (amplitude σ of 10) with a mean peak latency of 155 ms (peak latency σ of 25 ms) and a mean duration of 130 ms (duration σ of 25 ms). The scale of the arbitrary units was defined with respect to the standard deviation of the originally generated WN signal, equal to 10 units. The final signals with the added ERPs were then used as the artificial dataset to analyze.

In the main simulations, Gaussian WN, with equal power at all frequencies, was used as a proxy for brain activity, as described above. Using pink noise instead of WN could be deemed more biologically plausible; however, pink noise (or 1/f) signal is characterized by a power spectrum decreasing as a function of frequency, meaning that differences of signal-to-noise ratio between frequencies could then have confounded our ability to detect phase effects. Nonetheless, we also performed control simulations using pink noise, to verify that the conclusions held with more biologically plausible input data. The same parameters as above were used. These control simulations gave comparable results, which are not presented here for the sake of brevity.

Extracting the latency of the phase difference

Once the final artificial datasets had been created, the time-frequency information was extracted using a wavelet decomposition, varying the numbers of cycles (logarithmically) from three cycles at 3Hz to 8 cycles at 100 Hz for each of the 50 (logarithmically spaced) frequency steps. The time course of the phase modulation was evaluated by looking at the time course of significant phase opposition using the *p* values extracted (see above, Measuring phase differences). For the purpose of these simulations, we assume that the rhythmic modulation frequency is known, and we aim to derive the latency of the effect. To this end, we restricted our analysis in time and frequency to an analysis window spanning 800 ms around the true latency of the phase modulation (i.e., from −360 to 440 ms) at the actual frequency at which the phase modulation had been introduced in the dataset. For each of the

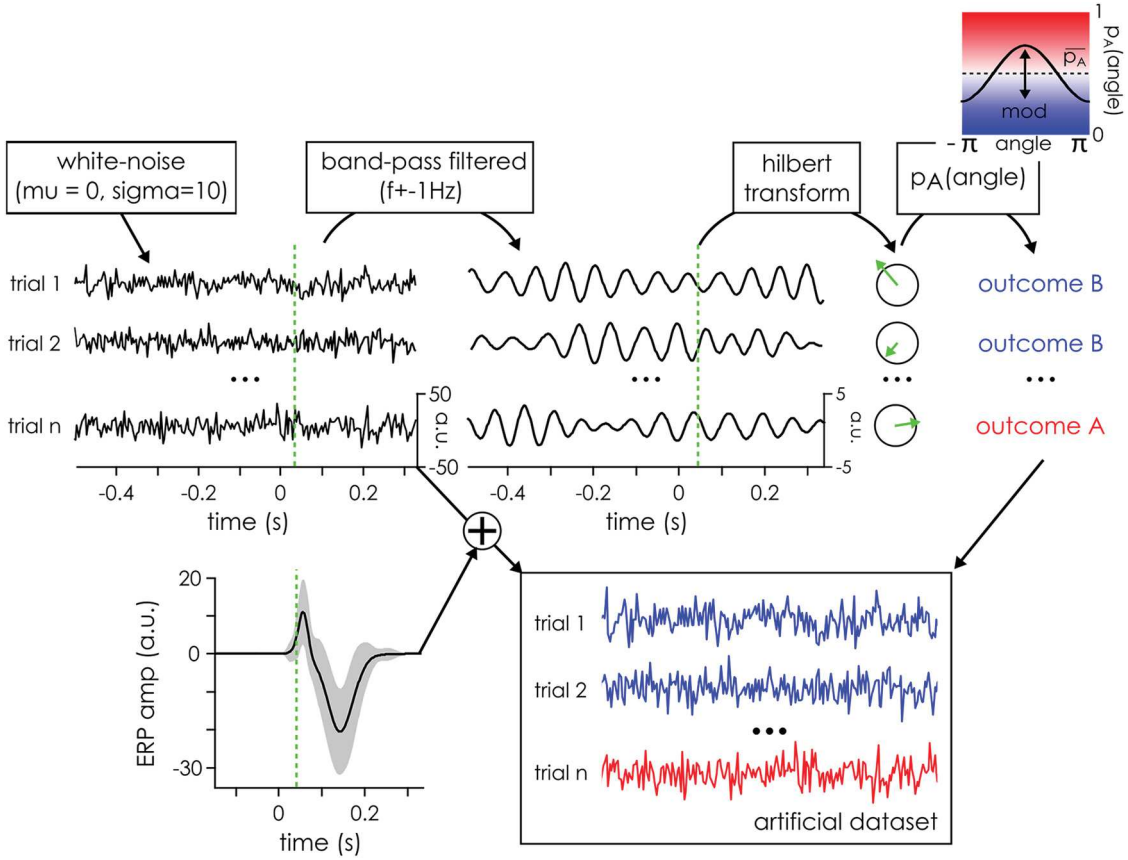

Figure 1. Illustration of artificial datasets creation for the simulation. The artificial signal was initialized using WN drawn from a Gaussian distribution with $\mu = 0$ and $\sigma = 10$ arbitrary units. These random data were then bandpass filtered at the frequency of interest plus or minus 1 Hz, and a Hilbert transform was applied to extract the phase at 40 ms after time 0, the time of target presentation. The phase angle at this time was then used to separate the trials between outcome A and B, with a given probability following a cosine function. Finally, an outcome independent ERP wave form (with slight random variations between trials) was added to each trial's signal to create the final artificial dataset.

100 artificial datasets, the time course of significance of the POS was evaluated by only keeping *p* values reaching or exceeding a Bonferroni threshold computed so as to correct for multiple comparisons across the 170 time points of the analysis window. This was taken as evidence for a significant phase difference between the two conditions at that particular latency. The time courses for each of the artificial datasets were then aggregated by computing the percentage of the simulated datasets which showed a significant POS at each time point. From this, we also computed the mean latency of the largest cluster for each of the time courses of the 100 datasets as a measure of central tendency. These latencies were then aggregated across all artificial datasets by looking at the confidence intervals (CIs) for the median latency of the phase modulation across artificial datasets, using the following formula by McGill et al. (1978), where q1, q2, and q3 are, respectively, the 1st, 2nd (or median), and 3rd quartile of the distribution and *N* is the number of artificial datasets.

$$95\%CI = q2 \pm 1.7 * \frac{1.25(q3 - q1)}{1.35\sqrt{N}}$$

We also computed the significance of the temporal distortion effect using a Wilcoxon sign rank test comparing the observed median latency to the true modulation latency of 40 ms, using an α level of 0.01.

The WN paradigm

In the experimental part, we introduce a paradigm based on linear-systems analysis methods, to help us uncover the true latency at which phase influences perception, avoiding the pitfalls induced by target-evoked responses.

Participants, stimuli, and procedure

Twenty-one subjects (mean age of 28.04 years, SD: 3.97 years, 23–39 years old) took part in the experiment after giving written informed consent. One subject was removed due to technical issues during the EEG recording, thus 20 subjects were analyzed (10 women, 15 right handed, normal or corrected-to-normal vision, no history of epilepsy). The experiment consisted of two testing sessions of ~1 h each, composed of eight blocks of 48 WN (random luminance) sequences (WN). Each WN sequence lasted 6.25 s and had, on average, a flat power spectrum between 0 and 80 Hz. They were presented in a

peripheral disk at 7° of visual angle from fixation point and subtending 7° of visual angle on a cathode ray monitor (resolution of 640 by 480 pixels and refresh rate of 160 Hz) situated 57 cm from the chin rest. In both sessions, participants had to detect near-perceptual threshold targets embedded in the stimuli sequences. In each trial, there were two to four targets composed of a lighter disk at the center with a darker surrounding annulus. They were presented within the peripheral flashing disk on a medium gray background disk for one frame only. The mean luminance of the target was always a medium gray, identical to the mean gray level of the WN sequence; only the contrast within the target was manipulated. A staircase procedure was conducted over the first 100 targets (i.e., ~30 trials) of each session using the *Quest* function (Watson and Pelli, 1983) in the PsychToolBox (Brainard, 1997). The contrast between the darker annulus and the lighter inner circle was adjusted to converge to the luminance contrast at which people perceived 50% of targets on average. This contrast was then kept for the remainder of the session. Session 1 data revealed that the visibility of the target was influenced by immediately surrounding luminance values (see below, Classification image). This influence could have masked the (potential) effect of oscillatory phase. Therefore, in session 2, we decided to remove any luminance fluctuations around targets, setting 14 frames before (i.e., 87.5 ms) and 11 frames after (i.e., 68.75 ms) the target to the same medium gray value.

Using a similar design and protocol, we ran a control experiment to assess whether the fluctuation-free periods could be used by subjects ($N = 6$) to detect the targets. The only difference was that for 1/3rd of the suppressed luminance time windows no targets was presented, creating "catch trials" of sorts. As with the main experiment, any button press after a target (from 150 to 800 ms) was counted as a hit. The same response window was used for catch trials: any button press within this 650-ms time window after the moment where a target would have been was counted as a false detection. The percentage of catch trials where a subject pressed the button was computed, similarly to the percentage of detected targets. Separately, we also assessed the false alarm rates of subjects, i.e., button presses falling outside of response time windows for either presented targets or catch trials. To get an estimate of the false alarm rate that is comparable with the other detection rates, we counted the number of 650 ms long time windows outside of target times and divided the number of false alarms by this number.

EEG recording and preprocessing

During the first session, the EEG to WN sequences was recorded using a 64-channels Biosemi system (1024 Hz sampling rate) with three external ocular channels recording the electrooculogram. Only the behavioral responses (to new WN sequences) were recorded in the second session. The preprocessing of the EEG data from session 1 was conducted using the EEGlab toolbox (Delorme and Makeig, 2004) in Matlab, with the following steps: (1) rejection and interpolation of noisy channels if necessary, (2) down-sampling of the EEG signal to 160 Hz to match the presentation rate of stimuli and thus facilitate the cross-correlation of the two signals, (3) notch filtering (between 47 and 53 Hz) to remove any artifacts due to power line, (4) average-referencing, (5) high-pass filtering (1 Hz) to remove any drifts in the signal, (6) creating epochs from −0.25 to 6.5 s around each WN sequence, (7) removing the baseline, i.e., the mean activity from −0.25 to 0 s before trial onset, and (8) manual artifact rejection where whole epochs were removed (as needed) to get rid of eye blinks and muscular artifacts. Once the data had been preprocessed, the IRF (also called VESPA by Lalor et al., 2006; or perceptual echoes by VanRullen and Macdonald, 2012) were extracted by cross-correlating the preprocessed EEG data with the WN luminance sequences, yielding 64 IRFs for each of the 20 subjects (one IRF for each EEG channel). These functions were then used to reconstruct the brain activity ("reconstructed EEG") to the new WN sequences presented in session 2, by convolution of the IRF with the WN. The target luminance (medium gray) was included in the WN sequence used to reconstruct the EEG. However, this value was no different from the surrounding values in the sequences, and thus there was no ERP evoked by the target in the reconstructed EEG. In session 2, the same random sequences (different from those in Session 1) were presented to all subjects in a randomized order so as to compare the visibility of the same targets across all subjects. This reconstructed EEG was epoched around each of the 821 presented targets ([−800 to +794 ms]). Finally, a wavelet transform was applied to extract the oscillatory characteristics of the reconstructed EEG at each frequency band (two to eight cycles, 50 log-spaced frequencies from 3 to 100 Hz). Note that we limited our analysis to frequencies below 80 Hz as our signal was sampled at 160 Hz (Nyquist frequency limit).

Measuring the phase difference between conditions

Further analyses of the phase differences between detected and undetected targets were evaluated only on the reconstructed EEG signals using the phase opposition method (see above, Measuring phase differences). Before extracting the z-scores, we computed the grand average POS values (whether real or permuted) by first summing the POS across electrodes and subjects to aggregate information along these dimensions. To increase the robustness, this step was repeated for the surrogate POS by randomly selecting (without replacement) surrogates for each subjects and summing across subjects again, a large number of times, thus yielding 100,000 grand average POS surrogates values. Using these grand averaged POS values, the z-scores and p values were then extracted as previously described (see above, Measuring phase differences). A false discovery rate (FDR) correction was applied across frequency and time points, using an α level of 0.05. Once a time-frequency region of interest was found, we extracted the topography of the effect by going back to single channel data and summing the phase opposition values across subjects and the time and frequency points composing the largest significant cluster. Here, again, to increase the robustness, we randomly

selected (without replacement) surrogates for each subject and summed the POS across subjects again. This was conducted 10,000 times before extracting the z-scores as previously described.

Measuring phase-dependent performance

To evaluate the strength of the effect and compare it with previous reports (Busch et al., 2009), the performance variability attributable to phase changes was computed on the fronto-central channel. The normalized hits ratio was extracted, for each subject and for each of the 11 phase bins as the proportion of hits in each bin, normalized by the mean performance of the subject. These were then averaged across subjects. A cosine fit was applied to the data and the resulting amplitude is reported as the amount of performance modulation for each study.

Classification image

We sought to test whether the luminance values of the WN sequences at any specific time point around the target time had an impact on perception. To this end, we used the classification image method (Ahumada, 2002), which can help identify which stimuli parameters have an impact on performance. For both sessions, we computed the mean luminance values across trials for hits and missed targets for each subject separately. An independent samples *t* test was then applied for each subject to compare the distribution of luminance values for detected versus missed targets. A FDR correction was applied across time and subjects to correct for multiple comparisons in each of the sessions.

Evoked response

We also wanted to confirm that the reconstructed EEG did not in fact show any sign of an evoked response to the target presented in the WN sequence. We computed the ERP for the EEG recorded in session 1 (as a control) and for the reconstructed EEG for session 1 and session 2. This was done by first removing the baseline activity ([−200 to 0 ms]) to each time course and then averaging the different signals across trials and subjects. Because the signal was visual in nature, we looked at the ERP over the central parieto-occipital channel (POz).

Measuring the correlation between recorded EEG and reconstructed EEG

To test how well the reconstructed EEG modeled the recorded EEG, we used the data from session 1, for which the recorded EEG was available. First, the total number of trials for each subject (which was variable due to artifact rejection) was sub-sampled to the trial count available for the subject with the fewest trials. Then, using a 10-fold cross-validation strategy, the model (i.e., IRF) was computed on 90% of trials, and consequently used to reconstruct the EEG for the remaining 10% of trials. This reconstructed EEG was then correlated (using the Pearson correlation) with the EEG actually recorded on the remaining 10% of (independent) trials. This ensured that both the recorded and reconstructed EEG were available for the same sets of trials, while avoiding any circularity in the analysis: using the same trials to compute the IRF and the reconstructed EEG could have led to spuriously high correlations between reconstructed and recorded EEG. Moreover, this also reduced the influence of a sampling bias: using a cross-validation strategy in correlating the signals allowed us to get a better estimate of the true underlying correlation between the two signals.

We also tested whether the EEG modeled using the IRF was more accurate for certain frequency bands by correlating the two signals (with the same cross-validation approach) filtered in five different frequency bands (Finite Impulse Response filter, δ: 2–4 Hz; θ: 4–8 Hz; α: 7–14 Hz; β: 14–28 Hz; γ: 30–60 Hz).

Since the correlation coefficients were not normally distributed (as revealed by a one-sample Kolmogorov-Smirnov test), a Fisher Z transform was applied to the data. It was computed as follows:

$$Z = 0.5 * ln\left(\frac{1 + r}{1 - r}\right)$$

Consequently, the mean transformed coefficients across all repetitions were extracted for each subject and channel and a one sample *t* test against zero was applied. The *p* values were corrected for multiple comparisons using a FDR correction across channels. To extract the correlation coefficient values for plotting purposes, the inverse Fisher Z transform was applied. Only the correlation coefficients at the strongest channel are presented across subjects. Note that using a nonparametric test (instead of a Fisher Z transform and a *t* test) gave equivalent results.

As a comparison, we also tested how much of the variability in the signal could be explained by the ERP to the target embedded in WN. Here, we used target ERPs (rather than IRF) as a model for the target-evoked activity. Finally, for comparison purposes, we also wanted a measure of how "noisy" EEG data normally is (i.e., in typical ERP paradigms without ongoing WN sequences). We tested this using a separate dataset from Busch et al. (2009), in which isolated targets were presented on a static background. The target-evoked ERP was again used as a model of evoked EEG activity. In both conditions (ERPs from targets embedded in WN; ERPs from isolated targets), the ERPs were extracted on 90% of the trials and then convolved with a sequence of target onsets for the remaining 10% of trials, in a 10-fold cross-validation approach. Consequently, the same correlation method between recorded and reconstructed EEG (modeled here by the ERP) was applied as described above for the IRF-based EEG reconstruction model.

Results

Simulations

Using artificially created datasets with a true phase modulation introduced at 40 ms after stimulus onset, we evaluated how early a significant phase opposition would be detected. We measured the time course of the *p* values of the POS (see Materials and Methods), aggregated across 100 simulated datasets (for more reliability). In

particular, we tested various frequencies of phase modulation from 4 to 100 Hz to see if this factor influenced the latency at which a significant effect could be measured.

At all frequencies of phase modulation, we did find evidence for a phase difference between conditions (as expected given that this phase modulation had been explicitly introduced in each dataset). But at frequencies of the phase modulation below 30 Hz, the observed phase modulation appeared to peak well before the true latency of the phase modulation. At lower frequencies (roughly below 20 Hz; Fig. 2*A*), this effect was even visible in the prestimulus time window. This temporal displacement seemed to be frequency dependent: the lower the frequency, the earlier the phase modulation was detected (Fig. 2*A*). In fact, the median latency of observed phase opposition effects was significantly different (Wilcoxon sign rank test) from the true phase modulation (at 40 ms) at all frequencies between 3.99 and 39.44 Hz. At 3.99 Hz, the median measured phase modulation latency peaked at −143 ms in the prestimulus window (95% CI: [−151 to −135 ms]). At 39.44 Hz, the median latency of the effect was at +37.5 ms, i.e., 3 ms before the true effect (95% CI: [35–39 ms]). At frequencies higher than 40 Hz, the median latency of the observed phase modulations effects was not significantly different from the true latency.

Interestingly, these temporal displacements seemed to be closely linked to the time-frequency content of the ERP: the latencies at which the phase difference could be measured were pushed back in time by an amount commensurate with the temporal spread of ERP spectral power at the corresponding frequency (Fig. 2*B*). As a matter of fact, when the ERP was removed from the simulated data, and all the analyses were performed again, the median latency of the phase modulation was not shifted (Fig. 2*C*). The only effect seen in this case was a smearing around the true latency of the phase modulation due to the window-based time-frequency decomposition. In fact, this smearing of information can account for the specific shape of the temporal displacement created by the ERP: there is a strong correlation (Pearson correlation, $r = 0.81$, $p < 0.05$, 95% ci for $r = 0.799–0.827$) between the error in the measured latency of the phase modulation effects (Fig. 2*A*) and the length of the window function at that frequency. Here, the smearing had a chimney-like shape typical of wavelet analysis: the smear of information varied from 750 ms at lower frequencies (at 4 Hz: window function of three cycles of 250 ms each) to 80 ms at the highest frequency (100 Hz) around the region of interest. The absence of latency distortion in the absence of ERP confirms the notion that it is the stimulus-evoked activity (and its temporal smear caused by any window-based time-frequency decomposition) that is likely responsible for the prevalence of EEG phase modulations reported to peak in the prestimulus time window in the relevant literature (for a review, see VanRullen, 2016b).

We showed that a phase difference between the two conditions in our artificially created datasets could be measured well before the latency at which the actual phase modulation was introduced. This temporal displacement could be quite dramatic at lower frequencies, with peak effects being apparently pushed by almost 200 ms. It was evident that the time-frequency content of the ERP was responsible for this apparent shift for at least two reasons: first, the measured latency of the phase modulation directly followed the left edge of the time-frequency content of the ERP (Fig. 2*B*); second, the temporal displacement disappeared when the analysis was replicated without an ERP (Fig. 2*C*).

Our simulations highlight a large uncertainty regarding the latency of EEG phase modulation. Importantly, this uncertainty is expected whenever ERPs are present and signal filtering is used, and would appear to be unavoidable for any experimental measure of phase-dependent perception. Is there a way, then, of uncovering the true latency at which ongoing phase affects stimulus detection?

The WN paradigm

Here, we sought to experimentally investigate the time course of phase effects using the WN paradigm, i.e., using WN sequences to constrain ongoing brain oscillations in a predictable manner. The IRF can serve to model the relationship between fluctuations of luminance values in the sequence and the brain response at different delays (Fig. 3; Ringach and Shapley, 2004; Lalor et al., 2006; VanRullen and Macdonald, 2012). It is extracted by cross-correlating WN sequences with the concurrently recorded EEG. These IRFs can then be used to reconstruct the brain activity to any new WN sequences presented. This reconstruction was done by convolving the IRFs with the new WN sequences presented in the second session (Fig. 3). We call this signal the reconstructed EEG: a model of the brain activity in response to the WN sequences.

Before investigating the time course of phase effects, we verified that our new paradigm met specific requirements. First, we sought to quantify the correlation between the modeled brain activity (the reconstructed EEG) and the real data (recorded EEG).

Correlating the reconstructed and recorded EEG

We first quantified the similarity between the recorded EEG (in session 1) and the EEG reconstructed using the IRF as a model of the background activity in the EEG. This was achieved by correlating the two signals using a 10-fold cross-validation approach: the IRF was computed on 90% of trials, and then used to model the EEG response for the remaining 10% of trials, which we then correlated with the recorded EEG.

A one-sample *t* test revealed that the distribution of the mean (Z transformed) correlation coefficients using the raw signal was significantly different from 0 on all channels, with the mean correlation coefficient across subjects reaching a maximum on electrode Oz with a r of 0.091 ($t_{(19)} = 7.78$, $p = 2.51*10^{-7}$, 95% CI for $r = 0.066–0.115$). We also investigated whether filtered signals would yield a better correlation depending on the frequency band. We found that the correlation strength was strongest for the α-band on channel Oz (mean $r = 0.163$, $t_{(19)} = 8.21$, $p = 1.14*10^{-7}$, 95% CI for r: 0.121–0.204), closely followed by the θ-band on channel POz (mean $r = 0.125$, $t_{(19)} =$

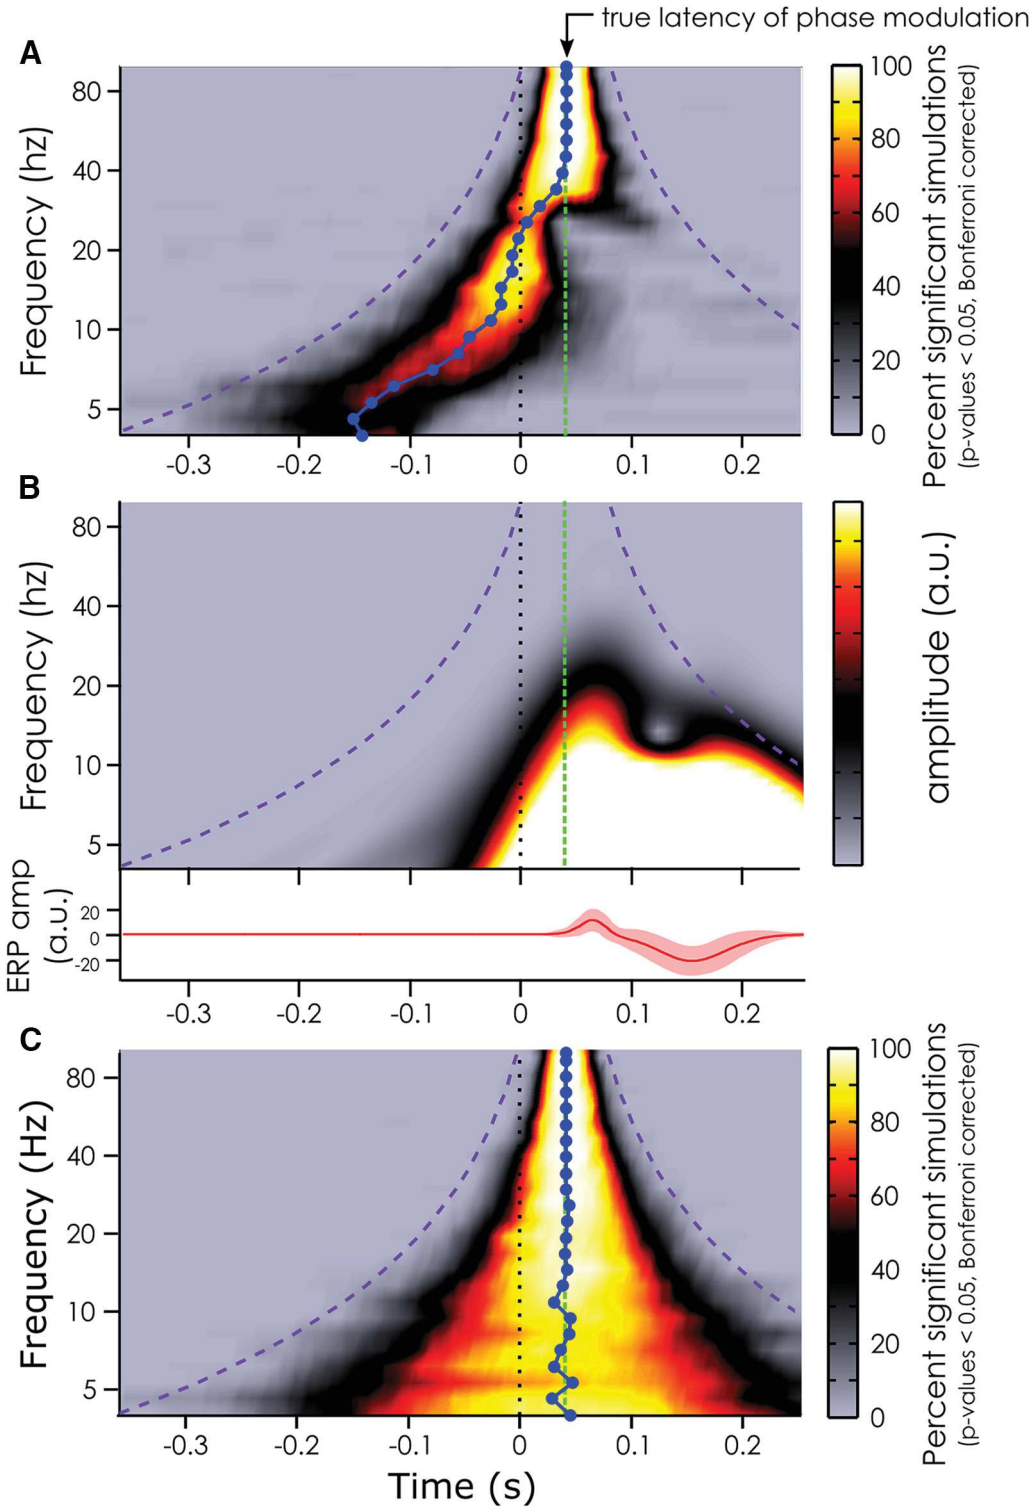

Figure 2. Simulation results. ***A***, Median latency (blue line) at which phase modulation effects can be measured depending on the frequency of the phase modulation introduced at 40 ms (true latency represented by the green dashed line) when the ERP is included in the artificial datasets. The color bar indexes the percentage of significant datasets at each time point after Bonferroni correction. The purple dashed line represents the outer edge of the window function (Morlet wavelets). ***B***, Representation of the evoked response (lower panel) included in the artificial data in ***A***, and its time-frequency content (upper panel). The color bar represents the oscillatory amplitude (arbitrary unit) at each frequency and time point. **C**. Same analysis as in ***A*** but without ERP included in the artificial datasets. Note the absence of temporal latency distortion in this case. Note also that the temporal smearing created by the window function is slightly shorter than the window duration (i.e., it does not reach the purple line) as Morlet wavelets have Gaussian tapers on either end.

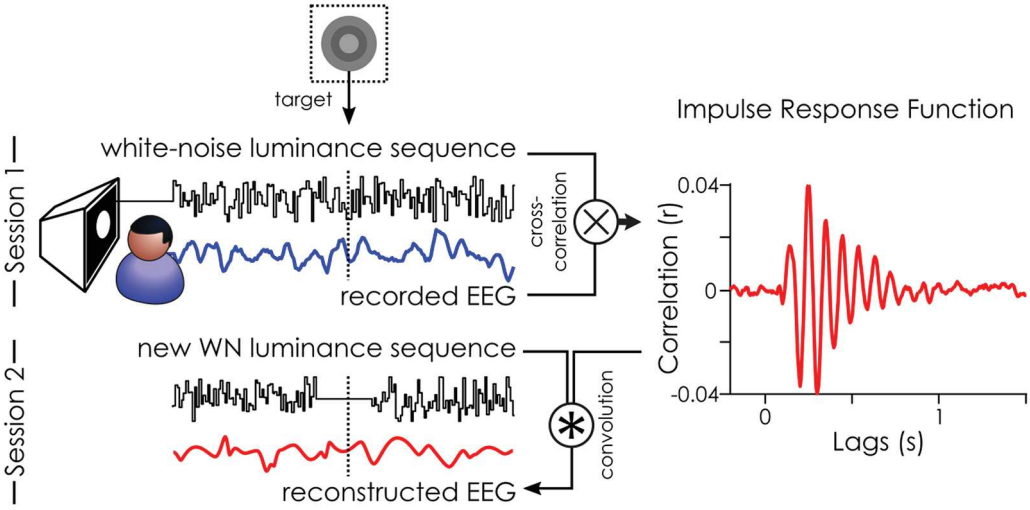

Figure 3. White-noise (WN) paradigm. The IRF to WN sequences can be extracted by cross-correlating the stimuli sequence with the recorded EEG: here, an example IRF is shown from one subject on electrode POz. This is what we did in the first session of our experiment. This IRF can, in turn, be used to reconstruct the brain activity (reconstructed EEG) to any new WN sequence by convolution. This was done in the second session of the experiment. The stimulus fluctuations around the target were removed in the second session to avoid any target masking by the luminance (Fig. 5).

13.69, $p = 2.69*10^{-11}$, 95% CI for r: 0.105–0.144). The other frequency bands had lower correlation coefficients, with, respectively, a mean *r* of 0.058 (channel POz) for the δ-band ($t_{(19)} = 7.28$, $p = 6.60*10^{-7}$, 95% CI for r: 0.041–0.075), a mean *r* of 0.071 (on channel Oz) for the β-band ($t_{(19)} = 7.80$, $p = 2.42*10^{-7}$, 95% CI for r: 0.052–0.090) and a mean *r* of 0.027 (on channel POz) for the γ-band ($t_{(19)} = 6.16$, $p = 6.37*10^{-6}$, 95% CI for r: 0.018–0.037).

To better evaluate the quality of our EEG reconstructions, we decided to run two control analyses, as a basis for comparison. First, we wanted to know how much of the variability in the signal could be explained using the ERP instead of the IRF as a model of the EEG activity to targets presented in WN (for details, see Materials and Methods, Measuring the correlation between recorded EEG and reconstructed EEG). Using the ERP as a model for target evoked activity led to a correlation strength on par with that obtained using the IRF as a model of background activity. The correlation coefficient reached its maximum across subjects over the left central parietal channel (CP1) with an r of 0.099 ($t_{(19)} = 10.92$, $p = 1.25*10^{-9}$, 95% CI. for $r = 0.080–0.118$, see "targets in WN"; Fig. 4). Secondly, we also evaluated the amount of signal variability that can be explained by the ERP to isolated targets in a more "typical" visual-evoked potential paradigm (i.e., without the concurrent WN stimulation), using the data from Busch et al. (2009). This was done as a way to estimate the "noise" in a typical EEG setting, for comparison purposes. A one-sample *t* test revealed that the distribution of the mean (Z transformed) correlation across subject reached a maximum over channel (Pz) with an r of 0.115 ($t_{(13)} = 9.57$, $p = 2.97*10^{-7}$, 95% CI for r: 0.089–0.141; Fig. 4, isolated targets). In conclusion, using the IRF as a model for reconstructing the EEG is a useful, although far from perfect (i.e., $r < 0.2$), characterization of the recorded EEG; although the obtained EEG reconstruction was clearly noisy, the reconstruction error was not much higher than that observed with ERP-based models of target-evoked activity.

Behavioral results

In both sessions, the participants had to detect targets embedded in the WN luminance sequences. The contrast of the target with regard to its medium gray background was adjusted to achieve 50% performance at the beginning of the experiment. The mean hit rate across the 20 participants for session 1 was 50.38% (SD: 9.17%) and 45.76% (SD: 10.88%) for session 2. For both sessions, the false alarm rate was relatively low with a mean false alarm rate of 1.65% (SD: 1.35%) and 1.65% (SD: 1.49%) for sessions 1 and 2, respectively. A control experiment revealed that the suppressed luminance fluctuations could not be used by subjects to detect the targets. The mean response rate to catch trials across subjects (i.e., suppressed luminance without targets presented) of 8.94% was not significantly different from the mean false alarm rate of 6.30% (Student's *t* test, $p = 0.646$). Moreover, both were much lower than the mean detection rate across subjects of 52.80% (Student's *t* test, $p < 0.002$).

A classification image analysis of the data from session 1 showed that the luminance values immediately surrounding the target had a large impact on target detection. Systematically, higher luminance values just before and after the target led to decreased visibility while lower luminance values led to increased visibility (Fig. 5*B*). In session 1, all 20 subjects showed a significant difference in the mean luminance values between detected and missed target trials (independent sample *t* test for each subject, with FDR correction, all 20 peak *p* values below $p < 0.0008$), and this difference affected at least 3 separate time points for all subjects (Fig. 5*B*). For this reason, we decided to remove luminance fluctuations around the target. In the second session, 87.5 ms (i.e., 14 frames) before the target and 68.75 ms (i.e., 11 frames) after the

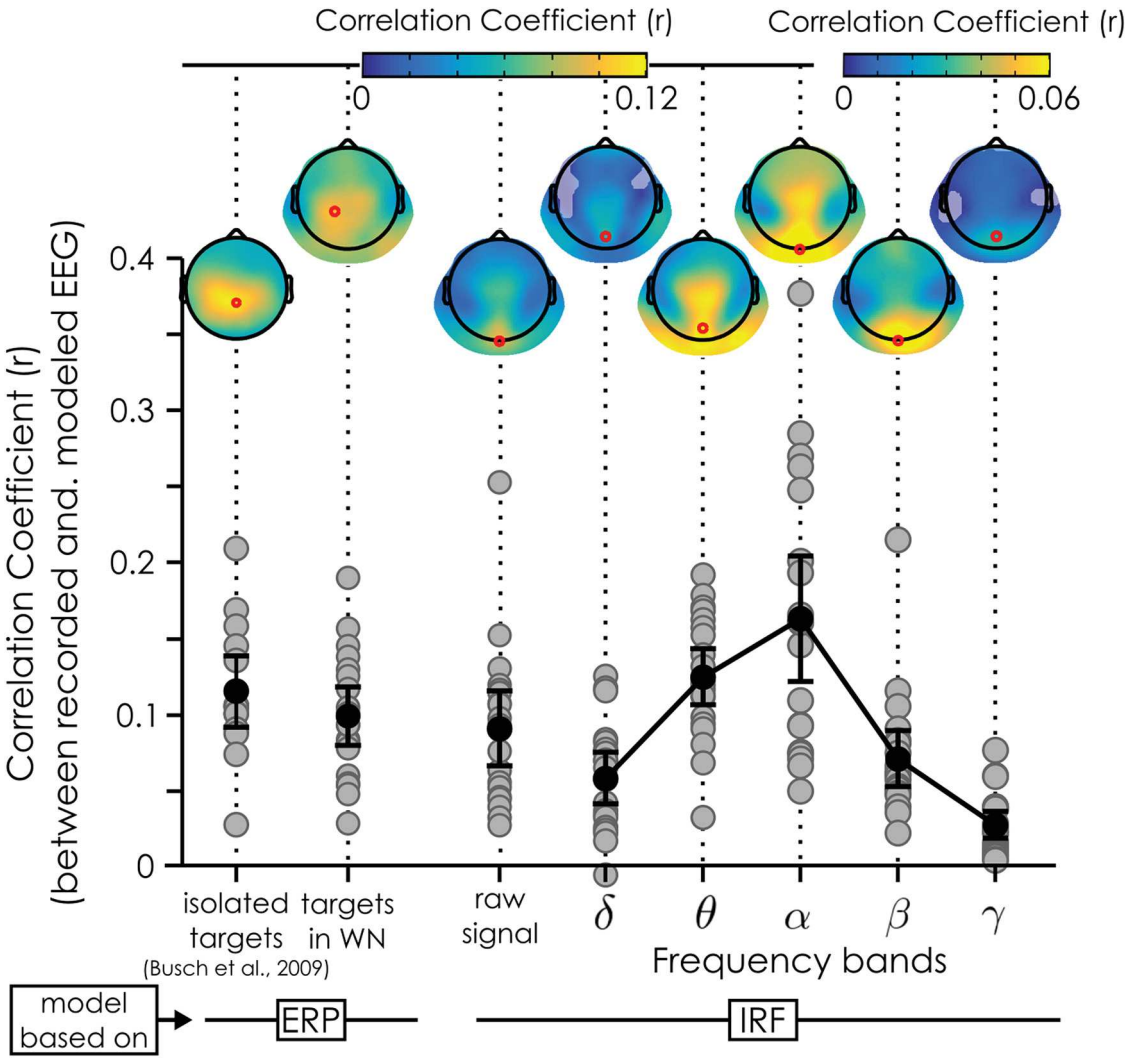

Figure 4. Correlation between reconstructed and recorded EEG. Two models of brain EEG activity were tested: one based on the ERP to targets (left) and the other based on the IRF to WN sequences (right). In both cases, we systematically correlated the prediction of the model with the single trial recorded EEG. For the ERP models, the "isolated targets" coefficients are based on the dataset presented in Busch et al. (2009), while the "targets in WN" are from our own dataset, with the ERP extracted relative to the targets embedded in WN (Fig. 5*C*). For the IRF models, we used the IRF to model the brain response to the WN sequences, and then correlated reconstructed and recorded EEG data using the raw signals, as well as signals filtered in different frequency bands (δ: 2–4 Hz; θ: 4–8 Hz; α: 7–14 Hz; β: 14–28 Hz; γ: 30–60 Hz). The gray dots represent the mean coefficient for each subject (1 dot per subject) across cross-validation runs at the maximum electrode (red dot on the topographies); the error bars represent the 95% CI of the mean (black dot) coefficient across subjects. The topographies represent the mean correlation coefficients across subjects and cross-validation runs. Shaded areas represent channels not significant after FDR correction. Note the difference of color scales for the β and γ correlation coefficients relative to the other topographies.

target were replaced with medium gray values (Fig. 5*A*), the same medium gray used as the target's background. In session 2, this manipulation effectively cancelled the apparent masking or facilitation effects observed in session 1 (Fig. 5*B*): 18 of the 20 subjects showed no significant luminance difference between hits and misses at any time point; the remaining two subjects showed a significant difference at only a single time point. Note that we also verified that all results described in the following section (Fig. 6) remained valid when these two subjects were discarded from the analysis.

The major advantage of using reconstructed EEG for our purposes is that there is no evoked response to the targets embedded in the stimulation sequences. This is because the luminance of the target frame (medium gray), which is taken into account when the convolution is done, is identical to the luminance of the surrounding frames. In other words, the presence of the target is invisible to the EEG reconstruction process (convolution). In fact, we can see that there is no evoked ERP to the target in the reconstructed EEG for session 2 (Fig. 5*C*).

Latency of measured phase modulation effects in the absence of ERP

Because the reconstructed EEG is virtually blind to the presence of target-evoked activity, we can estimate the

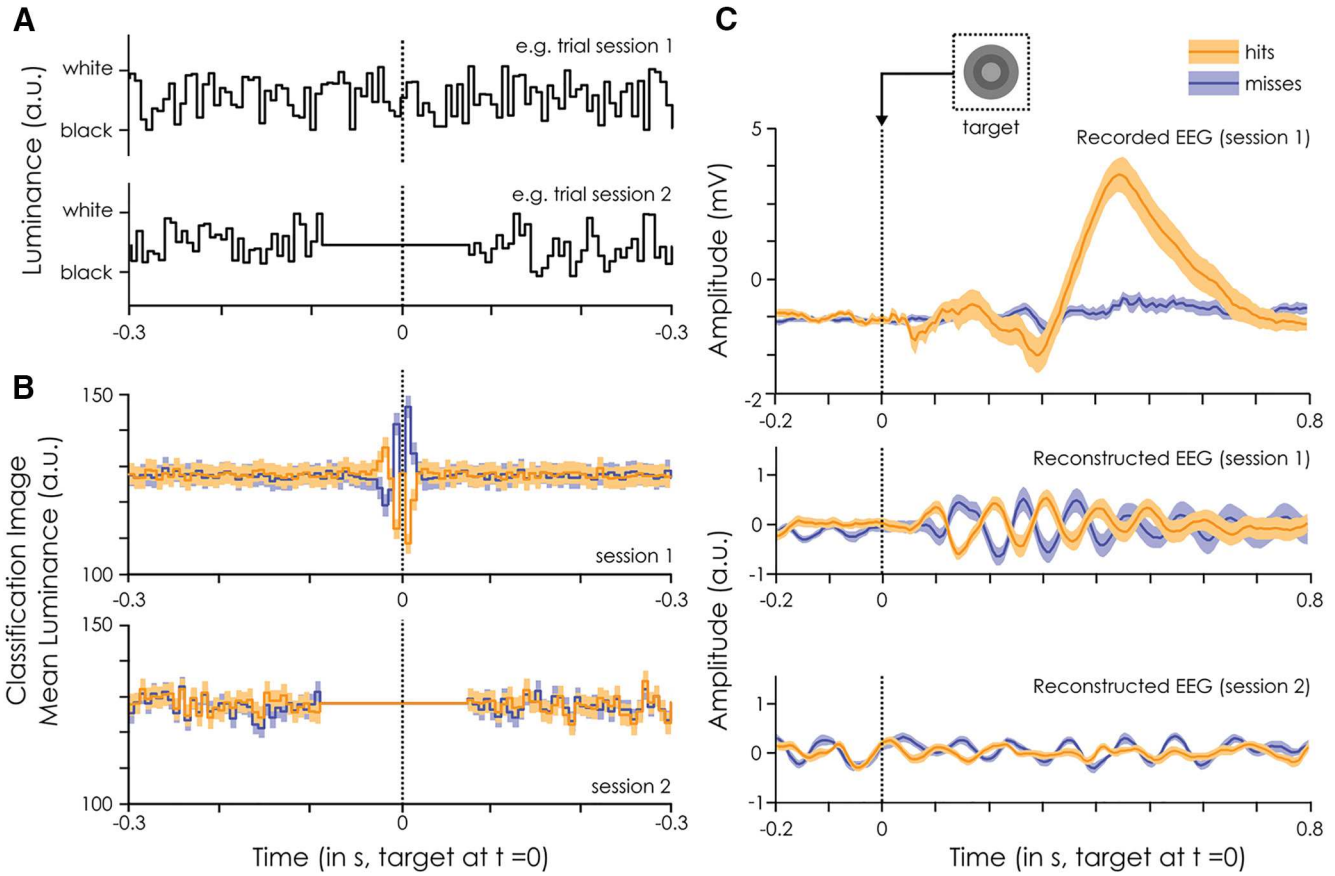

Figure 5. Classification image and evoked response. ***A***, Example stimuli sequences for session 1, when the EEG was recorded, and session 2, when only the behavioral response was recorded. The sequences are centered on the embedded target presentation (at $t = 0$ ms). ***B***, Classification image of the luminance values around the target (at $t = 0$ ms) for detected (orange) and missed (blue) targets. Darker colors represent the mean across subjects and trials. The lighter shades represent the standard error of the mean across subjects. The dotted line represents target presentation. Note that the same sequences were shown to all subjects in session 2. ***C***, ERP evoked by the detected (orange) and missed (blue) targets embedded within the WN sequences. These are computed separately for the recorded EEG for session 1 (top) and the reconstructed EEG for session 1 (middle) and 2 (bottom). Note the absence of visible target-evoked ERP in reconstructed signals for session 2 (bottom). While there seems to be a very strong phase opposition between hits and misses in the reconstructed EEG to session 1, this is likely to be an artifact of the strong relationship between luminance value and behavioral outcome illustrated in ***B***, rather than a direct relationship between phase and perception. Darker colors represent the mean across trials and the lighter shades represent the standard error of the mean across subjects. The dotted line represents target presentation.

background or "ongoing" signals around the time of target presentation. Thus, we can causally investigate the dynamics of the ongoing oscillations' influence on perception, without having to restrict our analysis to the prestimulus window due to evoked response contamination.

We applied a standard procedure to evaluate the oscillatory "phase opposition" among groups of trials in which the target was detected versus missed (VanRullen, 2016a); but in this case, the input data were the reconstructed EEG (obtained by convolution with the IRF) rather than an actually recorded EEG signal. The phase of ongoing oscillations in the θ-band (~6 Hz) seemed to be significantly different for detected versus missed targets, compatible with findings from previous studies (Busch et al., 2009). However, here the largest effect was present at ~75 ms after target onset ($p = 1.2242*10^{-6}$; Fig. 6*A*). A scalp topography revealed that phase opposition was maximal over frontal and occipital channels (Fig. 6*B*).

We then evaluated the amount of performance modulation that could be explained by the ongoing oscillatory phase at the peak of significance (i.e., 6 Hz and 75 ms), by looking at the normalized hit rate averaged across subjects for electrode Fz (see Materials and Methods). We found that the phase accounted for 11.62% of the performance modulation (Fig. 6*C*, left hand side). As a comparison, the results from Busch et al. (2009) have also been replotted here. In that experiment, after realigning phase bins across subjects to align maximal performance at zero phase (a procedure which was unnecessary in the present experiment), the ongoing EEG phase at 7 Hz, −120 ms before stimulus onset explained 15.62% of the performance fluctuations (Fig. 6*C*, right hand side).

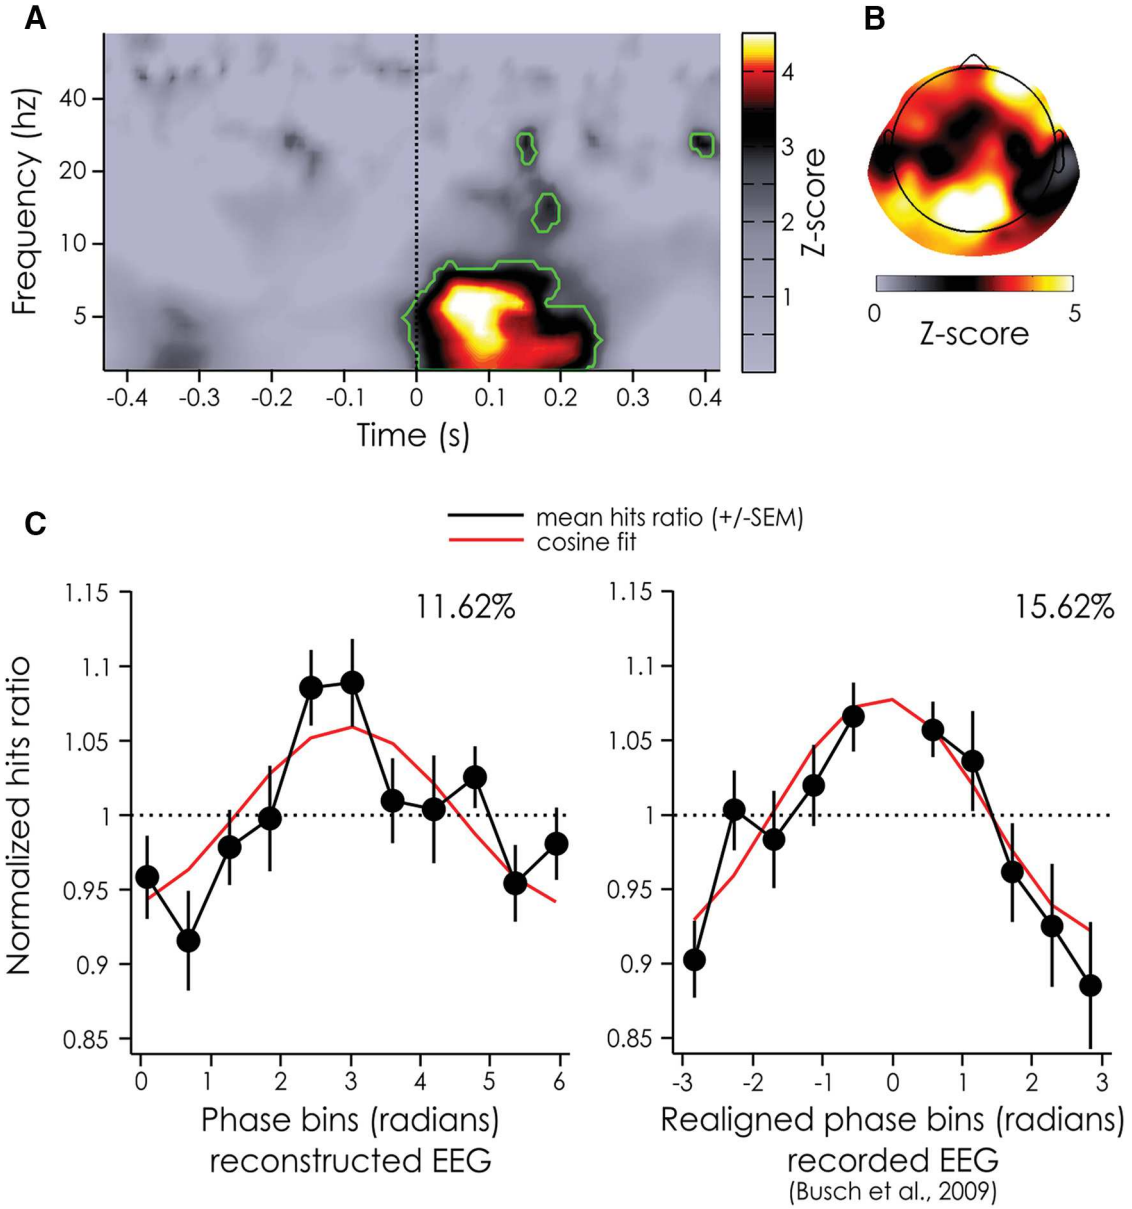

Figure 6. Z-score of POS for the reconstructed EEG and phase dependent performance. ***A***, Z-scores of the POS values aggregated across electrodes and subjects for the reconstructed EEG for each time and frequency point. Significant time points after a FDR correction ($\alpha = 0.05$) are outlined in green. ***B***, Z-scores of the POS values aggregated across subjects for the reconstructed EEG at the frequencies and time points inside the largest significant cluster (from 3–8.17 Hz and from −12.5 ms to 250 ms, ***A***). ***C***, Normalized hits ratio depending on the phase bin for the reconstructed EEG at the peak significance (left hand side: 6.1 Hz, +75 ms). For comparison, previously published results linking detection performance and prestimulus phase of actually recorded (rather than reconstructed) EEG signals are replotted on the right hand side (7 Hz, −120 ms; taken from Busch et al., 2009). Both phase dependence curves are measured on the same fronto-central channel. Note that in our experiment the phase bins were not realigned between subjects, but that they were realigned (to produce peak performance at 0 for all subjects) for the experiment by Busch et al. (2009). The black line represents the mean across subjects (left hand side: 20 subjects, right hand side: 14 subjects), the bars represents the SEM, the red line represents a cosine fit, whose modulation amplitude (in %) is noted above the curve.

To summarize, in this experiment, we presented a method to study the dynamics of ongoing oscillations' influence on perception, through the use of WN sequences. We found that the phase of θ-band ongoing oscillations is causally linked to the detection of near-perceptual threshold targets. This rhythmic influence of θ oscillations extended in time up to 200 ms after target onset, but was the largest at ~75 ms after stimulus and could explain up to ~12% of the performance modulations.

Discussion

In this paper, we sought to reconcile the apparent gap between prestimulus EEG phase effects reported in the literature and the idea of phase as actively shaping perception.

In a first section, we used simulations to show that a phase modulation, artificially introduced 40 ms after stimulus onset, could be reliably detected in the prestimulus period. This depended on the ERP time-frequency content. In fact, we found that, at lower frequencies, the peak phase difference between conditions happened up to 200 ms before the true phase modulation.

The measured latencies for a phase modulation introduced at 7.08 Hz are consistent with the results of Busch et al. (2009), who found an effect of the ~7-Hz EEG phase on the detection of near-threshold peripheral stimuli about −120 ms before stimulus onset. Of note, the 7.08-Hz simulation resulted in a median latency of −79 ms (95% CI: −88.5 to −69.5 ms) with 48/100 simulated datasets having a significant POS at −120 ms. Turning our result around, the prestimulus latency observed by Busch et al. (2009; −120 ms) could in fact be compatible with the notion that the critical phase for stimulus detection was the phase that occurred during (rather than before) stimulus processing.

It is important to highlight that the latencies reported here via simulations should not be used as a definite guide to determine the latency of the true phase modulation in a new EEG experiment. This cannot be done, because various parameters in the new EEG experiment would differ from those used in our simulations, such as the exact ERP shape on each trial, the parameters of the filters used, or the true latency at which the phase modulation occurs in the brain. All these could have a large impact on the measured latency of the phase modulation. In other words, this simulation was only intended as a proof of concept and cannot serve as an exhaustive or quantitative evaluation.

The implications of our simulations are twofold. First, any study of phase modulation of perception involving sizeable stimulus-evoked activity (that is most, if not all, studies) should be susceptible to temporal distortions of the oscillatory phase effects. Second, the numerous previous reports of prestimulus phasic influences on perception (such as Busch et al., 2009; Mathewson et al., 2009) may still be compatible with a true latency of phase modulation at or even after the stimulus onset.

Although the evoked activity leads to a temporal distortion of oscillatory activity, the true latency of phase modulations cannot be retrieved, in our case, by simply subtracting the stimulus-evoked activity (i.e., the ERP) from the EEG signal –although this approach has been successfully used in other studies (Tallon-Baudry and Bertrand, 1999) to untangle the relative contribution of “evoked” (i.e., phase locked) and “induced” oscillations (i.e., nonphase locked). In our case, the postulated phase modulation effects are actually phase-locked to the stimulus onset (i.e., we hypothesize that all hits tend to show a similar oscillatory phase relative to stimulus onset, and likewise for the misses). Subtracting condition-specific ERPs from each trial would effectively flatten any systematic phase differences between the two groups. On the other hand, simply subtracting the common ERP for hits and misses from single-trial data are also not a viable option, because it is likely to introduce artefactual phase opposition when the ERPs between conditions differ.

In a second section, we presented an experimental method based on linear-systems analysis methods (Ringach and Shapley, 2004; Lalor et al., 2006; VanRullen and Macdonald, 2012): the WN paradigm. This paradigm effectively allows us to bypass the effects of the target-evoked ERP. Thus, we could estimate, for the first time, the true latency at which ongoing EEG oscillations influence perception. Using the EEG IRF (VanRullen and Macdonald, 2012) to WN luminance sequences, we reconstructed (rather than recorded) the brain activity of subjects to new sequences, by convolving the IRF with the stimuli sequences. The reconstructed EEG allowed us to estimate the ongoing (background) oscillations before and after each presented target without the evoked response which could have biased our ability to measure the phase modulation (as shown in our simulations).

We found that the reconstructed EEG is a (relatively) good model of the recorded EEG: the two signals were significantly correlated, with the highest correlation in the α- and θ-bands. The coefficients were relatively small, meaning that at most $r^2 = 0.163^2 = 2.7\%$ (at the maximum channel in the α-band) of the variance of the EEG signal was driven by the luminance sequence in a way that we could predict with the IRFs.

On the one hand, it might seem that this is a very small amount of variance explained. For example, using a similar approach, Kayser et al. (2015) predicted the stimulus-driven response of neurons to auditory stimulation (Kayser et al., 2015). They constructed various models of spike production based on linear spatiotemporal-response filters but also added various nonlinear factors on the prediction (Kayser et al., 2015). The best model, which included phase dependent variations in the sensory gain and the background firing of cells, could explain up to 30% of the variance (r^2) in the original signal. This increase of variance explained (relative to the present findings) can be attributed to the method used: LFP and multi-unit activity recordings are less subject to noise than scalp measurements like EEG. In fact, Lalor (2009), who also evaluated correlation strength between reconstructed and recorded EEG using the VESPA as a model for brain activity, found correlation values similar to ours: he reported a mean correlation coefficient of 0.084 over channel Oz (Lalor, 2009). Moreover, using a nonlinear (quadratic) model for the VEP only marginally improved the amount of signal explained by the model (mean r across subjects of 0.097; Lalor, 2009).

On the other hand, it is important to remember, that EEG is a noisy method of recording, especially when looking at single-trial data (Picton et al., 2000). The small amount of signal variability explained in the present and Lalor’s experiments (Lalor, 2009) might thus be due to the inherently noisy nature of EEG recordings. The rest of the EEG signal variance could reflect other cognitive functions (such as endogenous attention or arousal levels) or noise in the system, in the EEG recording or in the EEG reconstruction procedure. Of more interest is the fact that the IRF-based models reported here are on par with ERP-based models, explaining just as much variability in

the signal, if not more. ERPs are however routinely used in experiments and considered a meaningful measure of brain activity.

Finally, despite the apparently small percentage of variance of the original EEG signal explained in our study, it could well be that this small portion of the signal is the only oscillatory activity whose phase actually modulates visual perception. In other words, we may only be able to predict a small portion of the recorded EEG signal, but our prediction was sufficiently accurate to capture most of the existing relation between EEG oscillatory phase and perception (Fig. 6*C*, ~12% modulation in our case, compared with 16% modulation in the study by Busch et al., 2009). Using the reconstructed EEG, we found that the phase of θ (~6 Hz) oscillations was related to the detection of near-perceptual threshold targets, as suggested before by Busch et al. (2009). But crucially, the phase between detected and missed targets was significantly different from 50 to 150 ms after stimulus presentation, with a peak at 75 ms over fronto-occipital channels. This new approach thus allowed us to uncover the true latency at which ongoing EEG oscillatory phase influences visual target detection. This could not be done before, as the presence of the ERP usually biases the detection of phase effects toward prestimulus time windows, as demonstrated in the first section. This, however, is not a problem within the present approach because only the ongoing oscillations are modeled, not the ERP.

This paradigm allowed us to explore the true latency of the effect of oscillatory phase, unbiased by evoked responses, but it also opens up a wide range of avenues of investigation. For example, it would be possible to investigate the effect of stimulus-driven ongoing oscillation amplitude on perception. In particular, this could help us untangle the contributions from stimuli driven versus top-down driven oscillations in different tasks. Indeed in our paradigm, only the ongoing oscillations directly driven by the stimuli are modeled, and thus we could evaluate their influence on performance, and compare this to the performance modulation of actually recorded EEG activity, in which top-down effects are also present.

Interestingly, this paradigm can also provide a bridge between EEG findings and other findings linking spiking activity to ongoing oscillations at the level of local neuronal populations. If the oscillations reflect the excitatory state of the population, the phase that should matter is the one expressed in the cortex at the exact moment when the stimulus information is processed. In fact, our results support this hypothesis. Further, our study lends support to the hypothesis that perception is rhythmic, and that ongoing oscillatory phase marks the underlying sampling mechanism of the environment. Here, the phase of the reconstructed EEG influences perception in a causal manner: it shapes our visual world as soon as the target enters the brain.

References

Ahumada AJ (2002) Classification image weights and internal noise level estimation. J Vis 2:121–131. CrossRef Medline

Brainard DH (1997) The psychophysics toolbox. Spat Vis 10:433–436. CrossRef

Busch NA, Dubois J, VanRullen R (2009) The phase of ongoing EEG oscillations predicts visual perception. J Neurosci 29:7869–7876. CrossRef

Busch NA, VanRullen R (2010) Spontaneous EEG oscillations reveal periodic sampling of visual attention. Proc Natl Acad Sci USA 107:16048–16053. CrossRef

Callaway E, Yeager CL (1960) Relationship between reaction time and electroencephalographic alpha phase. Science 132:1765–1766. Medline

Chakravarthi R, VanRullen R (2012) Conscious updating is a rhythmic process. Proc Natl Acad Sci USA 109:10599–10604. CrossRef Medline

Delorme A, Makeig S (2004) EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J Neurosci Methods 134:9–21. CrossRef

Drewes J, VanRullen R (2011) This is the rhythm of your eyes: the phase of ongoing electroencephalogram oscillations modulates saccadic reaction time. J Neurosci 31:4698–4708. CrossRef

Dugué L, Marque P, VanRullen R (2011) The phase of ongoing oscillations mediates the causal relation between brain excitation and visual perception. J Neurosci 31:11889–11893. CrossRef

Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cogn Sci 9:474–480. CrossRef

Fries P, Schröder J-H, Roelfsema PR, Singer W, Engel AK (2002) Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection. J Neurosci 22:3739–3754.

Haegens S, Barczak A, Musacchia G, Lipton ML, Mehta AD, Lakatos P, Schroeder CE (2015) Laminar profile and physiology of the α rhythm in primary visual, auditory, and somatosensory regions of neocortex. J Neurosci 35:14341–14352. CrossRef

Haegens S, Nácher V, Luna R, Romo R, Jensen O (2011) α-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. Proc Natl Acad Sci USA 108:19377–19382. CrossRef

Hanslmayr S, Volberg G, Wimber M, Dalal SS, Greenlee MW (2013) Prestimulus oscillatory phase at 7 Hz gates cortical information flow and visual perception. Curr Biol 23:2273–2278. CrossRef

Jacobs J, Kahana MJ, Ekstrom AD, Fried I (2007) Brain oscillations control timing of single-neuron activity in humans. J Neurosci 27:3839–3844. CrossRef

Kayser C, Wilson C, Safaai H, Sakata S, Panzeri S (2015) Rhythmic auditory cortex activity at multiple timescales shapes stimulus-response gain and background firing. J Neurosci 35:7750–7762. CrossRef

Lachaux J, Rodriguez E, Martinerie J, Varela FJ (1999) Measuring phase synchrony in brain signals. Hum Brain Mapp 8:194–208.

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.

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.

Lalor EC (2009). Modeling the human visual system using the white-noise approach. In: *Proceeding of the 4th International IEEE/EMBS Conference on Neural Engineering*, Antalya, Turkey, pp 589–592.

Lalor EC, Pearlmutter BA, Reilly RB, McDarby G, Foxe JJ (2006) The VESPA: a method for the rapid estimation of a visual evoked potential. NeuroImage 32:1549–1561. CrossRef

Marmarelis P, Marmarelis VZ (1978). The white noise method of system identification. In: Analysis of physiological systems: the white-noise approach. New York: Plenum Press.

Mathewson KE, Gratton G, Fabiani M, Beck DM, Ro T (2009) To see or not to see: prestimulus alpha phase predicts visual awareness. J Neurosci 29:2725–2732. CrossRef Medline

McGill R, Tukey JW, Larsen WA (1978) Variations of box plots. Am Stat 32:12–16. CrossRef

McLelland D, Lavergne L, VanRullen R (2016) The phase of ongoing EEG oscillations predicts the amplitude of peri-saccadic mislocalization. Sci Rep 6:29335. CrossRef

Nunn CM, Osselton JW (1974) The influence of the EEG alpha rhythm on the perception of visual stimuli. Psychophysiology 11:294–303. Medline

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. CrossRef

Ringach D, Shapley R (2004) Reverse correlation in neurophysiology. Cogn Sci 28:147–166. CrossRef

Tallon-Baudry C, Bertrand O (1999) Oscillatory gamma activity in humans and its role in object representation. Trends Cogn Sci 3:151–162.

Tallon-Baudry C, Bertrand O, Delpuech C, Pernier J (1996) Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. J Neurosci 16:4240–4249. Medline

VanRullen R (2011) Four common conceptual fallacies in mapping the time course of recognition. Front Psychol 2:365. CrossRef

VanRullen R (2016a) How to evaluate phase differences between trial groups in ongoing electrophysiological signals. Front Neurosci 10:426. CrossRef

VanRullen R (2016b) Perceptual cycles. Trends Cogn Sci 20:723–735. CrossRef Medline

VanRullen R, Koch C (2003) Is perception discrete or continuous? Trends Cogn Sci 7:207–213. Medline

VanRullen R, Macdonald JSP (2012) Perceptual echoes at 10 Hz in the human brain. Curr Biol 22:995–999. CrossRef Medline

Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. Percept Psychophys 33:113–120. CrossRef

4.3 CONCLUSIONS

This paper provided evidence for one of the methodological considerations presented by VanRullen (2016b), namely explaining pre-stimulus effects reported in most papers linking ongoing brain oscillations and visual perception. If oscillations reflect the neuronal excitability, their phases modulating the responses of the underlying neuronal populations, then the critical phase for perception should be the one present in the cortex *while* the information is being processed. Here, using complimentary approaches, we show that: first, these effects are due to the presence of a target-related evoked activity. This large signal with similar phases across trials effectively destroys the visibility of phase differences between conditions. Second, using the white noise paradigm, we show that when only background oscillations are modelled, without the event-related activity in the EEG signal, then we can the latency of phase effects we recover is more biologically plausible.

Moreover, we showed that the reconstructed EEG is in fact linked with the visual outcome: the phase of the theta band reconstructed EEG oscillation has a causal influence on visual perception. Interestingly, we can explain about as much variability in the performance of subjects as what had been reported previously using "spontaneous ongoing oscillations". This phase-dependent performance across subjects didn't need to be realigned, whereas most previous studies have had to realign the phases across subjects. This is likely due to the latency at which the effects could be measured. In fact, it is possible that the subject specific frequency of the effect might be slightly different for each subject, even if it lies within the theta band. These slight frequency difference mean that 100 to 150ms before the target is presented, from 150ms to 200ms before the true phase, each subject would have slight differences in which phase was correlated with their optimal output. Consequently, phases would need to be realigned to extract the average performance.

Finally, it also allowed us to disentangle the relative contribution of the different signals described in chapter 1 to the pre-stimulus phase effects reported in the literature (Busch et al., 2009; Drewes & VanRullen, 2011; Mathewson et al., 2009; VanRullen et al., 2011). In our model, the detection of the target can be explained by any of 3 relevant signals which are either spontaneous signals (S2), phase locked (S4) and or induced (S6) (as by definition the "irrelevant signals" S1, S3 and S5 cannot influence perception). In the case of pre-stimulus phase effects, the induced signal (S5) can be discarded, as they are, by definition, not phase locked to the targets. However, the analysis of pre-stimulus effects relies on phase

similarity across trials. A jittered phase with regards to the targets will results in a jittered phase across trials, and thus, these signals will be ignored in the analysis. In the case of signals S2 and signals S4, we expect the maximum influence of phase to be post-stimulus, which we can't see due to the evoked responses to the targets. In our taxonomy (presented in chapter 1), the *targets*' evoked responses are not part of signal S4. With regards to the WN sequences, the targets' evoked responses are (tentatively) part of signal S2: their phase occurs randomly. Since the reconstructed EEG won't be contaminated by S2 signals, and if the phase of signal S4 influences perception, then we can recover the post-stimulus effects. Based on results described above, we can conclude that the phase of the perceptual echoes influences perception (and at the right post-stimulus moment). This leaves open the question of whether the phase of spontaneous signals (S3) also influences perception. In fact, we have seen in our experiment that the phase modulation of perception by the S4 phase is almost as strong as the one found in normal EEG experiments. This suggests that the effects of S3 phase on target detection are probably small(er).

Chapter 5. AMPLITUDE AND PERCEPTION

5.1 INTRODUCTION

In the last chapter, we saw that the phase of background oscillations (as modelled by the reconstructed EEG) could modulate perception after target onset. Now, we ask whether this is also the case for the amplitude of oscillations. In particular, we investigate how, by dissociating the spontaneous ongoing signals (S2) from the constrained signals (S4), the white noise paradigm can allow us to better understand the influence of oscillations on visual perception. In the reconstructed EEG, only the stimulus constrained background oscillations are present (i.e. signal S4). The spontaneous alpha oscillations which have usually been found to influence perception as a function of attentional allocation (e.g. Rihs et al., 2007; Worden et al., 2000) are not captured by the reconstructed EEG. Thus, we use it to investigate the relationship between fluctuations in alpha amplitude and target detection outside of any fluctuations in alpha amplitude due to endogenous factors. This approach is complementary to non-invasive brain (such as TMS and tDCS) and visual stimulation techniques which have been used in the past to resolve this issue. One advantage of the WN paradigm is that it does not require the use of pre-defined frequency of stimulation, which is the case in most other stimulation techniques paradigm.

5.2 ARTICLE 3: ALPHA POWER MODULATES PERCEPTION INDEPENDENTLY OF ENDOGENOUS FACTORS (IN PREP)

Abbreviated Title (5 words)
Alpha power modulates perception causally

List of all Author Names and Affiliations

Sasskia Brüers[1,2] and Rufin VanRullen[1,2]
[1]Faculté de Médecine, Université de Toulouse Paul Sabatier, Toulouse, France
[2]Centre de Recherche Cerveau et Cognition, CNRS, UMR 5549, Toulouse, France

Author Contributions

RV Designed research, SB performed research, RV and SB analysed data and wrote the paper

Correspondence should be addressed to (include email address)

Rufin VanRullen, Centre de Recherche Cerveau et Cognition (CerCo), CNRS, UMR 5549, Pavillon Baudot, CHU Purpan, BP 25202, 31052 Toulouse Cedex 03, (rufin.vanrullen@cnrs.fr)

Number of Figures: 4
Number of Tables:
Number of Multimedia:
Number of words for Abstract: 321
Total number of words:

Acknowledgements

Conflict of Interest
Authors report no conflict of interest

Funding sources
This research was supported by the ERC grant P-CYCLES (N°614244) to RV.

Keywords (5 compulsory)
Alpha oscillations amplitude; EEG; vision;

Abstract

Oscillations are ubiquitous in the brain. Alpha oscillations in particular have been proposed to play an important role in sensory perception. Studies show that the amplitude of ongoing EEG oscillations in the alpha band is negatively correlated with visual outcome. Moreover, it also co-varies with other endogenous factors such as attention, vigilance or alertness. In turn, these endogenous factors influence visual perception. Therefore, it remains unclear how much of the relation between alpha and perception is indirectly mediated by such endogenous factors, and how much reflects a direct causal influence of alpha rhythms on sensory neural processing. Following the approach used in brain and visual stimulation study, we propose to disentangle the two causal routes (direct/indirect) by introducing modulations of alpha amplitude, independently of any fluctuations in endogenous factors. To this end, we use white-noise sequences to constrain the brain activity of 20 participants. One advantage of using white noise sequences is that, instead of having to choose one stimulation frequency (as previously done), we can use our brain's resonance properties which echoes the stimulation in a subject specific manner. Performing a cross-correlation between the white-noise sequences and the concurrently recorded EEG, we first model the systematic relationship between stimulation and brain response by extracting the impulse response functions. These are then used to reconstruct (rather than record) the brain activity linked with new random sequences (by convolution). Interestingly, this reconstructed EEG only contains information about oscillations directly linked to the white-noise stimulation; fluctuations in attention and other endogenous factors may still modulate EEG alpha rhythms during the task, but our reconstructed EEG is immune to these factors. We found that the detection of near-perceptual threshold targets embedded within these new white-noise sequences depended on the amplitude of the ~10Hz reconstructed EEG over parieto-occipital channels (*a priori* region of interest): around the time of presentation, higher amplitude led to poorer performance ($p=0.00018$, cluster correction). Thus fluctuations in alpha power (induced here by random luminance sequences) can directly influence perception. It is possible that these fluctuations could causally modulate attention, nevertheless, we can still rule out a causal "primacy" of endogenous factors: the relation between alpha power and perception is not a mere consequence of fluctuations in endogenous factors.

5.2.1 Introduction

When recording the electro-encephalography (EEG) in humans, the most prominent rhythm is the ~ 10Hz oscillations over the occipital cortex. The exact role of alpha oscillations in visual perception is still a matter of debate. Nevertheless, most current theories agree that alpha oscillations play an active inhibitory role in shaping our visual experience (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007; Mathewson et al., 2011; VanRullen, 2016b).

Evidence in line with this has come from studies linking the instantaneous amplitude of alpha oscillations over the occipital channels to visual outcome. For example, Ergenoglu and colleagues (2004) showed that if the amplitude of the ~10Hz ongoing EEG oscillation was higher just before a threshold target, subjects were more likely to miss it (Ergenoglu et al., 2004). This correlation between trial-by-trial variability in alpha amplitude and visual performance has since then received support for various other experimental paradigms (Busch et al., 2009; Mathewson et al., 2011; van Dijk et al., 2008; Wyart & Tallon-Baudry, 2009).

At the neuronal level, these changes in alpha amplitude reflect modulations of excitability. In fact, the instantaneous level of alpha oscillatory amplitude was a good index of the excitability of the cortex measured by single pulse transcranial magnetic stimulation (TMS) between (Romei, Rihs, et al., 2008) as well as within subjects: participants were more likely to report a "phosphene" (illusory percept) when the ongoing EEG alpha amplitude was (relatively) lower (Romei, Brodbeck, et al., 2008, replicated by Dugué, Marque, & VanRullen, 2011 and Samaha, Gosseries, & Postle, 2017). The functional role of these spontaneous fluctuations in excitability of the cortex is still unresolved: they could be used for the detection of unpredictable events in the visual field through a spatial scanning mechanism (Romei, Brodbeck, et al., 2008).

In fact, the amplitude of alpha oscillations also co-varies with endogenous factors (such as attention), as uncovered using the classical Posner paradigm: a central cue (usually an arrow) is used to induce the deployment of covert attention (i.e. without eye movement) towards a visual hemi-field in preparation for the upcoming peripheral target (Posner, 1980). Using this paradigm, studies have shown that the instantaneous strength of the EEG alpha amplitude follows the deployment of spatial attention reflected in two complementary effects with regards to the location of the target: a decreased alpha amplitude over the contra-lateral sensors (Sauseng et al., 2005; Thut et al., 2006; van Diepen, Cohen, Denys, & Mazaheri,

2015; Yamagishi, Goda, Callan, Anderson, & Kawato, 2005) and/or an increased amplitude over the ipsilateral sensors (Cosmelli et al., 2011; Kelly et al., 2006; Rihs et al., 2007; Worden et al., 2000). Thus the locus of visuo-spatial attention is often reflected in an alpha "amplitude lateralization". Furthermore, this amplitude lateralization is retinotopic (Kelly et al., 2006): the locus of spatial attention can be successfully decoded from the topography of the amplitude of alpha oscillations (Samaha et al., 2016).

Importantly, the Posner paradigm has also been used to highlight the effects of endogenous factors on behaviour. The attended targets (i.e. targets presented in the cued hemi-field) are detected faster (Posner, 1980), and more accurately (Posner, Snyder, & Davidson, 1980) than uncued targets. These effects have been extensively reviewed (e.g. Carrasco, 2011; Petersen & Posner, 2012) and replicated: attention acts as a selective tool to narrow the amount of information and optimize the use of our limited brain resources (Carrasco, 2014).

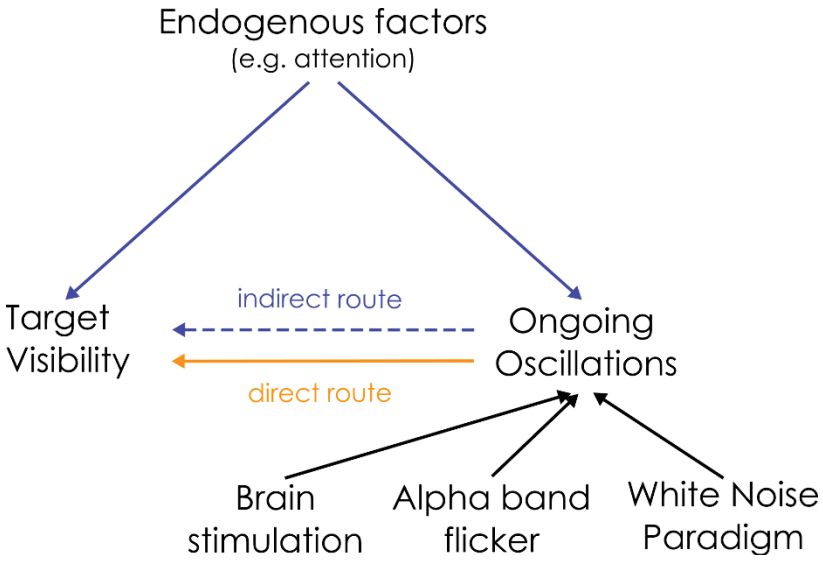

Figure 5-1. Dual route between the ongoing oscillation amplitude and visual perception. The relationship between alpha oscillations and visual perception could be mediated by an indirect influence of endogenous factors on both ongoing oscillations and target detection independently. To disentangle the relative contribution of both routes, the state of the ongoing oscillations can be manipulated by either brain or visual stimulation to then test the effect on perception.

To summarize, both visual detection and the amplitude of alpha oscillations are directly influenced by attention: this create an indirect link between ongoing oscillations and perception (see Figure 5-1). How can we disentangle the relative contribution of the two (direct/indirect) causal routes on sensory neural processing?

It is possible to do so by directly modulating the state of the ongoing oscillations outside of any fluctuations in endogenous factors. By synchronizing the ongoing oscillations with an external driving rhythm, the functional relevance of oscillations can be tested by probing perception as a function of this entrainment (Thut, Schyns, & Gross, 2011). This entrainment was done using two methods: either through transcranial stimulation techniques or through the stimulation of sensory pathways (Thut, Schyns, et al., 2011). In fact, transcranial stimulation (Helfrich et al., 2014; Thut, Veniero, et al., 2011) as well as visual flickering stimulus (Halbleib et al., 2012; Spaak et al., 2014) have been shown to entrain ongoing oscillations even for several cycles after the stimulation is stopped. Crucially, these entrained oscillations have been found to causally influence visual target detection and discrimination.

Using rhythmic TMS, Romei and colleagues (2010) applied a train of pulses over the occipital and parietal cortex in the theta, alpha or the beta band. They found a frequency specific modulation of performance by the entrained alpha rhythm (compared to the other two frequency bands): the detection of contralateral (near-perceptual threshold) visual targets was decreased concurrently with increased alpha amplitude by the rTMS (Romei et al., 2010). In another study, Kanai and colleagues found that stimulation at alpha frequency (versus stimulation frequency) using tACS was the most effective in inducing an illusory perception (i.e. phosphenes) in the dark (Kanai, Chaieb, Antal, Walsh, & Paulus, 2008). In addition, Helfrich and colleagues (2014) showed that tACS successfully entrained the alpha oscillations over the parietal occipital activity. This entrained oscillations successfully modulated behaviour (Helfrich et al., 2014).

Using four flashes of light to entrain the oscillations, de Graaf and colleagues (2013) showed that entrainment of the alpha oscillations (compared with other frequencies) resulted in a specific impairment of the usual cueing benefit (de Graaf et al., 2013). In separate task, they also found that the discrimination of a target presented after the entraining rhythm was modulated during 3 cycles of the alpha rhythm (de Graaf et al., 2013). This effect was also replicated by Spaak and colleagues (2014) who showed that this rhythmic modulation of performance was supported by a neural entrainment of alpha oscillations.

Thus, these non-invasive brain stimulations have been successfully used to study the direct causal influence of ongoing oscillations on sensory processing. Nevertheless, they require an *a priori* hypothesis about the stimulation frequency (Dugué & VanRullen, 2017) : the above mentioned studies used 10 Hz (Kanai et al., 2008; Romei et al., 2010; Spaak et al.,

2014) and 10.6 HZ (de Graaf et al., 2013) respectively. There is, however, a difficulty in choosing the exact stimulation frequency: each subject has their own individual alpha peak frequency which varies as a function of the task demands (Haegens et al., 2014). Ideally, the stimulation would be tailored to the individual subject's alpha peak frequency to create a 1:1 frequency locking, a condition "ideal" for entrainment (Thut, Schyns, et al., 2011).

In this study, we use an alternative method which overcomes this limitation of choosing a frequency of stimulation. By means of the White Noise Paradigm (Brüers & VanRullen, 2017), identical stimulus can be used to constrain the state of background oscillations in a subject specific manner. White-noise (random luminance) sequences are used as the driving visual stimulus. These have been shown to create perceptual echoes in the brain (VanRullen & Macdonald, 2012). Crucially, their individual peak frequency is highly correlated with the individual alpha peak frequency (VanRullen & Macdonald, 2012).

Therefore these WN sequences allow us to modulate the alpha amplitude independently of any fluctuations in endogenous factors. We use them to investigate whether changes in the state of alpha oscillations can influence target visibility directly, regardless of the attentional state.

5.2.2 Method

In this study, we re-analysed data from a previously published experiment (Brüers & VanRullen, 2017), using an identical paradigm but a different analysis pipeline. For the convenience of the reader, the necessary information is described again here.

5.2.2.1 Participants

Twenty-one volunteers were included in the experiment. The data from one subject had to be discarded entirely because of technical problems during the EEG recording session, thus 21 participants (aged 23 to 39 years old, 10 women) took part in the experiment. All subjects gave written informed consent before taking part, reported normal or corrected to normal vision and no history of epileptic seizures or photosensitivity. This study was carried out in accordance with the relevant ethic committee of the Centre de Recherche Cerveau et Cognition at the University of Toulouse.

5.2.2.2 Stimuli and protocol

The experiment was composed of two sessions of 8 experimental blocks (of 48 trials), lasting about 1 hour each (depending on the duration of the self-administered rests). The stimuli were 6.25s long random luminance (white-noise) sequences presented on a cathode ray monitor positioned 57cm from the subject and a resolution of 640 x 480 pixels and refresh rate of 160Hz. The white-noise sequences had a flat power spectrum up to 80Hz. The stimuli were created using custom script in MATLAB and presented using the Psychophysics Toolbox (Brainard, 1997). The sequences were presented inside an overhead peripheral disk with an angular size of 7°, with its centre at 7° of eccentricity from fixation on a black background. Participants initiated the beginning of each trial and block by a button press. The task was to maintain covert attention to the white noise sequences and report the presence of targets (from 2 to 4 targets per trial) embedded within them by pressing a button. The targets (a lighter disk with a darker surrounding annulus) were presented for 1 frame only on a medium grey background disk. The target always had a mean medium grey luminance, identical to the mean medium grey level of the white-noise sequence. Using a staircase procedure, we manipulated the visibility of targets by changing the contrast between the outer (darker annulus) and inner (lighter disk) parts to reach the contrast at which participants perceived about 50% of the first 100 targets presented (i.e. about 30 trials), keeping the resulting contrast constant for the reminder of the session. The perceptual threshold was computed for each session independently using the *quest* function (Watson & Pelli, 1983). During the second session, we removed the luminance fluctuations around the presentation of the target, as classification images revealed that these fluctuations influenced the visibility of the target (Brüers & VanRullen, 2017). Thus, 156.25ms long fluctuations-free periods were created around each target by setting 14 frames before and 11 frames around the target to the same target background medium grey value. A control experiment (on 6 subjects) revealed that these fluctuations could not be detected by the subjects (Brüers & VanRullen, 2017).

5.2.2.3 Electro-encephalography recording and pre-processing

During the first session, we recorded the electro-encephalography (EEG) using a 64 channels (active BioSemi system, 1024Hz digitizing rate, 3 ocular electrodes) to white-noise sequences. The following pre-processing steps were applied to all subjects using the EEGlab toolbox (Delorme & Makeig, 2004) in Matlab. Once the noisy channels had been rejected and interpolated (if necessary), the data was down-sampled offline to 160 Hz, to match

presentation rate of stimuli and thus facilitate the cross-correlation of the two signals. A notch filter was then applied (between 47 and 53 Hz) to remove power line artifacts. An average-referencing was applied and slow drifts were moved from the data (>1Hz high-pass filter). Data epochs (384) were created around each white-noise sequence (from -0.25 to 6.5s) and the baseline activity was removed (i.e. mean activity from -0.25s to 0 before trial onset). Finally, the data was screened manually for eye movements, blinks and muscular artifacts and whole epochs were rejected as needed: on average 20 /384 trials were rejected per subject.

5.2.2.4 Extracting IRF and reconstructing the EEG

Once pre-processed, the EEG data was cross-correlated with the white-noise sequences (see Figure 5-2) to extract the impulse response functions (IRF, also called VESPA by Lalor, Pearlmutter, Reilly, McDarby, & Foxe, 2006; or "perceptual echoes" by VanRullen & Macdonald, 2012). For each subject, this yielded one IRF for each of 64 EEG channel.

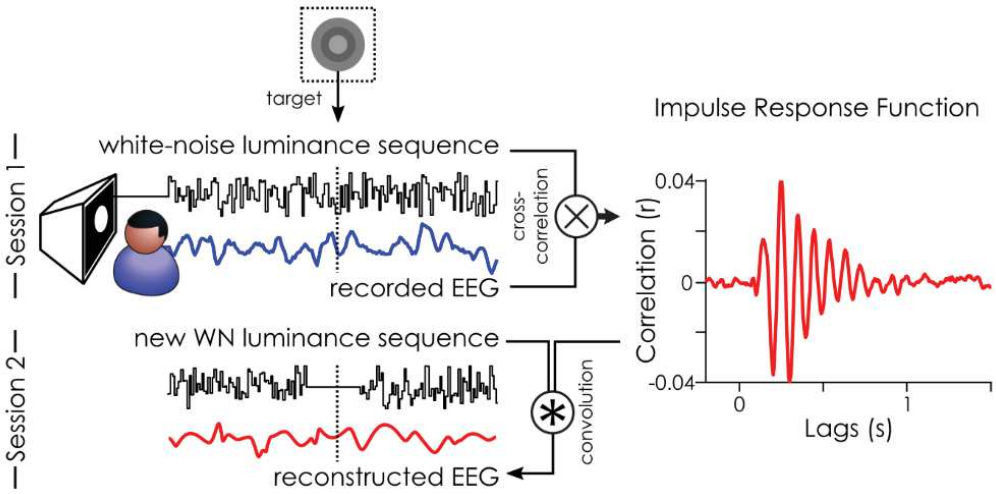

Figure 5-2. Illustration of the White-Noise Paradigm. The impulse response function (IRF) to white-noise sequences can be extracted by cross-correlating the stimuli sequence with the recorded EEG (done in session 1). Here, an example IRF from one subject on the parieto-central channel. This IRF can, in turn, be used to reconstruct the brain activity (reconstructed EEG) to any new white-noise sequence by convolution (done for session 2). Figure reproduced from Brüers & VanRullen (2017).

In session 2, only the behavioural responses (to new white-noise sequences) were collected. All subjects saw the same random luminance sequences (different from session 1) in a randomized order, in order to compare the visibility of targets across subjects. Instead of recording the brain activity ("recorded EEG"), we estimated it ("reconstructed EEG") by

doing a convolution between the IRFs as a model of brain response and the white-noise sequences. The average target luminance (medium grey) was included in the reconstruction. However, it was no different from the surrounding values in the sequence, and thus, no ERP was evoked by the target in the reconstructed EEG (see figure 5 from Brüers & VanRullen, 2017). We created 1.6s long epochs of reconstructed EEG around each target ([-800ms to +794ms]). Any epoch where the target lasted for more than one frame or the staircase had not converged was removed from analysis for all subjects, yielding 821 acceptable trials. The oscillatory characteristics (amplitude and phase) of the reconstructed EEG were extracted using a time frequency transform (using 50 wavelets varying from 3 to 100 Hz in log-spaced frequency steps with 2 to 8 cycles). Note that due to Nyquist frequency limit, our analysis is limited to frequencies below 80 Hz (our signal is sampled at 160Hz).

5.2.2.5 Statistical analysis: computing the amplitude difference between conditions

To test whether the amplitude of the reconstructed EEG had an impact on behaviour, we evaluated the difference between the amplitude of hits and missed target trials. Based on previous studies, we created a region of interest (ROI) to channels where previous effects of amplitude have been reported. The analysis was limited to this ROI composed of 22 channels in the occipital-parietal region (purple dots on Figure 5-3) including all parietal, parieto-occipital and occipital channels. To this end, the decibel difference (dB) was computed for each subject at each channel, frequency and time point as the log transformed ratio between mean amplitude of undetected target trials (misses) and the mean amplitude of detected target trials (hits). We compared this "real" amplitude difference to a "surrogate" distribution. We created 1000 "subject level" surrogates by systematically switching labels between conditions (hits and misses), and computing the amplitude difference between these arbitrary trial groups for each channel, time and frequency point. First, to evaluate the location of the effect in terms of latency and frequency, we averaged the decibel (dB) amplitude difference across all subjects and channels in the ROI. The strength of this "real" amplitude difference was then statistically assessed by applying a nonparametric randomization method allowing the identification of clusters on time-frequency points to control for multiple comparisons (Maris & Oostenveld, 2007). For surrogates, we randomly picked, without replacement, one surrogate for each subject and then aggregated the information for each subjects and ROI channels. This surrogate amplitude difference was then added to the "group level" surrogates distribution and this step repeated 100 000 times. For both the "real" and "group level" surrogates, the time-frequency clusters were extracted by applying a two-sided 2.5%

(arbitrary) threshold based on the distribution of "real" dB values. Finally, we extracted the p-values for each "real" cluster as the proportion of "surrogates" cluster sums in the distribution above the real value.

5.2.2.6 Amplitude dependent performance

In order to quantify the effect of the amplitude difference, we also tested how much variability in performance could be explained at the time and frequency point of the peak amplitude difference. To this end, for each subject and channel, we sorted the instantaneous reconstructed EEG amplitude and split it into 5 (equally spaced) bins across trials. We then extracted the modulation of hit rate for each bin which we corrected by the overall mean performance of the subject. A linear fit was applied to the mean vector across subjects for each channel, and the percent signal modulation was extracted based on the beta. We also report the R-square value for the goodness of fit of the linear line.

5.2.3 Results

To disentangle the relationship between alpha amplitude and visual perception, we used white-noise sequences to constrain the state of background oscillatory activity. Instead of recording the EEG, we reconstructed the background oscillations by doing a convolution between the white-noise sequences and the impulse response functions (see Method) recorded in a separate session. This allowed us to evaluate how alpha amplitude might be related to visual perception independently of any impact of endogenous factors (not present in the IRF). Once the EEG had been reconstructed, we evaluated how the amplitude of this signal might be related to the detection of near-perceptual threshold targets embedded in the white noise sequences.

The target visibility was adjusted using a staircase procedure over the first 100 targets (~30 trials) to reach a 50% average detection rate, the achieved luminance contrast being kept for the remainder of the experiment. During the remainder of the session, the hit rate stayed relatively stable, with subjects reaching an overall mean performance of 45.76% (SD: 10.88%).

We then evaluated whether missed and seen targets trials had different mean amplitudes across trials. Note that we limited our analysis to occipital and parietal channels, an a priori region of interest (see Method). For each 22 channels in this ROI, we computed the amplitude difference (in decibel) between the seen and unseen targets for each subject, channel and time and frequency point. First, we averaged the amplitude difference across subjects and channels to get an overall idea of when amplitude influence target detection. Using a permutation test and cluster correction, we found a significant difference between the mean amplitude of trials were the target was missed versus when the target was detected (significant after cluster correction, $p < 0.00018$) in the alpha band (from 6 Hz to 15 Hz), just around the target presentation (from about -150ms to 200ms). The largest effect is at 12.55HZ and -75ms just before stimulus onset (see Figure 5-3).

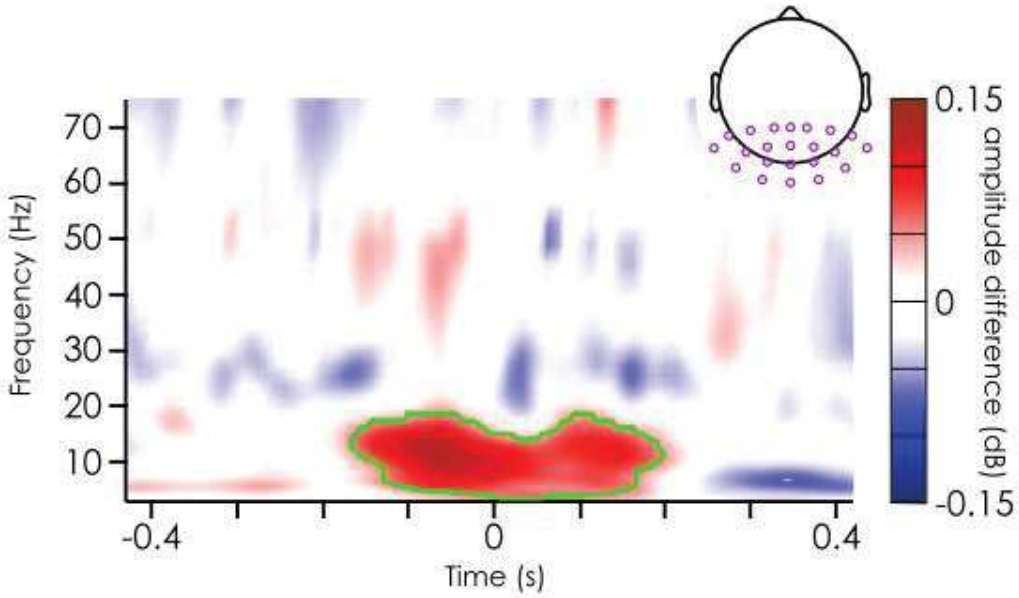

Figure 5-3. Mean amplitude difference (in dB [unseen-seen]) averaged across subjects (N=20) and channels in the ROI (purple dots on the topography). The green outline represents the significant cluster after cluster correction (see Methods section).

Next we sought to quantify the amount of variability in the behavioural response that this amplitude difference could explain. Thus, we computed the normalized amplitude dependent performance (i.e. percent hit rate corrected by the average hit rate across all bins) for each subject and channel in the ROI at the time of maximal amplitude difference (12.55 Hz and -75ms).

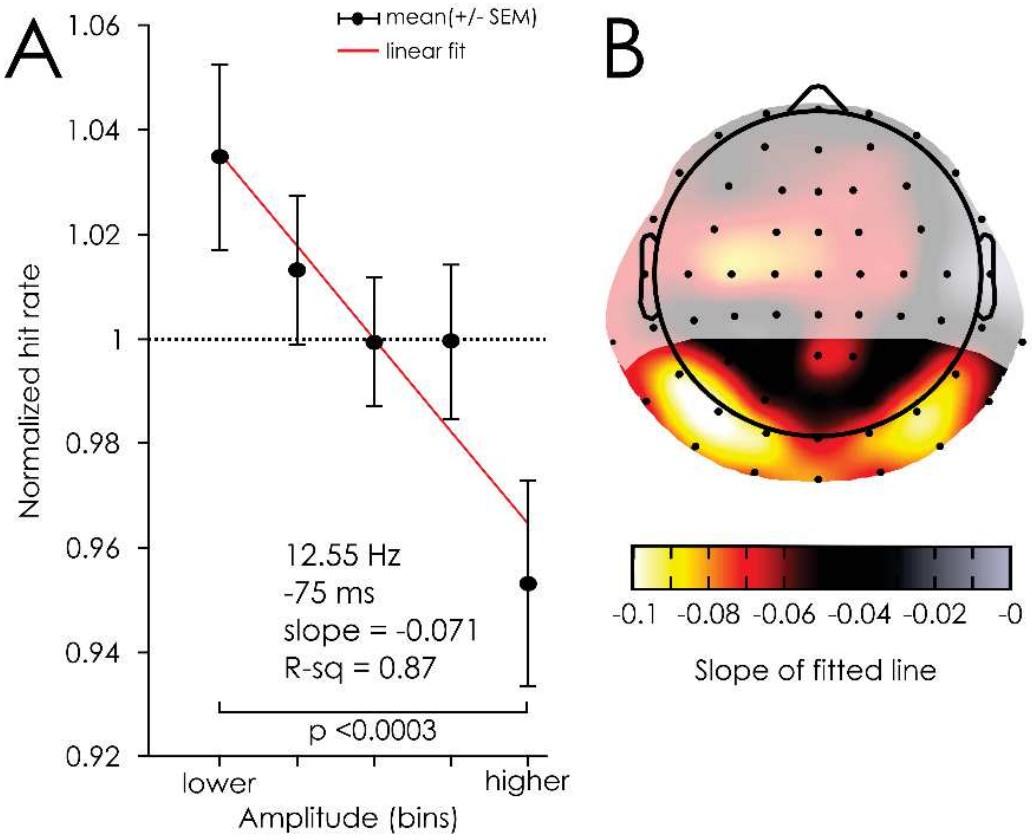

Figure 5-4. Amplitude Dependent Performance. **A**. Mean percent change averaged across all 22 channels in the ROI (black) and 20 subjects for each of 5 amplitude bins. The bars represents the 95% confidence interval across subjects. The red line represents the linear fit. **B.** Amount of performance modulation for each channel. The mean amplitude dependent performance was computed for each channel (averaged across subject) and the coefficient of the linear fit was corrected to extract the modulation of performance. The largest effects are found over the occipital channels. Shaded areas represents channels outside of the ROI.

There is a negative relationship between reconstructed EEG alpha amplitude and target detection. A paired t-test revealed a significant difference in percent change between the bins with the highest and lowest amplitude (see Figure 5-4A) across all channels in the ROI explaining ~7% of the variability in the behaviour. Note that the data was well approximated by a linear trend as the fit had an r-squared of 0.87. When looking back at the single channel level, we can see that this effect is maximal over bilateral occipital channels, where up to 10% of the variability in performance can be explained on channel PO7.

5.2.4 Discussion

In this study, we examined the relationship between the amplitude of ongoing oscillations and visual perception to better understand the correlates of sensory neural processing. We wished to disentangle the contribution of the two (direct/indirect) causal routes linking alpha oscillations and visual detection. To this end, we used the white-noise

paradigm to introduce modulations of alpha amplitude independently of any fluctuations in endogenous factors. This paradigm is based on linear system analysis and uses the properties of the impulse response functions. By presenting white-noise sequences, we are able to constrain the instantaneous state of background oscillations. Using the mathematical properties of IRFs, we can reconstruct, rather than record, the state of these background oscillations. Any effect linking reconstructed EEG and visual perception would necessarily depend on the constrained background oscillations. Thus, this allows us to test the link between alpha amplitude and visual detection without any confounding factor of attention.

Using this method, we find that trial-by-trial variability in reconstructed alpha oscillations amplitude around the target embedded in the white noise sequences was linked with the perceptual outcome. More precisely, the peak amplitude difference between seen and missed targets was found at 12.55Hz over the occipital channels (a priori region of interest) at -75ms just before the target is flashed.

Our results fit with the general findings in the literature which find that pre-stimulus alpha oscillations are correlated with the detection of visual targets at threshold (Ergenoglu et al., 2004; Hanslmayr et al., 2005, 2007; Romei, Brodbeck, et al., 2008; Romei, Rihs, et al., 2008; van Dijk et al., 2008). More generally, it also supports the notion that alpha oscillations play an inhibitory role in visual perception (Jensen & Mazaheri, 2010).

Furthermore, the effects we report, linking alpha amplitude to visual detection, are causal in nature. In the reconstructed EEG, any modulation in alpha amplitude is, by design constrained by the white noise sequences. Any relationship we find with performance is thus necessarily a result of the white noise sequences entraining the background oscillations in a predictable way. That is not to say that endogenous factors played no role in this experiment. It is still likely that fluctuations in endogenous factors had an impact the behavioural response of subjects through changes in attentiveness or arousal level as the experiment unfolded. However, these would be visible only in the recorded EEG. The reconstructed EEG only captures oscillations that are directly phase locked to the background fluctuations in luminance values. It is thus virtually blind to these modulations.

It has been proposed that alpha oscillations play an inhibitory role through the modulation of the neuronal excitability. Spontaneous fluctuations could reflect a spatial scanning mechanisms which would allow the detection of unpredictable events in the visual field (Romei, Brodbeck, et al., 2008). Our data could speak to this interpretation: it is possible

that the background oscillations are present over the visual cortex even in the absence of the white-noise sequences to constrain them. They could in turn play a role in the scanning of visual field for unattended targets, and using continuous stimuli helps highlight this role. In fact, the fact that perceptual echoes could represent a sort of scanning mechanism of the visual cortex (Pitts & McCulloch, 1947) has been suggested before: two echoes showed a systematic phase difference between ipsi- and contra-lateral cortex supporting this mechanism (Lozano-Soldevilla & Vanrullen, 2016).

In a nutshell, we were able to show that alpha oscillations have a causal influence on visual target detection independently from the effects of endogenous factors (such as attention for example). This confirms the inhibitory role of alpha oscillations in visual perception.

5.3 CONCLUSIONS

In this chapter, we showed that the reconstructed EEG alpha amplitude had an influence on the detection of near-perceptual threshold targets embedded within the white noise sequences, outside of any influence of fluctuations in endogenous factors. The white noise paradigm allowed us to introduce fluctuations in alpha amplitudes which were not dependent on attentional modulations. Rather, these fluctuations were a direct consequence of the white noise sequence constraining the brain activity of subjects. Therefore, this study confirms the causal inhibitory role of alpha oscillations in visual perception (Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007; Mathewson et al., 2011), which was shown using other brain stimulation methods such as TMS (e.g. Romei et al., 2010) and visual stimulation (e.g. de Graaf, Koivisto, Jacobs, & Sack, 2014). The relative advantage of the white noise paradigm over traditional methods of visual stimulation is that it does not require choosing one specific stimulation frequency, but the subject specific alpha entrainment results from a broad-band visual stimulation. One further advantage of using non-periodic stimulation is that the oscillatory pattern seen in the "late" part (i.e. perceptual echo) of the impulse response function is necessarily constrained by the sequence and not an artefact of the stimulation (VanRullen & Macdonald, 2012). This is often a problem in flicker studies as it has been shown that the steady-state response to visual flicker could be explained by a superposition of ERPs (Capilla, Pazo-Alvarez, Darriba, Campo, & Gross, 2011). This does raise the question about what the relationship is between the signals that are constrained by the white noise sequence (namely the irrelevant S3 and relevant S4 "phase locked signals")

and the signals that can be constrained using more traditional methods of stimulations. The signals that are constrained by the White Noise Sequence cannot be, if we follow the taxonomy presented in chapter 2, signals involved in endogenous factors. Throughout the course of the experiment, the effects of these top-down modulations would not be systematically phase locked to the stimulus and thus these responses are completely absent from our signal (although they may still influence perception). Yet, we find that the alpha amplitude can determine fluctuations in visual performance. This suggests that, alongside the endogenous oscillations in the alpha band, another type of functionally different oscillations play a role in visual perception. Moreover, these oscillations can be constrained by the White Noise Stimulation. We will speculate further about the possible functional roles of these oscillations in the general discussion.

Chapter 6. UNIVERSAL FORWARD MODEL OF PERCEPTION

In the previous two chapters, we have seen that the oscillations constrained by the sequences have an impact on perception. Both the phase and the amplitude of reconstructed EEG causally influence target perception, and could explain (respectively) 11% and 7% of the variability in the perception of the near perceptual threshold target. In the next two chapters we take a different approach and ask whether this relationship between oscillation and perception can be generalized to new targets and new subjects.

6.1 EXTRACTING THE UNIVERSALS: BUILDING A UNIVERSAL FORWARD MODEL OF VISUAL PERCEPTION

As mentioned in chapter 2, we wish to build a model (or pattern) of the brain response to "seen" and "missed" targets that will be used to predict whether new (independent) targets will be perceived. Here, instead of using the subject's behavioural responses we will use their brain features. In chapter 2, we saw that a large part of the effects of white noise sequences in driving perception could be explained by a shared influence of the white noise sequences across subjects: when a target was presented in a given WN sequence, it was very likely that most subjects would either see or miss it (i.e. high between subjects agreement). Therefore, in this chapter we will learn the "universal" (i.e. subject independent) features of brain activity, which influence perception.

We reasoned that if the WN sequences drive perception in the same way across subjects, this would be reflected at the neuronal level by a shared impact of the WN sequences on the brain responses. Therefore, we decided to extract this subject independent activity and use it to our advantage. This shared influence of WN sequences on brain activity will be reflected in the average impulse response function, which we call the "universal" IRF. Instead of using the subject-specific IRF as a model for the relationship between brain and WN sequence, we extracted the universal IRF: the model of how, on average across all subjects, the luminance information in the sequence was processed. Arguably, if the IRFs reflect a key mechanism of brain function then there should be some similarities across subjects in these IRFs that could be exploited. In fact this similarity between subjects in the visual evoked potential has often been studied: the grand-average ERPs are often extracted to examine the temporal profile of the cognitive function under study across subjects (e.g. Lalor, Kelly, &

Foxe, 2012). It allows, at the cost of removing the inter-subject variability, the benefit of highlighting the similarities between subjects and removing some of the noise (Luck, 2005b).

First, we evaluate if there is any information in the universal reconstructed EEG by testing whether its amplitude and its phase could be linked to target detection across the whole group (using the same approaches described in chapters 4 and 5). Second, we build a forward model, which allows us to transform the raw, random luminance values presented around the target, to an output about the visibility of the target. If we take a new WN sequence, we will be able to predict the likelihood that the embedded target will be perceived or not by anyone, anywhere.

6.2 METHOD

The same participants, stimuli procedure and EEG recording method were used as in the main method (see General Method in Chapter 2).

We used the impulse response functions (IRFs) as a model of the brain's response to the WN sequences presented in session 2 rather than recording it. The IRF were extracted for each subject and channel by cross-correlation of EEG and WN sequences. The universal IRF was computed by averaging the single subject IRF across subjects (see Figure 6-2) thus yielding 64 universal IRFs (one for each channel). The universal reconstructed EEG to the session 2's WN sequences was then determined by convolution of the stimulus and universal IRF. It is this universal reconstructed EEG that was used for all subjects.

6.2.1 Time and frequency transforms

First we tested whether the universal reconstructed EEG contained information about target detection. For the analysis of the influence of phase and amplitude on perception at the whole group level (i.e. on 20 subjects), we extracted the time frequency information present in the universal reconstructed EEG by using a wavelet transform (46 log-spaced frequencies extracted from 3 to 75 Hz with 2 to 8 cycles).

For the forward model, using the wavelet transform would have yielded too many dimensions. Moreover, there would be a lot of redundant information at the different neighbouring frequencies. We therefore used a filter and Hilbert approach for both predictors (i.e. phase and amplitude based). First, the whole trials (i.e. 6.25s) data were filtered (FIR filter, *filtfilt* function in MATLAB) in one of 5 frequency bands (delta: 2-4 Hz, theta: 4-8 Hz, alpha: 7-14 Hz, beta: 14-28 Hz, gamma: 30-60 Hz). Then, we created 1.6s epochs around

each target (i.e. -800ms to +793.75ms), and only the acceptable trials were kept (trials which had been presented for 1 frame only for all 20 subjects). Finally, we applied the Hilbert transform.

6.2.2 Measuring the effect of amplitude & phase on visual perception

For the analysis of the difference in universal reconstructed amplitude between the detected and missed targets, we used a similar approach to what is described in chapter 5 (see 5.2.2.5). The amplitude difference was extracted using a decibel difference between the misses and hits for each channel, subject and time point and compare to a surrogate distribution with the means of a cluster correction (Maris & Oostenveld, 2007). The only two differences was that we used the universal reconstructed EEG and that we computed the amplitude difference across all electrodes rather than on the pre-selected ROI we used in chapter 5, as we wanted to construct the predictor on all electrodes as well.

For the analysis of the phase difference between the hits and misses, we used the same approach as described in chapter 4 (see Method section – measuring phase differences). The phase difference was estimated by means of the phase opposition sum (POS) and compared against the distribution of 1000 surrogates (VanRullen, 2016a).

In both cases, we extracted the significance of the effect in time and frequency, so the data was averaged across channels and subjects. These two analyses were used to verify whether there was any information about the detection of targets to be extracted from the amplitude and phase of the universal reconstructed EEG.

6.2.3 Constructing the classifiers

We built two classifiers based on the universal reconstructed EEG to predict whether a given target embedded in a new white noise sequence would be more likely to be seen (i.e. predict "good" perception) or less likely to be seen (i.e. predict "bad" perception). One classifier used the information present in the amplitude and the other one the information present in the phase.

For both (i.e. the phase-based and amplitude-based classifiers), we used a cross-validation strategy on both subjects and trials: the classifiers were constructed on 19 subjects with 90% of the targets, and tested on 10% of remaining targets for the left-out subject. In total we carried out the cross-validation across trials 10 times per subjects (i.e. 10-fold, thus 100 predictors created per left out subject). Note that to ensure that the algorithm was really

blind to the left out subject's data, we always computed the universal IRF on 19 subjects only and recomputed the reconstructed EEG and time frequency transform at each iteration. The classifiers were constructed in the same way for each frequency, left out subject, time point and channels.

6.2.3.1 Amplitude-based classifiers

The amplitude-based classifier was constructed in the following way (see Figure 6-1). First, we extracted the amplitude of the Hilbert transform by taking the absolute value of the complex across the pooled data from the 19 subjects. Then, we learned the relationship between amplitude and performance for each electrode and time point of interest (64 by 256). To do this, we sorted the amplitude into 5 bins (corresponding to the 20^{th}, 40^{th}, 60^{th}, 80^{th} and 100^{th} percentile). The hit rate was computed for each bin and normalized by the mean hit rate across all bins. We applied a linear fit using the *regress* function to extract the features of the classifier (see Figure 6-1A). For each channel and time point, we learned the sign of the coefficient (i.e. whether the low amplitude was linked with a good or a bad detection) and the coefficient of the slope (reflecting the amount of performance explained by the linear fit, see Figure 6-1A). The sign was used to invert the amplitude of the test trials so that all electrodes and time points were re-aligned and had the same relation between what amplitude led to "good" (positive predictor amplitude) and "bad" (negative predictor) performance (see Figure 6-1). This was used to increase the signal to noise ratio: aggregating information would smooth out any spurious effect from noisy trials or channels. The coefficient was used as a measure of how much information that particular electrode and time point explained. Coefficients were transformed into a weight using a sigmoid function ($\mu = 0.04$, $\sigma = 0.005$) so that a coefficient of 0.03 would result in a weight of 0.02 and a coefficient of 0.05 would give a weight of 0.98. This was defined based on the distribution of the coefficients computed on 20 subjects and used to weigh the coefficients so that only 5% of coefficients (i.e. all coefficients above 0.05) had a significant impact in the decision of the classifier. Finally, we also extracted the mean amplitude for the lowest and highest bin for that fold. This was used to normalize the test trials' amplitude so that their amplitude would be in the same scale as the train trials. Once these features had been extracted, we applied the classifier to the test trials (see Figure 6-1B). The aim was to extract a "predictor amplitude" for the left out trials to predict whether the target had been seen or missed by the left out subject. Once the amplitude of the left out trials had been normalized, invested and weighted, we computed the average

amplitude across all electrodes and time points and used this as the "predicted amplitude", which, as a result of the re-aligning of the amplitude by the sign inversion, was directly interpretable as the prediction value (see Figure 6-1C). If the predicted amplitude was positive, we predicted a hit (i.e. "good" perception) while if the predicted amplitude was negative, we predicted a miss (i.e. "bad" perception). We applied a threshold so that we did not make any predictions for trials which had a "predicted amplitude" close to 0. To do this, for each predictor, we extracted the values corresponding to the 40^{th} and 60^{th} percentile of the distribution of the predicted phase for the training trials. These were then used as edges for the test trials: any predicted amplitude value below this was counted as a "bad" trial for perception and any predicted amplitude value above this threshold led to a "good" perception prediction (se Figure 6-1C).

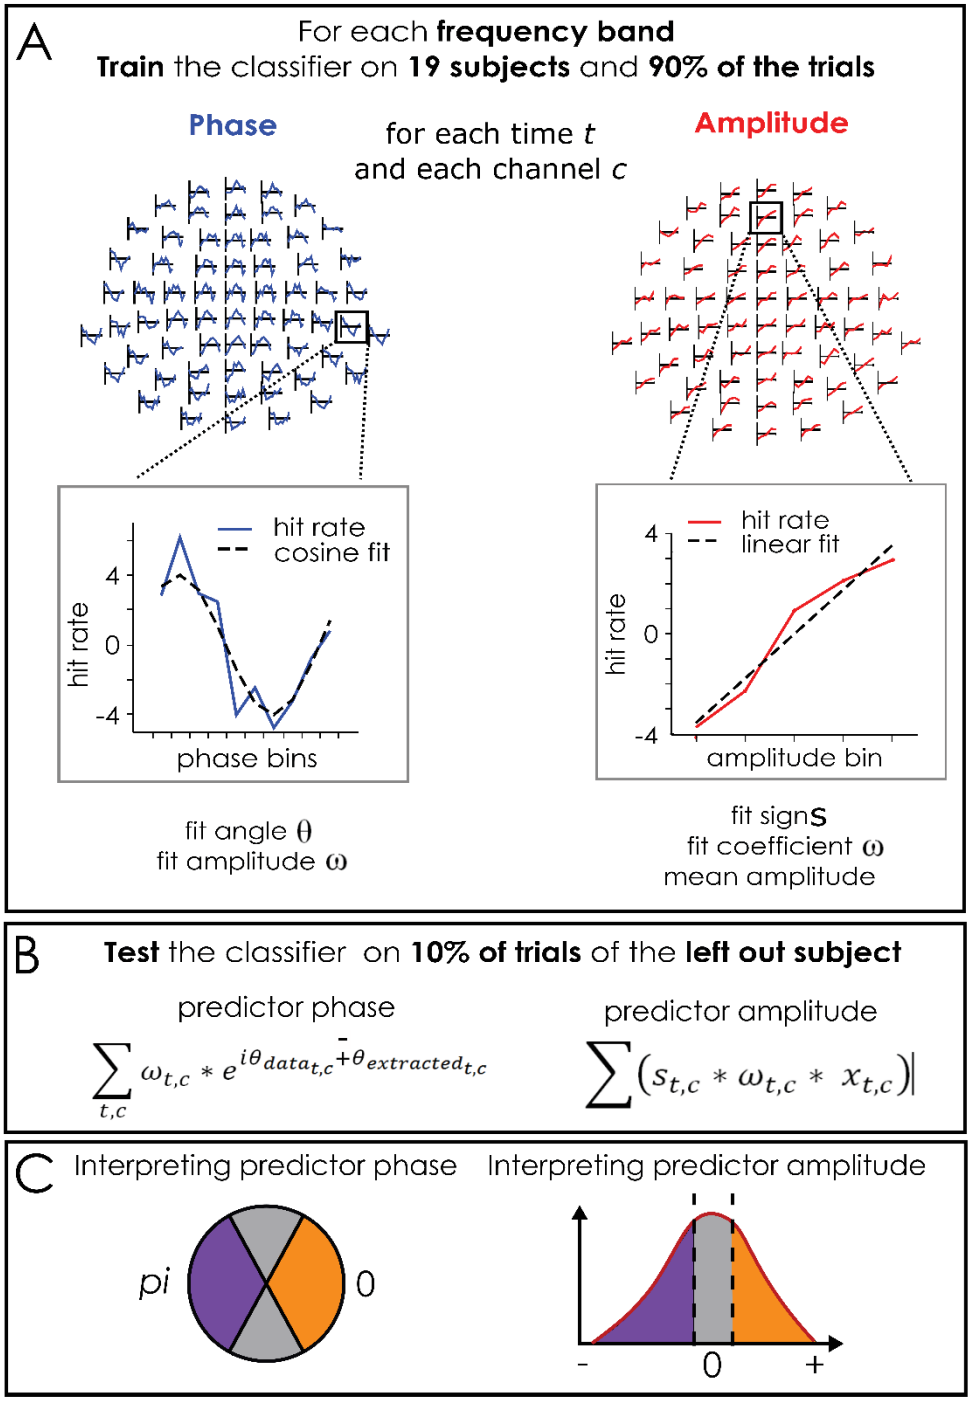

Figure 6-1. Illustration of the construction of the phase and amplitude based classifiers. A. The classifier was trained on 19 subjects and 90% of data. For each of 5 frequency bands, 256 time points and 64 channels, we extracted the phase- and amplitude-dependent performance, which was then fitted with a cosine or a linear fit (for the phase and amplitude respectively). We extracted the cosine angle and amplitude (for the phase based classifier) and the sign and slope of the linear fit as well as the mean amplitude of the lowest and highest amplitude bins and the amplitude distribution. B. These features were applied to the remaining 10% of trials for the left out subject as a weighted sum across all channels and time points, which was realigned so that the "good" and "bad" amplitude and phase were the same across channels and time point. C. Consequently, the predicted phase and amplitude

could be directly read out as the prediction of the model. The phase falling within 2 phase bins (i.e. 1.25 radian) from 0 corresponded to a predicted "good" perception, while the predicted phases falling within 2 phase bins from pi corresponded to a predicted "bad" perception (in purple). For the amplitude, the negative amplitudes were ascribed to the predicted "bad" perception (purple) while the positive amplitudes were ascribed to the predicted "good" amplitude. Similarly to the phase classifier, 20 % of the data was not included in the prediction because reflecting noise (amplitudes close to 0).

6.2.3.2 Phase-based classifiers

The phase-based classifier was constructed in the following way. First, we normalized the amplitude of the complex to only keep the phase information. Then, we sorted the trials into 10 bins computed the normalized hit rate for each bin by subtracting the mean hit rate (see Figure 6-1A, left). This was fitted with a cosine by applying an FFT and extracting the amplitude and angle of the 1 cycle component. We learned the angle to re-align the test trials so that the "good" phase and "bad" phase would be the same (0 and pi, respectively) for each channel and time point. The amplitude of the cosine fit was used to extract a weight using a sigmoid function (μ=0.06 and σ=0.005). An amplitude of 0.05 was given a weight of 0.02 and an amplitude of 0.07 was given a weight of 0.98. Here also, these thresholds were extracted based on the distribution of cosine fit amplitudes computed for all trials and subjects. Once extracted, the predictor was applied to the remaining 10% of trials not included in the classifier for the left out subject. We wanted to extract one predictor phase per test trial which would predict whether the target had been seen or not. To increase the signal to noise ratio and improve our prediction, we decided to aggregate information across all channels and time points. The complex values of the test trials were shifted and weighted and then averaged across channels and time points to extract one predictor value. The predictor was built so that the predicted phase could be directly read out as a prediction for the detection of the target: a predictor phase within 2.5133 radian (i.e. 2 phase bins on each side) of *pi* would predict a "bad" performance while a phase within 2.5133 radian (i.e. 2 phase bins on each side) of *0* and *2pi* would predict a "good" performance (see Figure 6-1). We did not make any predictions about the values falling in between (i.e. two bins were rejected), because these could not be easily classified as hits or misses.

6.2.3.3 Evaluating the classifier's performance

To evaluate the performance of the classifier, we computed the "performance modulation" (as described in chapter 2): from the predictions of the models, the number of targets actually seen by the left out subject in the phase and amplitude predicting a "good"

and “bad” perception was computed. From this, the performance modulation was extracted as the difference between the hit rates for “good” minus “bad” perception predictions. This performance modulation was computed for each classifier constructed, for the phase and amplitude modulation, and then averaged across 10-fold repetitions. A t-test was applied to test whether the distribution of performance modulation across subjects was significantly different from 0. Furthermore, we have seen in chapter 2 that the detection of a large part of the targets presented in this experiment could not be explained based on the WN sequences. Between subjects, the perception of these targets was random: we cannot hope to make any prediction about their detectability using a universal forward model based on the luminance of the WN sequence as the WN sequence does not influence their perception. Therefore, we also computed the amount of performance modulation that could be expected from a subset of high-agreement targets. The targets with high agreement rate were recomputed for each 19 subject combinations, and defined as the targets for which at least 65% of subjects agreed.

6.3 RESULTS

In this experiment, we wanted to see whether “universal” features in oscillatory activity influenced the perception of the near-perceptual threshold target embedded within the white noise sequences. We did this using two approaches. First, we used “classical” analyses methods to link the trial-by-trial amplitude and phase of the universal reconstructed EEG to the visibility of the target, then we constructed classifiers so that, for any new target embedded in a white noise sequence, we would be able to predict whether it would be seen or not for anyone anywhere (i.e. universally).

6.3.1 Extracting the subject-independent IRFs

First, we extracted a model of the “universal” brain response to visual stimulation - the universal Impulse Response Functions (see Figure 6-2) by averaging the individual IRF across subjects. The average IRF have a similar topography to the “single subject” IRF: a strong amplitude in the alpha range. However, they also show a strong theta amplitude over the fronto-central channels (see Figure 6-2). Moreover, it seems that there is power in the beta band over the occipital channels (which was not apparent in the individual subjects’ IRF). This however is only a qualitative appraisal of the effects. We will see in the classifier whether the information present in the IRF actually influences perception.

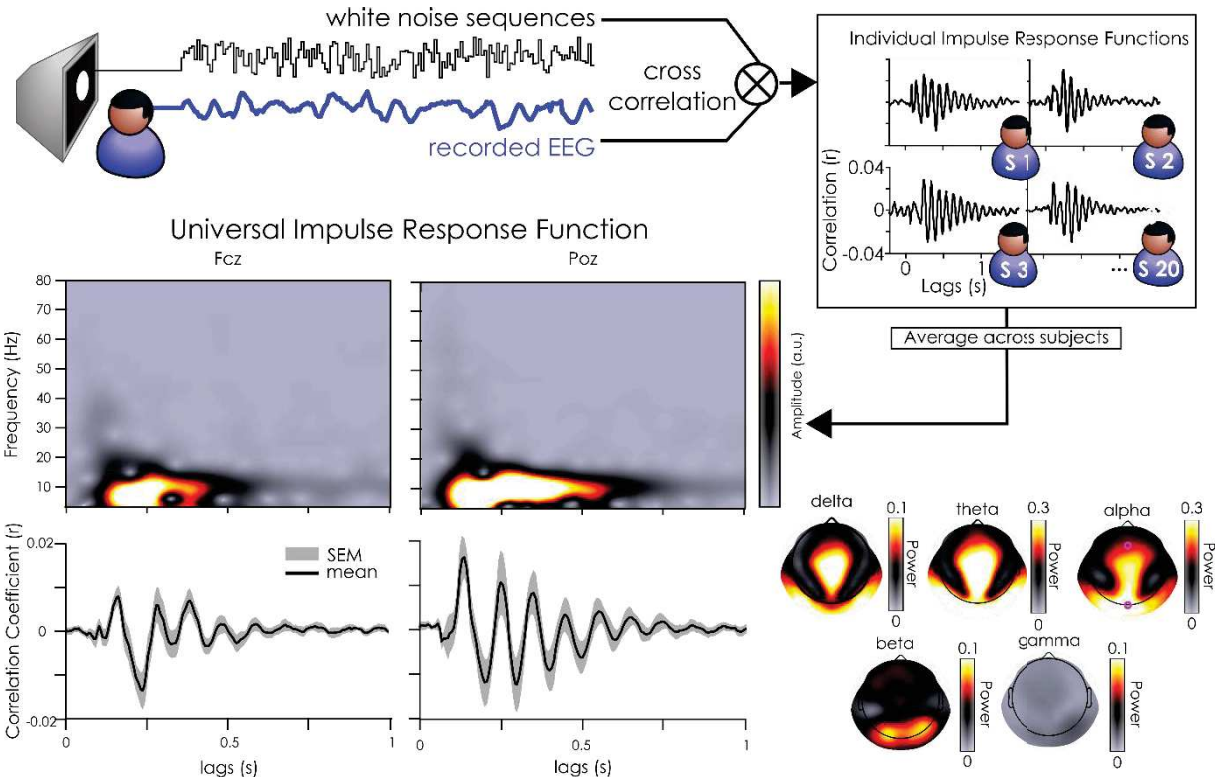

Figure 6-2. Illustration of the extraction of the universal impulse response function (IRF). The individual subject's IRF are extracted by doing a cross-correlation between the WN sequences and the concurrently recorded EEG signal for each subject. Top right, the IRF from 4 subjects on the parietal-occipital central channel are represented. The universal IRF is then extracted for each channel by averaging the individual responses across subjects. Two universal IRF (computed over 20 subjects) are represented here measured over the parieto-central (right) and the fronto-central (left) channel, with the corresponding time and frequency transform. The topographies on the right represent the power in each of the 5 frequency bands used in this experiment measured from the "late" part of the IRF i.e. from 250ms to the end. As can be seen, there is a strong (frontal) theta and (occipital) alpha component to these IRF.

6.3.2 Universal amplitude and phase have an impact on performance

First, we tested whether the phase and the amplitude of the universal reconstructed EEG contained any information that could be extracted at the group level to explain the detection of the target. This approach is identical to that taken in chapter 4 and 5 to relate the features of the subject specific EEG to perception. The only difference is that, instead of using the subject specific IRF, we use the universal IRF to reconstruct the EEG.

Using the universal reconstructed EEG we computed the amplitude difference between detected and missed target trials for each frequency and time point around the presentation of the target (at time 0ms).

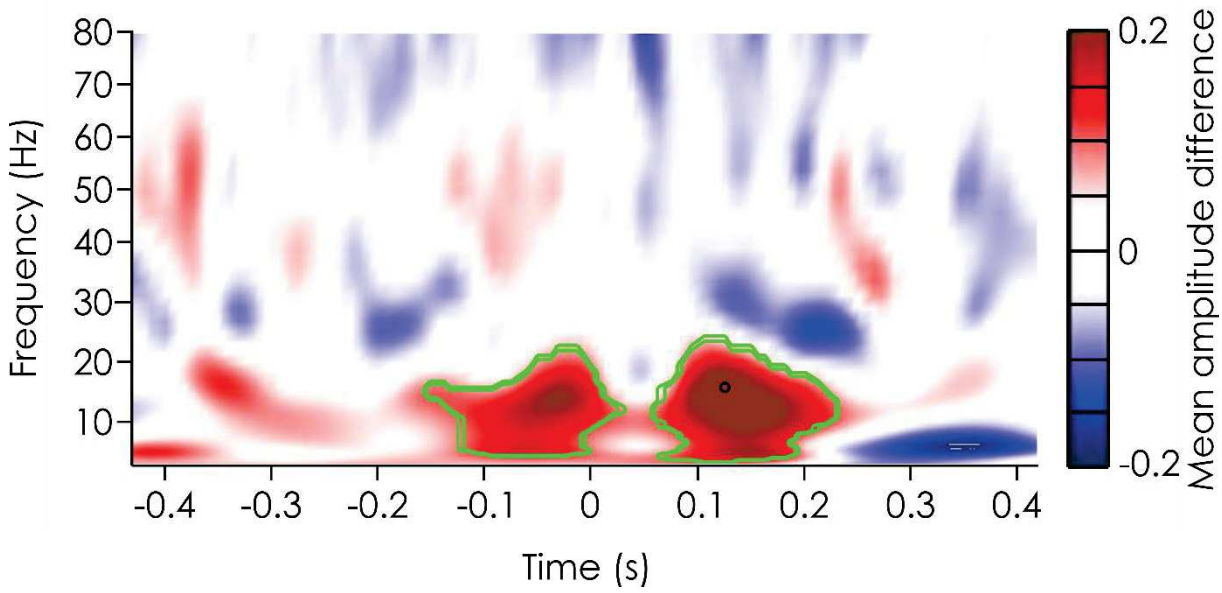

Figure 6-3. Amplitude difference (dB) between the mean amplitude of misses minus the mean amplitude of hits. The decibel differences were averaged across electrodes and subjects and cluster corrected across time and frequency points. The alpha band oscillation around the presentation of the target (at time 0ms) is significantly correlated with target detection. The green lines represent the significant clusters and the black dot represents the largest amplitude difference at 15.55Hz and 125ms post stimulus. The approach used here is almost identical to that used in chapter 5 to relate the amplitude of the (single subject) reconstructed EEG to target detection. The only two differences are that we used the universal IRF to reconstruct the EEG and that we computed the average across all electrodes instead of only the parieto-occipital channels in chapter 5. This explains the similarity of the results.

There was a significant difference between the mean power of missed versus detected targets in the alpha band oscillations around the presentation of the target (see Figure 6-3). Two clusters survived correction for multiple comparisons: cluster 1 in the post-stimulus time window had a p-value of $p = 0.0116$, while cluster 2 in the pre-stimulus time window had a p-value of 0.0347. In terms of frequency, although this effect has a peak in the alpha band, it is wide spread: from 5 to 20 Hz. The largest decibel difference between the mean power differences of misses versus hits was found at 15.5Hz and 125ms after the target was presented. Even though the effects seem to peak in the post-stimulus time window, the alpha power determines perception in the pre-stimulus time window, from -170ms before stimulus onset up to 225ms after the target. Note that this is very similar to the effects found in chapter 5, using the single subject IRF instead of the universal EEG. This is not surprising, as we used a similar approach in both cases. Moreover, since the universal IRF is the average of the single subject IRF, these two measures of brain activity while be correlated, although one will

emphasize the between subject similarities, while the other one will emphasize the between subject differences.

Next, we tested whether the universal reconstructed EEG phase was different depending on perceptual outcome.

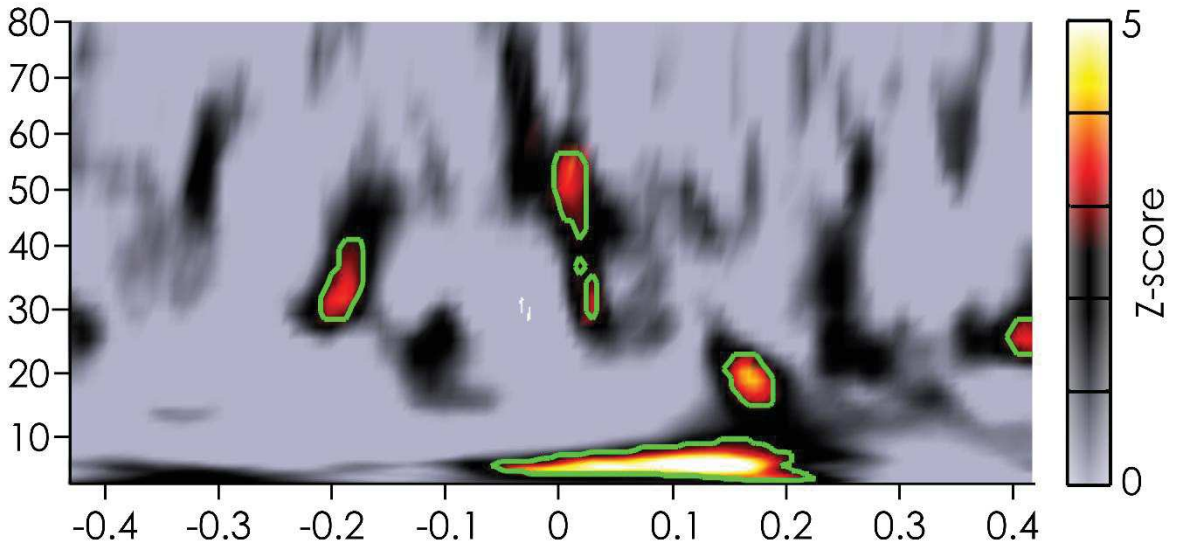

Figure 6-4. Z-scores for the phase opposition sum between the detected and missed trials for the universal reconstructed EEG features. The z-scores were averaged across subjects and channels and the green outline shows the significant time points after fdr correction (alpha level of 0.05). The approach used here is very similar to that used in chapter 4 to relate the phase of the (single subject) reconstructed EEG to target detection. We used the same approach and analysis, only the universal IRF was used instead.

To this end, we computed the phase difference (using the phase opposition sum) between the detected and missed targets, and compared its strength against a surrogate distribution (VanRullen, 2016a). This phase difference was averaged across electrodes and subjects to extract the time and frequency location of the effect. We can see in Figure 6-4 that the phase of the theta band oscillation around the presentation of the target (from -56.25ms to 225ms and from 3.46 Hz to 9.42Hz) is different between the target stimuli that were seen versus the targets that were missed. There are also a few significant (but much weaker) effects in the high beta band and in the gamma band. Note the similarity between this figure and the figure reported in chapter 4: this can be explained by the fact that identical approaches were used for both. The only difference was that we used the universal IRF instead of the single subject IRF to reconstruct the EEG here.

These two analyses confirm that there is information about the general perception of targets to be extracted in the amplitude and in the phase of the universal reconstructed EEG.

Next, we tested whether we could build universal forward models able to link the new WN sequences luminance values to perceptual outcome through phase and amplitude patterns.

In the next analyses, we built two forward models to predict the detection of targets with the features (either the phase or the amplitude) of the universal reconstructed EEG.

6.3.3 Amplitude based predictor

We started by building the amplitude-based predictor. To do this, we systematically learned the relation between moment by moment variability in the instantaneous amplitude from the reconstructed universal IRF and the target detection on 19 subjects and 90% of the targets. Then, we made a prediction about the remaining 10% of targets: could we predict what the perception of the left out subject was, based on the universal brain activity? The ability to predict was evaluated in terms of percentage of performance modulation, that is, the difference in hit rates between the amplitudes leading to "good" and "bad" perception predictions.

First, we applied the predictor on all trials (see Figure 6-5). We found no significant differences in any of the frequency bands: delta band (mean performance modulation = 0.4746%, $p = 0.1044$), theta band (0.0091%, $p = 0.9738$), alpha band (0.1475%, $p = 0.6399$), beta band (-0.0415%, $p = 0.8830$), gamma (0.1348%, $p = 0.6704$).

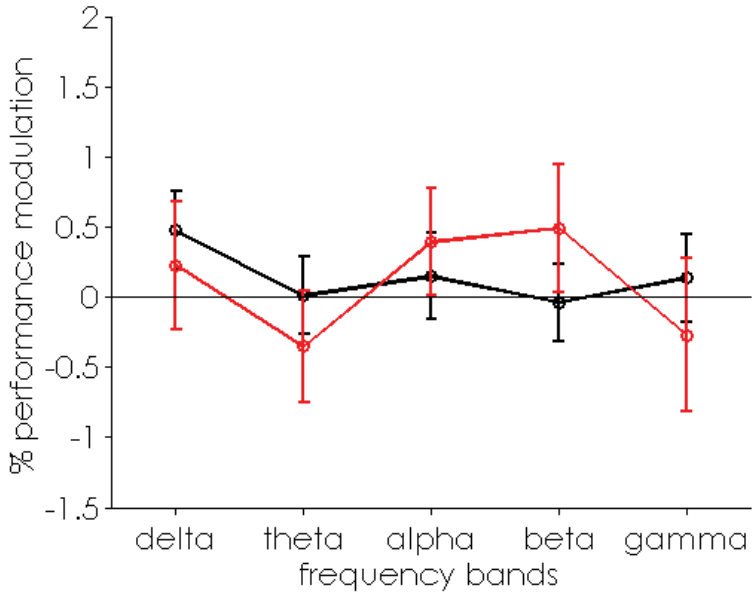

Figure 6-5. Performance modulation for the amplitude based predictor for all trials (in black) and trials with high agreement (in red). The circles represent the mean and the error bars represent the standard error of the mean.

As mentioned in chapter 2, a large part of the targets in the sequences don't have their visibility that depends on the white noise sequences. Therefore, we will never be able to explain any amount of performance modulation for any of these targets. We decided to compute the performance modulation only on targets which had a high agreement rate between subjects (see Figure 6-5, red line). Here again, however, we did not find any evidence for a modulation of performance. There were no significant effect in the delta (mean performance modulation = 0.2257%, p = 0.6251), theta (-0.3535 %, p=0.3871), alpha (0.3910%, p = 0.2977), beta (0.4899%, p = 0.2977), gamma (-0.2701%, p = 0.6288) band.

6.3.4 Phase based predictor

Next, we did the same approach, but using the phase instead of the amplitude. Looking at the percentage of performance modulation across all repetitions and subjects, we find a significant effect in the phase of the delta, theta, alpha and gamma band (see Figure 6-6A): the phase of the theta band explained 6.72% of performance modulation ($t(19)=6.36$, $p < 4*10^{-6}$, 95% confidence interval = 1.387 to 4.11), followed by the alpha band with a mean across subjects of 3.57% ($t(19)=3.74$, $p < 0.0014$, 95% confidence interval = 1.572 to 5.57), the delta band with a mean 2.75% of performance modulation ($t(19)=4.22$, $p < 0.00046$, 95%

confidence interval = 1.387 to 4.11) and the gamma band, which explained only 1.42% of performance modulation (t(19)=2.5251, p < 0.02, 95% confidence interval = 0.24 to 2.59). The phase of the beta band oscillations did not modulate performance significantly (mean = 0.319%, p=0.727).

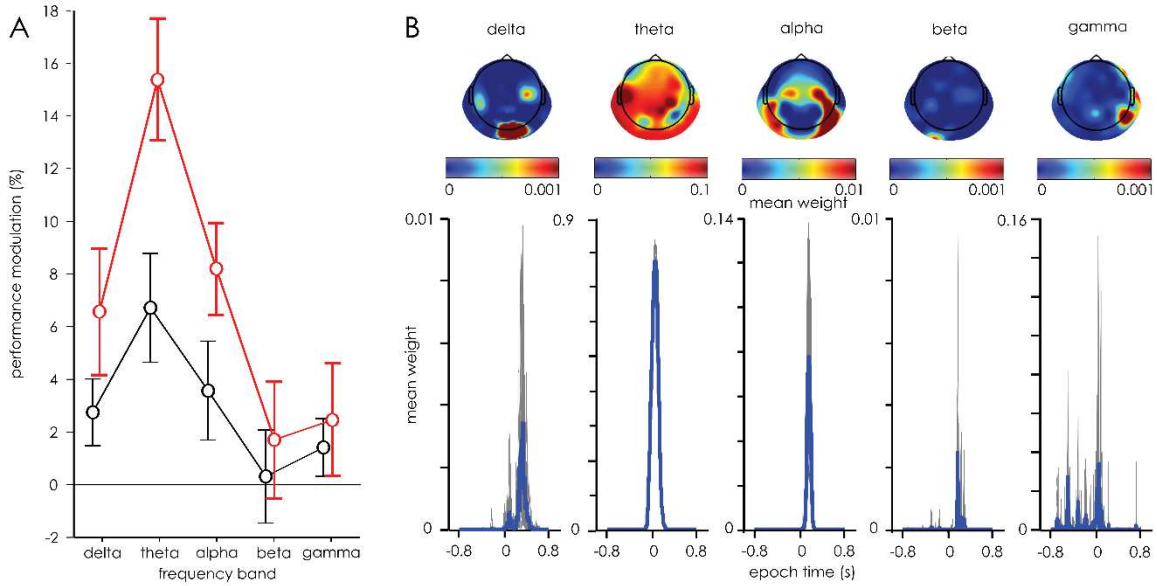

Figure 6-6.Performance modulation for the phase based classifier and average weights of the classifier. **A**. Results of the modulation of performance based on the universal reconstructed phase for all trials (in black) and for the subset of trials with high agreement (in red, see chapter 2). **B.** Average weights (across subjects and repetitions) of the classifier using phase for each frequency bands. At the top, the weights are averaged across time. Note the difference in colour scale. At the bottom, the weights are averaged across electrodes: the mean across subjects is in blue, and each predictor per left out subject is in grey. Note the difference in colour bars and axes: the strongest effects are in the theta band, followed by the alpha frequency band, the effects in the other frequencies are smaller.

Using the same approach as described above for the amplitude based predictor, we looked at the performance modulation expected from the high agreement trials. Again, the percentage of performance modulation across all repetitions and subjects was significantly different from 0 for the delta, theta, alpha and gamma band (see Figure 6-6). The largest effect was explained by the phase of the theta band: we could reach up to 15.4% performance modulation on average (t(19) = 13.03, p < $6*10^{-11}$, 95% confidence interval = 12.914-17.86). Next, the best prediction was achieved based on the phase of alpha band with a mean across subjects of 8.2% (t(19) = 9.23, p < 1*10^-8, 95% confidence interval = 6.3386 to 10.05), and the delta band with a mean 6.57% of performance modulation (t(19)5.37, p < 3*10^-5, 95% confidence interval = 4.0137 to 9.1276). The phase of the gamma band oscillation only

explained a very small amount of variability in the detection: about 1.7% ($t(19) = 2.226$, $p < 0.035$, 95% confidence interval = 018664 to 4.7519).

Next we looked at temporal and spatial dimensions of the predictors by plotting the weights created from the amount of phase-dependent performance modulation during the learning phase. The weights closer to one represent electrodes and time points which play a significant role in visual perception. We see that for the theta band, there is a widespread activity, centred roughly on fronto-central electrodes, and found around the presentation of the target. This topography for the theta phase makes sense when we compare it to the topography of the universal IRF theta power (Figure 6-2). For the phase predictor based on the alpha oscillation, the topography reveals that the effects are driven by bilateral effects over the occipital channels, just after target presentation. The weights of the other predictors are very small in comparison to the theta and alpha band.

6.4 DISCUSSION

In this chapter, we built two universal forward models to predict the perception of new targets embedded in the white noise sequences, using the "universal" reconstructed EEG features. We found that the phase (but not the amplitude) of the universal reconstructed EEG could be used to successfully predict the visibility of the targets embedded in the sequences. When taking only the high agreement trials, we could explain up to 15% of the variability in target detection across subjects. This was driven by the universal theta phase over fronto-central channels. In plain words, this means that at the "good" predicted phase for perception, a target would have about 57.5% of chance to be detected by any subject, while it would have only 42.5% chance of being detected at the "bad" predicted phase.

These results confirm that the phase of the universal reconstructed EEG oscillations plays a causal role in visual perception, and that this effect is generalizable across subjects. Since we were able to find significant effects in the analysis of the amplitude difference between seen and missed targets at the group level, it is surprising that we could not find any effect in the amplitude based predictor. Using a cross-validation strategy ensured that the results were generalizable, which might explain why we cannot find an effect of the amplitude of the universal predictor. In fact, the fact that the phase of the oscillation carries more information than that of the power has been suggested before (Montemurro et al., 2008;

Schyns et al., 2011). Although we did not formally define the amount of information present in the phase of the reconstructed EEG, as mentioned above, the cross-validation method ensures that the results can be generalized across subjects and trials. Moreover, the signal modelled by the EEG (i.e. "phase locked relevant signal") has its phase wholly determined by the white noise sequence (by its very definition). Thus, any phase effect found here must be causal in nature.

Moreover, we saw in chapter 2 that the perception of only a specific proportion of the targets could be influenced by the WN sequences, which ultimately created an upper limit to the amount of performance modulation we could explain. We found that the upper limit to the performance modulation we could expect from analysis of the between subject agreement rate was 18.18%. Here, the theta phase alone accounted for 15% performance modulation, that is, over 80% of the theoretical maximum. The delta and alpha band explain another 6.57% and 8.2% of modulation of performance, which might explain even more of the theoretical performance modulation. Yet, these different frequency bands are unlikely to reflect independent sources of information. In fact, the filter width was kept relatively large to keep the number of frequency bands to a minimum while covering a large part of the frequency spectrum. This also ensured that the number of features was kept relatively low to improve classification and avoid over-fitting. Moreover, these wide frequency bands were also used to avoid a decrease in temporal specificity with increased frequency resolution.

6.5 CONCLUSIONS

In a nutshell, we have shown in this chapter a general mechanism by which meaningful features of the reconstructed EEG activity can be learned to extract the relationship between the oscillatory characteristics and perception. In the next chapter we will apply the method to a different approach: we will try to predict perception in a subject specific manner.

Chapter 7. TOWARDS A NEURO-ENCRYPTION SYSTEM: PREDICTING SUBJECT-SPECIFIC PERFORMANCE

7.1 INTRODUCTION

The previous chapter has shown that the subject independent relationship between the white noise stimulation and brain response on the one hand and the brain response and perception on the other hand could be used to build a universal forward model of perception (see Figure 7-1, chapter 6). The perception of any new target embedded within a WN sequence can directly be predicted based on the universal reconstructed EEG phase. In this chapter, we will now test whether the same approach can be taken in a subject dependent manner.

In chapter 2, we saw that the relationship between WN sequences and brain response (i.e. the IRF or perceptual echoes) was characterized by individual differences: each subject had their own phase and amplitude relationship with the WN sequences. Moreover, we saw in chapters 4 and 5 that the phase and amplitude of this brain response was related to the detection of targets embedded within the WN sequences. Bringing these two together, we thus build a classifier which learns the stable relationship between the phase and amplitude of the single subject reconstructed EEG and perception (see Figure 7-1, Train Classif7). The difference here with chapter 6 is that the predicted EEG is reconstructed by using the individual IRFs instead of the universal IRF. Because the IRFs are different between subjects, the reconstructed EEG to the same WN sequences will be slightly different for each subject.

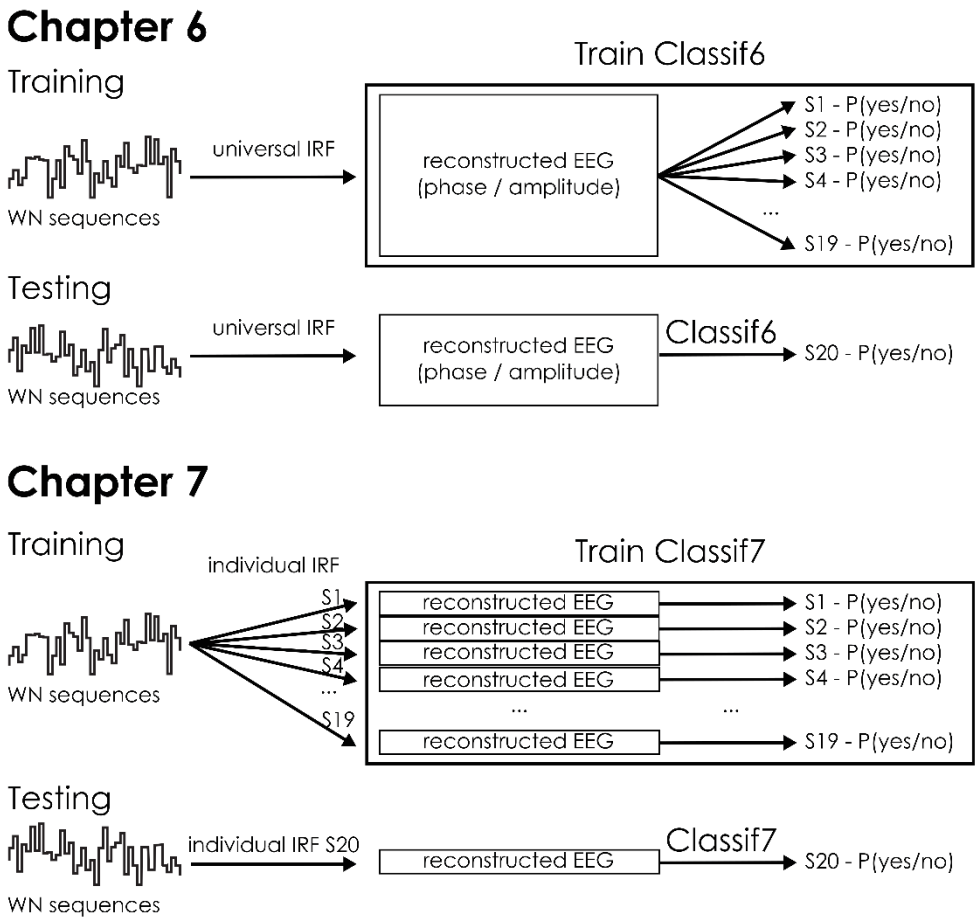

Figure 7-1. Classifier construction and prediction for the universal (chapter 6) and subject dependent (chapter 7) approaches. In chapter 6, the universal IRF is used as a model of the brain's response to WN sequences. Thus the universal reconstructed EEG is used to learn the spectral features (i.e. phase and amplitude) linked with good or bad perception (train classif6). During the testing stage, the classifier learned on 19 subjects and 90% of the data is applied on the universal reconstructed EEG for the 10% of trials left out in order to predict the behaviour of the left out subject. In chapter 7, the individual IRF are used as a model of the brain's response to WN sequences. Thus, the reconstructed EEG is different for each individual subject. Just as previously, the classifier is trained to learn the relationship between reconstructed EEG features and perception across all 19 subjects. During the testing session, classif7 is then applied to the reconstructed EEG extracted from the single subject's IRF and used to predict the perception of the left out subject.

As can be seen from figure 7-1, we decided to learn the relationship between reconstructed EEG and perception at the group level rather than building predictors for each individual subject. In other words, the individual reconstructed EEG of each subject was aggregated across 19 subjects and we learned how these signals related to target detection as described in chapter 6 (see predictor construction). This was done because we hypothesized that the relationship between phase/amplitude and excitability as well as that between

phase/amplitude and perceptual sensitivity was the same for all subjects. So that, for example, one phase can be systematically related to "good" perception and another to "bad" perception, and this would be the same for all subjects. Moreover, this would also decrease the learning time. We hoped that the relationship between perception and spectral features was generalizable from a group of subjects. Thus, we would only have to learn the relationship between reconstructed EEG features and target detection once on a representative group of subjects, and then apply this predictor to all new subjects for whom an IRF has been recorded. The individual differences in the prediction would only come at the "test" phase: here, since the individual IRF are used (instead of the universal IRF in chapter 6), the reconstructed EEG will be different for each subject. Therefore, even if the same classifier is applied to the individual reconstructed EEG of two subjects, the output/prediction of the classifier could differ. For example, for the same WN sequence and classifier, we might predict that subject A would see it, but not subject B. This would not happen with the universal IRF since the same reconstructed EEG is used for all 19 subjects (see Figure 6-2).

7.2 METHOD

The dataset used in this analysis is the same as described in chapter 2 and already analysed in chapter 4 and 5. The subject-specific IRF were recorded and extracted for 20 subjects on 64 channels (as described in chapter 2), and the session 2's brain responses were reconstructed by doing a convolution between IRF and WN sequences.

We then built two classifiers based on the phase and amplitude of the reconstructed EEG to extract the stable relationship between perception and spectral characteristics of the constrained brain activity. To do this, we used a procedure similar to that used in chapter 6, only the single subject IRF were used instead of the universal IRF (see Figure 7-1). We constructed a phase and amplitude based classifier based on the reconstructed activity filtered in 5 frequency bands. Just as before, the performance of the classifiers was evaluated by computing the "performance modulation". In other words, for the subject left out on the remaining 10% of trials, we computed the actual hit rate for the phase/amplitudes that were predicting a "good" performance versus the hit rate for those which were predicting a "bad" performance. We refer to the difference between the two as the performance modulation. We tested whether this was significantly different from 0 by applying a t-test.

Finally, we also computed the phase difference in the reconstructed EEG between subjects. For this, 384 new random WN sequences were created. For each of the 5 frequency bands, the reconstructed EEG of each subject was computed, filtered and a Hilbert transform was applied to extract the complex signal. The angle of the reconstructed EEG at time 0 and for electrode POz was chosen as the variable of interest, as this is the electrode which showed the most difference between subjects in terms of the perceptual echo. We then systematically tested, for each pair of subjects, what the phase difference was for each of 10 phase bins (same phase bins as used in the phase based predictor above).

7.3 RESULTS

In this chapter, we wanted to build a classifier which would predict whether a target embedded within a white noise sequence would be perceived by a new independent subject. For each frequency band, we systematically learned the stable relationship between the reconstructed EEG feature and perception on a subset of subjects (N=19) and trials (90%). We then applied this classifier to the phase or amplitude of the new trials (10%) for the left out subject to see if we could correctly predict whether that subject would see the target or not.

7.3.1 Amplitude based classifiers

The performance of the amplitude based classifier in predicting the perception of new targets for independent subjects can be seen from figure 7-1. When testing all trials, there are no significant effect (see figure 7-1, black line) in the delta (mean performance modulation of -0.70%, p=0.25), theta (0.06, p=0.91), alpha (-0.014%, p=0.98), beta (0.331, p=0.67) or in the gamma (0.339%, p=0.57) band.

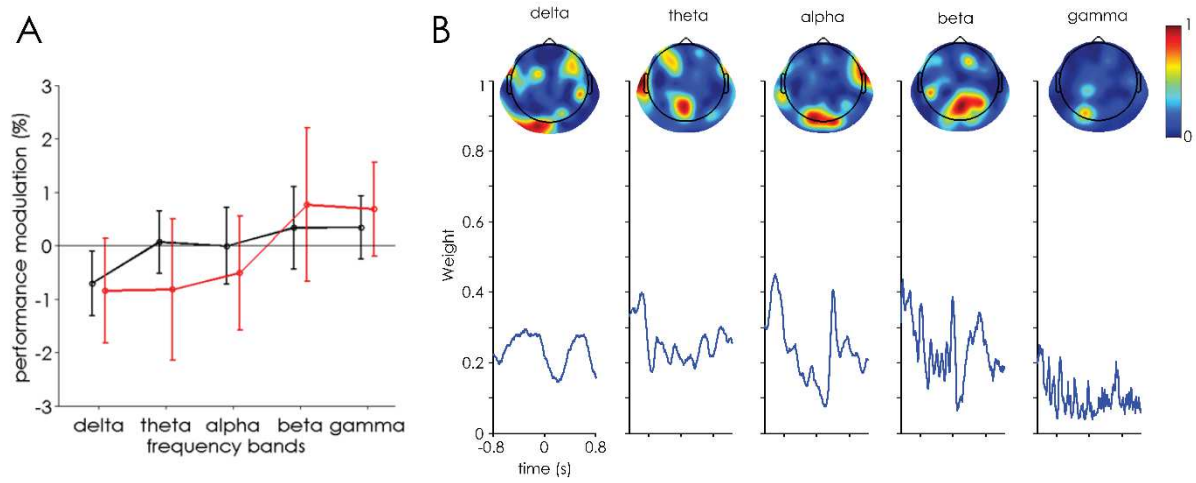

Figure 7-2. Results of the amplitude based classifier for the single subject analysis. A) Mean (and standard error of the mean) performance modulation across subjects for all trials (black) and high agreement trials (in red).B) Mean classifier weights across all classifiers (i.e. averaged across 20 left out subjects, and 10-fold cross-validation) averaged across time points (top) and electrodes (bottom) for each of the 5 frequency bands of interest.

The same analysis was applied to the high agreement trials. Here again, we found no effect in any of the 5 frequency bands: delta (mean performance modulation of -0.846%, p=0.397), theta (-0.825, p=0.54), alpha (-0.513%, p=0.63), beta (0.763, p=0.60), and the gamma (0.682%, p=0.44) band (see Figure 7-2A).

To better understand this (lack of) effect, we evaluated what information the classifier used for its classification (if any) by looking at the mean weights in time and space. By looking at the weights averaged across time (see Figure 7-2B), we see that the amplitude classifier is basing its decision output mainly on the amplitude of the occipital channels (in red) across all frequencies. This is consistent with the results from chapter 5, where the occipital channels showed the largest difference between behavioural outcomes (see Figure 5-3). Looking at the weights in time, the classifier is using information from across the whole epoch, but there seems to be a "dip" in weights just after time 0 (see Figure 7-1B), particularly visible for the delta and alpha band classifiers. This is likely due to the flattening of the white noise sequences that occurs for 150ms around the presentation of the target, which creates an amplitude decrease in the reconstructed EEG (see chapter 2).

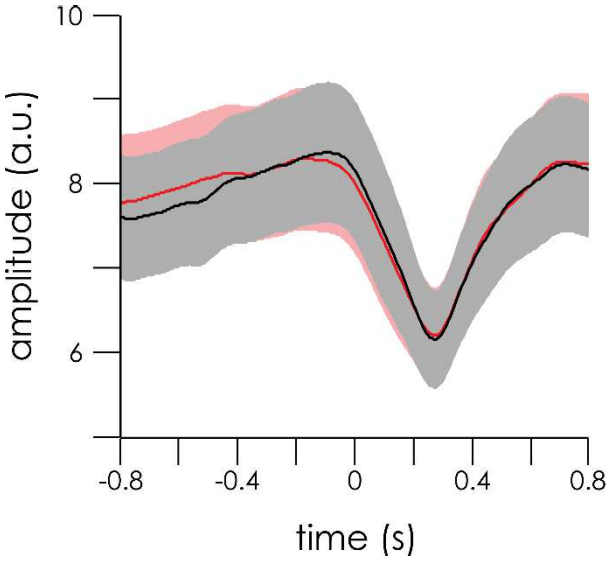

Figure 7-3. Average amplitude for hits (red) and misses (black) across subjects (mean +/- SEM between subjects) on electrode POz for the data filtered in the alpha band.

As we can see, removing the fluctuations in luminance around the target has an influence on the amplitude of the reconstructed EEG: there is a relative decrease in amplitude right after the target presentation (see Figure 7-3). Here only the alpha band is represented, but this effect can be seen at all frequencies.

7.3.2 Phase based classifiers

To evaluate the performance of the phase-based classifier, we computed the performance modulation (i.e. the difference in hit rates between the predicted "good" and predicted "bad" phases). When testing the classifier on all trials, we found that the performance was significantly modulated by the alpha and gamma band (see Figure 7-4). The phase of the alpha band explained 1.85% of modulation ($t(19) = 2.256$, $p = 0.036$, 95% confidence interval for the difference = 0.1338 – 3.569) and the phase of the gamma band oscillation explained 1.23% ($t(19) = 2.2652$, $p = 0.035$, 95% confidence interval for the difference = 0.09367 to 2.3705). There was no significant modulation of performance by the theta band phase (mean difference = 1.2%, $t(19) = 1.769$, $p=0.09$, 95% confidence interval = 0.21 – 2.62). Surprisingly, there were also negative effects: in the delta and beta band, it seemed that we systematically predicted the opposite of what the perception of the subject was. The beta band phase explained -1.45% ($t(19) = -3.7182$, $p = 0.001$, 95% confidence

interval for the difference = -2.26 to -0.633) and the delta band phase explained -2.08% (t(19) = -3.5, p = 0.002, 95% confidence interval for the difference = -3.3 to -0.8).

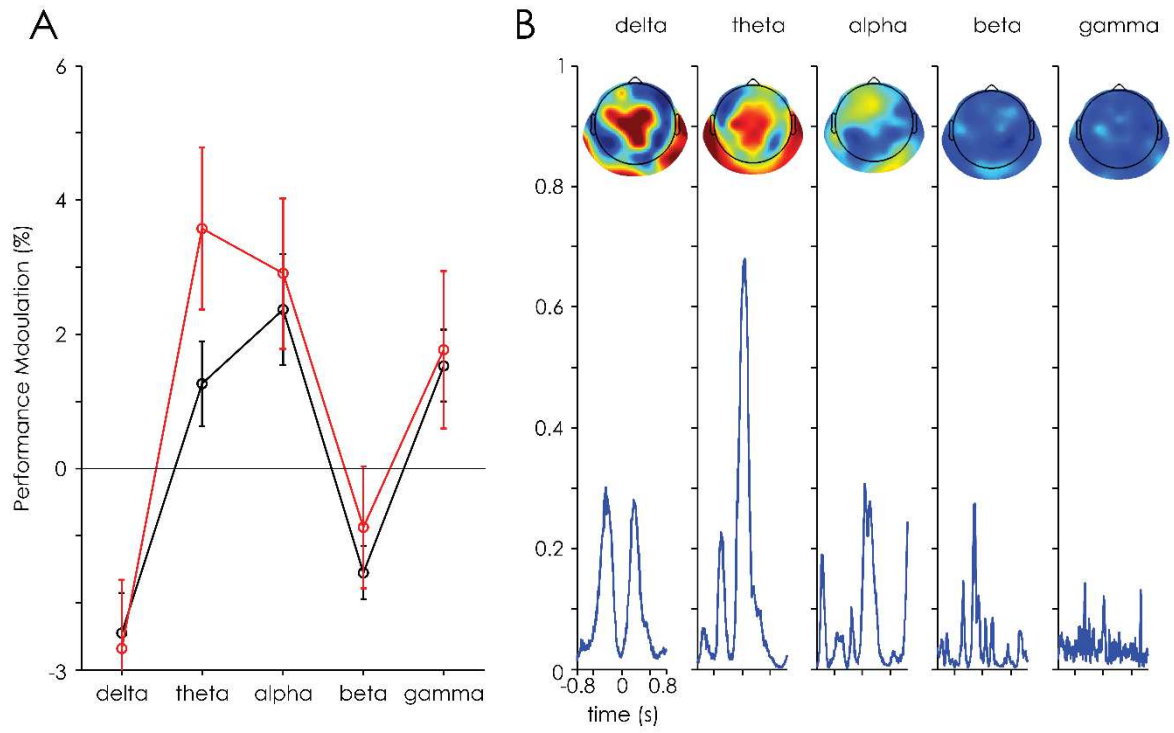

Figure 7-4. A. Performance modulation for the phase based predictor for all trials (in black) and for the high agreement rate trials only (in red). The circles represent the mean and the bars represent the standard error of the mean. B. Average weights of the phase based classifier based on the single subject reconstructed EEG for each frequency band of interest. At the top, the topographies represent the weights averaged in time and at the bottom the weights are averaged across electrodes.

When the predictors were applied only on the high agreement rate trials, we found that there was a significant effect in the theta band and alpha band (see Figure 7-4A). The phase of the theta band oscillation explained 3.58 % of variability in performance (t(19) = 2.9643, p=0.007, 95% confidence interval = 1.05 to 6.1097). The phase of the alpha band oscillation explained 2.91% (t(19) = 2.5933, p = 0.017, 95% confidence interval for the difference = 0.56044 to 5.25). The performance was not well predicted at higher frequencies: the performance modulation based on the phase of the beta band (p=0.349) and the gamma band (p=0.1481) were not significant. Furthermore, there was an effect in the phase of the delta band that was contrary to that expected: the delta phase explained -2.69% (t(19) = -2.6289, p = 0.0165, 95% confidence interval for the difference = -3.3 to -0.8).

Looking at the weights, we see that the effects in the theta band are mainly driven by larger weights in the post-stimulus time window over central electrodes. The alpha band effects are also in the post-stimulus time window, but over more frontal electrodes.

7.4 DISCUSSION

In this chapter, we assessed the possibility a "Neuro-Encryption" system. The aim was to build a classifier which could predict the perception of targets embedded within the white noise sequences. We rendered this prediction subject specific by leveraging the individual differences in the subject's responses to white noise sequences (i.e. in the impulse response functions). We wanted to present targets at key moments within a given white noise sequence so that the target would be seen by a subject with the "good" phase/amplitude for perception and missed by a second subject with a "bad" phase or amplitude for perception.

We find that although it was possible to predict the perception of subjects based on the learned relationship between the phase and perception, the size of the effect was small. The largest effect was found in the phase of theta oscillation: the targets presented at the "good" phase for perception" were more likely to be perceived than the ones presented at the "bad" phase for perception (by about 3.5%). Moreover, there was a negative effect in the delta band which suggests that, systematically, when we predicted a certain phase based on the group level data, the opposite was true for the left over subject. We do not see any obvious explanation for this unexpected finding.

The small size of the positive performance modulations observed in the theta and alpha bands (2-3%) is not entirely surprising. Ultimately, the amount by which we can hope to modulate the behaviour of subjects is bounded by the way in which white noise sequences drive the brain. We saw in chapter 2 that the detection of a limited number of targets was driven by the sequence. Moreover, we saw that a large part of this effect was actually explained by high between subject agreement rates and only a small amount of performance modulation could be expected based on this experiment. The "best" performance modulation we could hope to achieve, based on the within subject agreement rate, was 8% for the high agreement trials. Here, the reconstructed theta phase explains over $1/4^{th}$ of the overall possible performance modulation, which, although far from perfect, is relatively good.

Originally, we had the hope of building a single classifier (based on 20 subjects), which could be used to predict the performance of any new subjects. This would have considerably shortened the "encryption" process. For any new subject, we would only need to record their IRF, this would allow us to model their reconstructed EEG for any new random luminance sequence, and then we could directly derive their perception of embedded targets from the pre-created encryption algorithm. However, the results of the analyses presented in this chapter suggest that this will not work efficiently. Instead, it might be better to also learn for each new subject the relationship between its reconstructed EEG phase or amplitude and perceptual performance. This subject specific predictor could then be used to predict whether that particular subject would be able to see any new target presented within the WN sequences. This procedure, however, would require not only to record EEG while the subject observes random luminance sequences (approximately 1 hour of experiment), but also to systematically measure the relation between random luminance sequences and perception for this subject (at least another 1 hour of experiment per subject).

These results also raise the question of the type of task used. Throughout this thesis, we chose to use a simple detection task, where subjects had to respond to the presentation of a near-perceptual threshold target embedded within the white noise sequence. This ensured that both conditions (hits and misses) had relatively similar number of trials, and is a good paradigm to evaluate the effect of phase and amplitude on perception. From a theoretical point of view, the choice of this paradigm also was justified. We wanted to test the functional relevance of the perceptual echoes. In particular, it might reflect the driving of a local sampling mechanism in the alpha band (corresponding to the "occipital alpha" described in the introduction. If this were the case, we would expect perception to be modulated rhythmically as a function of the underlying fluctuations in neuronal excitability. If the perceptual echoes are a marker of the time course of this activity, then presenting near perceptual threshold target can directly test this. However, the white noise sequences create large masking effects on the detection of the target (see chapter 2, question 3: classification image): luminance values presented around the target had a direct influence on perception. Thus we had to remove them. As a consequence, there was 150ms in the middle of the epoch which had a reduced power in the reconstructed EEG. This can in part explain the "dip" seen in the analysis of amplitude difference (in chapters 5 and in chapter 6 more visibly) and from

the classifier weights presented above (see Figure 7-2). This is not a problem for phase analysis as the amplitude was only reduced and not 0, thus phase could be reliably estimated.

Consequently, it might be possible to boost the performance of a Neuro-Encryption system by using a different task instead. For example, using supra-threshold targets might reduce the masking effects of the WN sequences and thus remove the need to flatten the sequences. For example, temporal parsing, the ability of segregating two inputs in time, has been linked to the phase of the ongoing oscillation (Varela, Toro, John, & Schwartz, 1981) but most importantly, it has also been linked to individual differences in alpha peak frequency (Samaha & Postle, 2015). This might prove to be a more optimal task since the frequency of the echoes also vary between subjects. Using a temporal parsing task embedded within the WN sequences instead of a target detection task might take advantage of the differences in frequencies.

Once our prediction ability has been improved (either by building subject specific classifiers or by changing the task as discussed above), we will then be able to start building a true Neuro-Encryption system (see Figure 7-5).

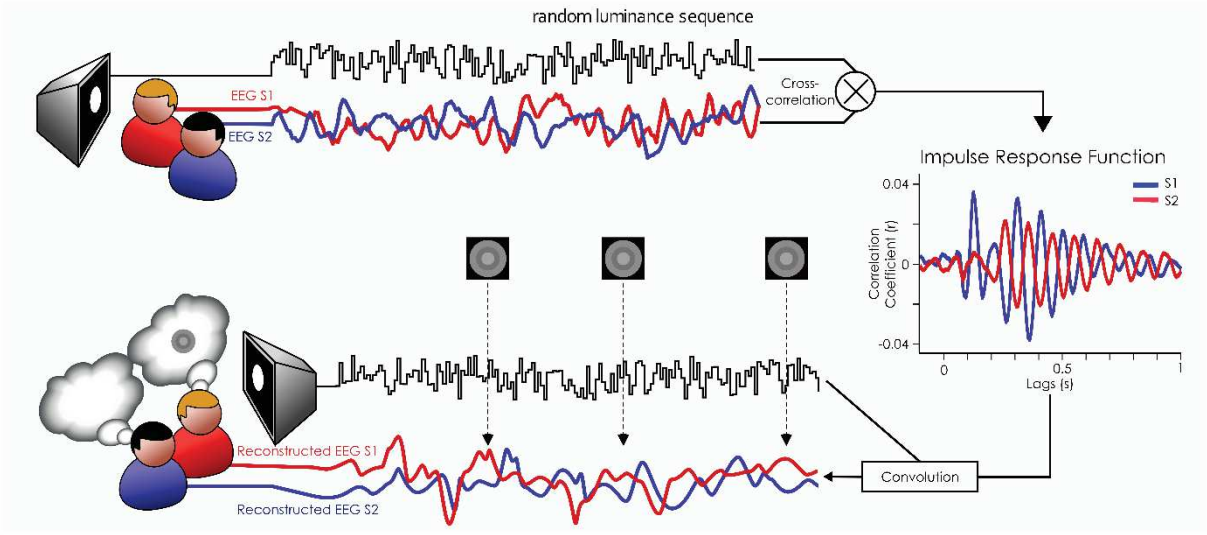

Figure 7-5. Illustration of the idea behind Neuro-Encryption. We know that the impulse response to white noise stimulation is composed by a long lasting reverberation of information in the brain, akin to a perceptual echoes. These perceptual echoes are subject specific, as can be clearly seen from these two subjects: at almost all time points between 300ms and 800ms, the IRF of the two subjects with regards to the flash is exactly opposite. Moreover, we know that there is a relation between the phase of the oscillation and

perception: stimuli flashed at a certain phase are more likely to be perceived than stimuli flashed at the other phase. Therefore, we can use the white noise sequences to constrain the state of oscillations. Then, we can present targets "optimized" for one subject only, so that, with the same stimulation, the subjects will have opposite perception.

Behind the idea of Neuro-encryption lies the hope that perception can be modulated for each subject independently (see Figure 7-5). We aim to use IRFs as encryption keys. Because of individual differences in their IRFs, for any WN sequences, two subjects will have different time courses for the phase and the amplitude of the constrained brain activity. Consequently, we can decide to present targets at moments which are optimal for the perception of one specific subject and thus boost their perception of the target. Importantly, we can also choose moments where two subjects have opposite predictions for the same target. Consequently, a target would be seen by the subject with the optimal brain activity at that time point, and missed by the subject with the non-optimal brain activity. Effectively, this would build a sort of "Neuro-Encryption" system, whereby targets could be hidden within the white noise sequences so that a message would be decrypted only by the person with the right "encryption" key (i.e. perceptual echo).

Chapter 8. DISCUSSION

Alpha oscillations are the hallmark of the visual system: from the very first electrophysiological recording (Berger, 1929) they have been consistently linked to the processing of visual information. They are thought to reflect an active inhibitory mechanism whereby task irrelevant brain regions are inhibited, allowing the gating of information (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007) and thus sensory selection (Schroeder & Lakatos, 2008). Alpha oscillations are pervasive in the visual system: alpha ringing has been shown to arise as a function of single flashes (Barlow, 1960) or even brain stimulation (Herring et al., 2015). Even when more complex stimuli are presented such as broad-band visual stimulation, our visual system responds with a long lasting reverberation of the stimulus information in the alpha band, a sort of "perceptual echo" (VanRullen & Macdonald, 2012). Better understanding these perceptual echoes was at the heart of this thesis. We devised a paradigm based on white noise stimulation to uncover their functional role: why do our brains reverberate in the alpha band in response to WN sequences? Do these alpha fluctuations have a functional role? In particular, we studied them as a possible reflection of the rhythmic nature of perception. By using both fMRI and EEG recordings, we investigated their neuronal basis as well as their influence on target detection.

8.1 SUMMARY OF THE CHAPTERS

Throughout this thesis, the WN paradigm was used to determine whether the perceptual echoes had any functional role in visual perception. We hypothesized that perceptual echoes were the direct reflection of the oscillations giving rise to perceptual cycles. We presented WN sequences to constrain these fluctuations. In chapter 1, we distinguished 6 types of signals which could be recorded in the EEG in response to WN sequences. The signals were either irrelevant (signals S1, S3, S5) or relevant (signals S2, S4, S6) for target perception, and either spontaneous (signals S1 and S2), phase locked (signals S3 and S4) or induced (signals S5 and S6) with regards to the WN noise sequences. A convolution of the WN sequence and the perceptual echoes allowed us to track their time course with regards to these random luminance sequences (i.e. signals S3 and S4), regardless of the state of any

other signals. Thus, we could probe their functional role by presenting near perceptual targets embedded within the WN sequences, which fell at random moments with regards to the state of these constrained fluctuations.

In chapter 2, we validated the WN paradigm as a tool to study the time course of perceptual echoes. First, we showed that a significant part of the recorded EEG can be successfully reconstructed from the convolution between perceptual echoes and WN sequences. Using the perceptual echoes as a model of the brain response to WN sequences is restrictive: it only captures the linear relationship between brain activity and stimulation, all the rest is removed. Yet, this restrictive model of brain response to flashes could still approximate relatively well the fluctuations in EEG signal. In fact, it did just as well as other (more standard) models of brain activity such as the ERP to isolated targets. Furthermore, the prediction ability was similar to that based only on the "early" part of the IRF (Lalor, 2009). Next, we showed that the perceptual echoes were stable in time by measuring, at intervals longer than 6 months, the correlation between echoes of the same or different subjects. There was a very strong relationship between two echoes recorded for the same subject ($r>0.8$), and a weak (but not null) relationship between the two echoes recorded for different subjects. However, we also saw a significant decrease in the correlation between two echoes of the same subject over time: it steadily decreased at a rate of 0.07 point per year. As a consequence, the difference in correlation strength between the echoes recorded on the same and different subjects was predicted to remain visible for about 12 years. Finally, we ensured that the targets were not masked by the luminance values of the WN, and that the WN sequences still played a role in their perception.

In chapter 3 we showed that the source of these perceptual echoes was likely from the early visual areas, areas V1 and V2. Interestingly, this activity was widely distributed throughout the early visual areas and not specific to the retinotopic location of the stimulation. This seems to support the notion that perceptual echoes act in travelling waves in the visual cortex (Lozano-Soldevilla & Vanrullen, 2016). Moreover, this is also an indirect supporting evidence for the functional role of perceptual echoes. The fact that we found activation in local areas of the visual cortex supports the idea that these could reflect the local sampling

mechanism of the brain (Lansing, 1957; Lindsley, 1952), which gates the incoming information as a function of the local excitatory activity.

In chapter 4 and 5, we showed that the perception of the embedded targets was linked to both the post-stimulus theta phase (chapter 4) and the alpha amplitude (chapter 5). Furthermore, in chapter 4 we also showed that the use of the white noise sequences allows us to recover, independently from the evoked response, the true latency of phase effects. We showed that the evoked response to the target could explain the pre-stimulus phase effects reported in the literature (Busch et al., 2009; Drewes & VanRullen, 2011; Mathewson et al., 2009; VanRullen et al., 2011). In chapter 5, we showed that the effects of amplitude could be shown independently from the effects of endogenous oscillations (such as attention).

Finally, in the last two chapters of this thesis, we evaluated whether we could use the fact that the perceptual echoes represent the local sampling rhythm of the brain to our advantage. We constructed classifiers based on the reconstructed EEG, which links its spectral content (i.e. phase and amplitude) to the visibility of the target with two different aims.

In chapter 6, we looked for the universal mechanisms (i.e. independently of the subject) which lead to target perception. What are the oscillatory features that will systematically give rise to visual experience across all subjects? We find that the phase of higher theta/lower alpha (from 5 to 9 Hz) is instrumental in perception.

In chapter 7, we asked whether we can leverage the subject specific relationship between white noise sequences and oscillatory brain activity on the one hand, and oscillatory brain activity and visual perception on the other hand to influence perception. This was met with limited success: we did find that the phase could be used to predict the perception of new subjects, but we also found that there were large between subject similarities in the phases evoked by the subject. Using a different type of task might be more promising (as we will discuss in section 8.4).

Taken together, the evidence gathered so far supports the idea that perceptual echoes play a direct functional role in perception. In fact, they directly influence the probability of perceiving near perceptual threshold targets. This supports the hypothesis that perceptual

echoes directly reflect the perceptual sampling occurring, while the visual processing is taking place (VanRullen, 2016b). In fact, it could be that the reverberations present in the perceptual echoes are evidence of natural fluctuations in local excitability which create perceptual cycles *during* stimulus processing.

8.2 PERCEPTUAL ECHOES AND PERCEPTUAL CYCLES

In this thesis, we found evidence that perceptual echoes are a direct reflection of the presence of perceptual cycles in vision: the detection of targets embedded within the white noise sequences fluctuated rhythmically as a function of the underlying constrained oscillation (see Figure 8-1). But how are these perceptual echoes related to the notion of perceptual cycles?

Recently, a review of the evidence for the presence of perceptual cycles in vision highlighted the presence of effects in two separate frequencies: one at 7 Hz and 11 Hz with seemingly distinct topographies and functional roles (VanRullen, 2016b). The "theta" band (7Hz) oscillations were more often found when tasks involving attentional demands were used. Moreover, they were characterized by effects appearing over frontal channels. The "alpha" band oscillations (11 Hz) were more likely to be reported in tasks concerned with sensory processing and conscious perception, often over occipital channels. It has been proposed that these mechanisms could reflect the relative contributions of an attentional and a sensory sampling mechanism of visual perception (Dugué & VanRullen, 2017; VanRullen, 2016b; Zoefel & VanRullen, 2017).

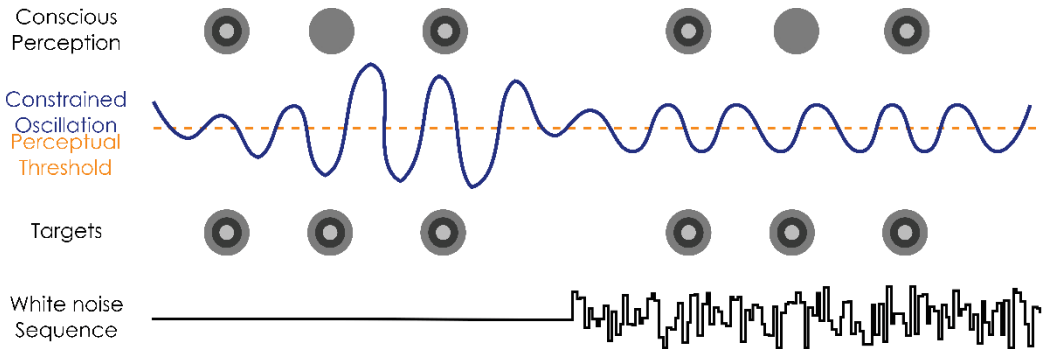

Figure 8-1. Illustration of the relationship between the oscillations constrained by the WN sequence and target detection. A target presented at key phases will be perceived only when it falls at the optimal phase for perception with regards to the constrained oscillations. These oscillations also exist when no WN sequence is presented.

Because they are directly linked with fluctuations in target detection, are in the alpha band, predominantly over the occipital channels, the perceptual echoes could therefore represent the fluctuations linked with the second oscillatory mechanism: the occipital alpha oscillations. This local alpha sampling mechanism reflects the oscillations carrying out the local sampling of incoming visual information, as proposed by the "cortical excitability cycles" hypothesis (Harter, 1967; Lansing, 1957; Lindsley, 1952). The perceptual echoes could be a direct insight into the time course of the excitability fluctuations reflected in the "local" occipital alpha sampling mechanism. They could directly reflect the gating of incoming information as it comes in.

This raises the question of which "phase" we are referring to when we talk about reconstructed EEG phase gating visual perception. The limitation of using such a global measure of brain activity as EEG is that the measured EEG phase becomes "meaningless" with regards to the firing at the single neuron level and the ongoing LFP phase as EEG integrates information across large portions of cortex. As such, EEG has a very poor spatial resolution. Thus, although we have shown that WN sequences can be used to track and constrain the activity of this "local" sampling mechanism, it is still unresolved as to where in the brain this mechanism happens. The fMRI experiment (chapter 3) points towards early visual area (i.e. V1 and V2), but further investigations are needed. In particular, it might be interesting to determine at what level of the information processing hierarchy this gating happens? In fact, this question is directly related to the issue of what type of activity leads to conscious perception. Is there any information that is processed in the brain before the activity is gated by the phase? Up to what level is the information processed?

It is notoriously hard to find evidence for an influence of phase on perception: the recent surge in studies reporting such effects are the direct consequence of technological advances and increased computational power. Finding effects of phase on perception requires the use of carefully controlled experimental paradigms, such as for example using near-perceptual threshold targets, and specific statistical tools (e.g. inter-trial phase clustering or phase opposition) which have been developed in the past 20 years. This raises the question of the validity and generalizability of these findings to more ecological settings are. Could we find an effect of phase on performance in more "realistic" situations? It may well be that this would not be the case, for the simple reason that the most readily apparent effects of phase on perception may well be for stimuli which are very briefly presented. This may be a direct consequence of the presence of a mechanism in lower visual areas, which '"buffers"

information arriving at moments of lower excitability so that it is actually perceived (Dubois, 2011). This buffering might be for example reflected in the content of the perceptual echoes them-selves: with regards to the initial impulse presented at time t, the echoes represent the reverberation of the initial information in time. Thus they could act as a buffer of information, whereby stimuli arriving at the wrong phase will be delayed and reactivated at the next cycle of the oscillation. This would explain why only very brief and weak stimuli will be missed. The others will be simply delayed and reactivated until the next cycle. In fact, this could be directly tested using the paradigm used in this thesis (this will be discussed in section 8.4).

If the presence of perceptual cycles might lead us to miss some information, why did this mechanism evolve? An easy answer to this, is the benefits in terms of energy cost. Biologically relevant motion is unlikely to change very much in the time frame of an alpha cycles (i.e. 100ms): the movement of predators, or human actions are spread over longer time spans. Thus if the visual information remains relatively stable, then subsampling the environment will not be a problem. In fact, our visual system has been shown to be robust to temporal subsampling (VanRullen et al., 2014). Thus, our visual system can take samples of the environment at various interval to extract an accurate depiction of the external world while decreasing the energy cost considerably. In fact, this sampling of information naturally occurs through eye movements: every 300 to 500ms we make saccades to different points of the environment. Our brains receive information in waves of activation, thus subsampling the information at an even faster rate than the saccade rate will not results in a large information loss (Zoefel & VanRullen, 2017). Finally, it might be that these fluctuations in excitability could be implemented to provide a frame of reference for the incoming information: by sampling the information rhythmically, the brain creates a reference frame, by which the relative timing of spikes could thus directly be readout as information (VanRullen et al., 2005).

8.3 DIFFERENT TYPES OF ALPHA OSCILLATIONS

Next we turn to evaluate the specific role of the fluctuations with regards to the other types of signals which were presented in the introduction through the specific taxonomy (see table 1). Reconstructing the EEG to WN sequences allowed us to extract the phase locked oscillations (i.e. signals S3 and S4), which reflect the direct time course of the perceptual echoes, reflecting a rhythmic sampling of the visual information. Here, we will consider how the phase and the amplitude of these fluctuations relate to other alpha oscillations reported in the literature.

In chapter 4, we showed (using simulations) that a phase difference between two conditions (e.g. seen and missed targets), inserted in the post stimulus time window (at 40ms after a target) was systematically detected in the pre-stimulus time window, up to 100ms before the stimulus presentation. This was especially true at lower frequencies (up to about 30 Hz). This "temporal shift" in measured latencies of phase modulation could be explained by the time frequency content of the evoked response to the "stimulus". Furthermore, we showed that using the reconstructed EEG allowed us to circumvent this influence of the evoked response and recover the true latency of phase effects: in the post-stimulus time window *while* the target stimulus is processed by the brain. As discussed in chapter 4, this is a direct benefit of reconstructing rather than recording the EEG. This allowed us to extract the time course of the constrained oscillations outside of any influence of the target evoked potentials.

In fact, using the WN paradigm also allows us to disentangle the relative contribution of the "relevant" signals that is the spontaneous oscillations (S2) and the phase locked oscillations (S4) to these phase effects. Note that signals S6 is not relevant here because induced signals have, by definition, no phase relationship with the target stimuli. However, this is a requirement for the usual methods of analysis of phase differences. In chapter 4, we showed that the effects of the constrained oscillations' phase (that is the phase of signal S4) on perception was almost as large as the effects usually reported in the literature from "spontaneous" fluctuations. We reported a modulation of performance by the phase of the reconstructed EEG of 11% while Busch and colleagues (2009) reported a modulation of about 15% based on the recorded EEG signal. This leaves open the question of the potential influence of the phase of spontaneous signals (S2) in perception. Does their phase play any role in target detection? In terms of phase effects, if these mechanisms are at play, they are

likely to be small (as S4 phase can explain almost as much modulation in performance). In fact, it is possible that all of the phase effects (e.g. Busch et al. 2009) previously reported could actually be explained, not by the effect of this top down mechanism, but rather by the same oscillations which are constrained by the WN sequences. In fact, these oscillations could also explain the effect reported by Mathewson and colleagues (as they have themselves suggested), whereby the mask and target are grouped together into one "snapshot" (Mathewson et al., 2009).

This distinction the alpha oscillations related to a local "sampling mechanism" and those related to top-down "sensory inhibition" was further explored in chapter 5. We showed the modulation of target detection depending on the alpha band power could occur outside of any influence of endogenous factors (such as attention). It is important to note that in the taxonomy we created, the top-down endogenous fluctuations in attentional sampling can be clearly classified as belonging to the "relevant spontaneous signals" category (i.e. signal S2). That is not to say that these signals do not play a role in the detection of the targets presented within the white noise sequence, but they will not be included in the reconstructed EEG. However, there might still be room for attentional mechanisms to be present in the reconstructed EEG: the WN sequences themselves could capture the exogenous component of the attentional mechanisms of our brains. During the stimulation, subjects report feelings of fluctuations in the strength of the luminance. Moments of stronger luminance could draw the exogenous attention to the WN sequences in a phase locked manner. This effect is likely to be small as the IRF are present even when higher luminance trials are discarded from the analysis (VanRullen & Macdonald, 2012). Nevertheless, it cannot be entirely excluded that an exogenous attentional modulation might be present. However, the influence of endogenous attention can be clearly rejected, as these fluctuations are unlikely to be phase locked to the WN sequences. In fact, this might, in part explain the relative weakness of the effects of amplitude reported throughout this thesis. We found that the amplitude of the reconstructed EEG was related to perception (chapter 4), however, these effects were smaller than those of the phase, and inconsistent. We were able to extract them at the group level, but repeatedly failed to find any generalizable link between amplitude and perception across subjects and across trials (i.e. chapter 6 and 7). Of course, these small effects of the alpha amplitude on perception should be taken with careful consideration. Caution is warranted in interpreting the absence of significant effect as a proof of the absence of an effect. Nevertheless, this is

coherent with the general idea that several alpha rhythms, with different roles might be at play here: a local "sampling" mechanism could be reflected in the alpha of the perceptual echoes while a "sensory inhibition" mechanism could be supported by the endogenously modulated alpha oscillations.

This difference between the effects of phase and amplitude effects might reflect an intrinsic difference in the temporality of the two mechanisms at play. The local "sampling mechanism" would predominantly act through the modulation of the instantaneous excitability (i.e. phase), while the top-down "sensory inhibition" mechanism would act through a more sustained inhibition of excitability (i.e. amplitude). This is also in accord with a recent study which showed that the power (but not the phase) of the alpha oscillations could be modulated by attentional demands and temporal expectancy (van Diepen et al., 2015).

Accordingly, the alpha oscillations present in the IRFs would be a different mechanism than the ongoing alpha oscillations which have been linked with the locus of spatial attention (Kelly et al., 2006; Romei et al., 2010; Samaha et al., 2016; Sauseng et al., 2005; Thut et al., 2006) and involved in inhibitory role (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Klimesch, Sauseng, et al., 2007; Mathewson et al., 2009).

In fact, this differentiation is further warranted when considering the effect of attention on the impulse response functions themselves. When subjects pay more attention to the WN sequences, the perceptual echoes show an increase in amplitude (VanRullen & Macdonald, 2012). Note that Jia and colleagues recently reported the opposite effect: they found an increased amplitude in the IRF computed to the unattended stimulus than in the IRF to the attended side (Jia et al., 2017). This effect, however was restricted to the first 200ms of the cross-correlation window. They did find a larger amplitude in the attended condition than in the unattended condition in the "late" part of the IRF (i.e. from about 200/250ms onwards; Jia et al., 2017). This increased amplitude in attended versus unattended stimuli might not be so surprising if the perceptual echoes do reflect a local sampling mechanism: the allocation of attention has been shown to increase the firing rate of cells in various visual brain areas, from the LGN to V4 (Bisley, 2010). As such, this also points towards the perceptual echoes reflecting the active processing of visual information.

In fact the suggestion of different co-existing alpha oscillations with different functional roles is coherent with findings of several alpha generators in various layers of the

cerebral cortex. Alpha generators are reported in both the infragranular layers (Spaak, Bonnefond, Maier, Leopold, & Jensen, 2012; van Kerkoerle et al., 2014), consistent with a feedback, inhibitory role of alpha oscillations and in superficial layers (Bollimunta et al., 2008; Haegens et al., 2015) consistent with the idea of a local sampling mechanism involved in the feedforward loop. In fact, Bollimunta and colleagues (2008) showed that these two generators were differently linked with perception. The alpha power in the infra-granular layers was negatively correlated to behavioural performance while the alpha power in the supra-granular layers was positively correlated with behavioural performance (Bollimunta et al., 2008).

8.4 PERSPECTIVES: TOWARDS A NEURO-ENCRYPTION SYSTEM

This thesis can be seen in some ways, as a proof of concept: the white noise paradigm was successfully used to study the time course of perceptual cycles. This opens up possibilities for testing new questions about the functioning and the role of these rhythmic fluctuations in perception.

In the previous section, we have made a distinction between different types of alpha oscillations present over the visual cortex. This could be even more firmly confirmed by a TMS experiment. If these perceptual echoes reflect fluctuations in the local excitability levels, then applying single TMS pulses at the phases pre-determined based on the reconstructed EEG should predict the perception of phosphenes. Interestingly, this creates the possibility of investigating both the "higher-level" frontal theta/alpha mechanism and the "lower level" alpha oscillations in conjunction. A simple attentional manipulation could be added to the task above by asking participants to pay attention to the WN sequence or not. The perception of the target could then be evaluated with regards to the constrained brain oscillations, but also with regards to the recorded oscillations over occipital channels or frontal channels.

An outstanding question remains about the content of these perceptual echoes. They reflect a reverberation of the information about the pulse intensity presented up to 1s earlier and which is rhythmically re-activated by the visual cortex. Could further information be carried in each of these cycles? Cross-frequency coupling between alpha and gamma

oscillations has been suggested as a potential mechanism through which the "content" of the representation (carried in the gamma cycle) is segregated and ordered by the alpha rhythm which provides a frame of reference (e.g. Jensen et al., 2012). This cross-frequency coupling between ongoing alpha and gamma oscillations has been shown to occur in the visual cortex (Spaak et al., 2012). Therefore, could these perceptual echoes also carry information? If so, what type of information?

Another avenue of research is the individual differences in perceptual echoes. As we have seen from this thesis, the individual differences in the phase of the IRF translate as subtle differences in the phase of the reconstructed EEG. However, there are other differences in the IRF: the length varies greatly between subjects and the strength also varies. Can these inter-individual differences be linked to differences in perceptual abilities? Furthermore, could a source of these individual differences be found in the specific brain patterns in response to stimuli which are randomly varied along another dimension? So far, perceptual echoes have been found only in response to random changes in light intensity. However, because the brain is a highly nonlinear system, these IRF only represent the linear brain response to this specific type of stimuli (i.e. intensity modulated sequences). Lalor and colleagues have shown that the VESPA evoked by contrast stimuli had a different spatiotemporal profile than that elicited in response to modulation of spatial frequency (Lalor, Lucan, et al., 2009). Therefore, this raises the question of what would happen to the perceptual echoes when other types of stimuli are used.

Further than this, there might be within-subjects variability which have yet to be explored. We showed in the longitudinal study that perceptual echoes are stable in time up to 7 years. How do they change as a function of time? It would be interesting to carry out a longer longitudinal analysis to see how the relationship between the ongoing alpha frequency and the echo frequency evolved. Since they might reflect to separate mechanisms, do the echo frequency also decrease with age as do the alpha peak frequency?

Finally, as already mentioned in chapter 7, we could use the WN paradigm to develop a Neuro-Encryption system. Visual information present in the sequence would only be detected by the subject with the right "encryption" key (i.e. IRF, see Figure 7-5). This, however, was met with limited success so far. This might be due to the type of task used in this thesis: a detection task might not be the best suited. We chose this task because it was easy to implement, and yielded a relatively equal number of "missed" and "seen" targets, which is necessary for an unbiased measure of the phase difference between groups (the inter trial phase clustering is highly influenced by the trial count). Moreover, using a detection task would have allowed a true encryption: in the ideal case where phase modulated performance completely, the targets presented at the "wrong" phase for perception would never have been perceived by the subject.

A possible alternative would be to use a task which uses the gap property, which arises as a consequence of perceptual cycles introduced at the beginning of the thesis. Varela and colleagues reported in 1981 that the perception of two flashes presented close together in time was influenced by the phase at which they were presented. The two flashes were more likely to be perceived as simultaneous if the stimuli were flashed at the trough of the alpha oscillation (Varela et al., 1981). On the other hand, an apparent motion was created between the two flashes if they were flashed at the opposite phase (Varela et al., 1981). We could use this to present supra-threshold letters, which would need to be segregated to be perceived correctly and thus for the "encrypted" message to be read out (see Figure 8-2).

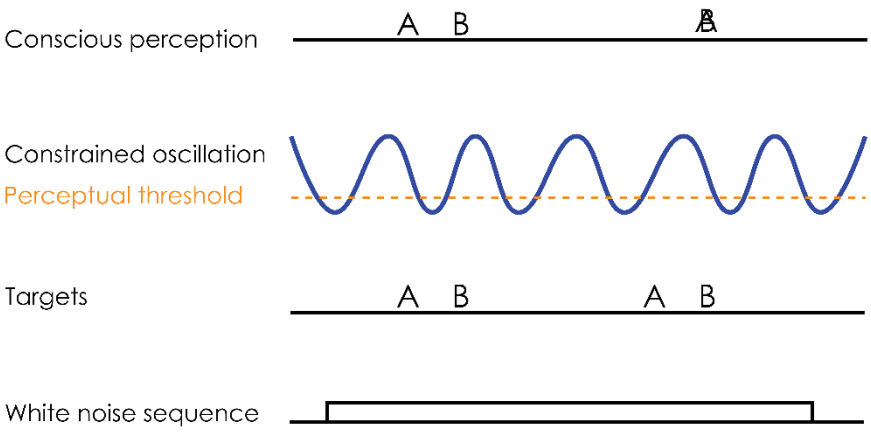

Figure 8-2. Temporal framing and WN paradigm. The phase of the constrained oscillations is determined by the WN sequences. Concurrently, targets can be presented, which will be perceived either as separate entities (i.e. the first two targets) or merged into one percept (the next two targets) depending on the phase of the constrained oscillation.

This would also allow us to directly test the idea of temporal framing in relation to both the phase and the frequency of perceptual echoes. We have seen in chapter 6, the WN sequences constrain the phase of the oscillation in a similar way across subjects. Therefore, we could extract the reconstructed EEG phase and present two stimuli at different phases to see if they are integrated or not. Furthermore, the WN paradigm allows us to make predictions about the frequency at which the effect will take place. Recently, Samaha and Postle revealed that the individual alpha peak frequency reflected the temporal resolution of subjects (Samaha & Postle, 2015). If the oscillation giving rise to the perceptual echoes plays a causal role in this perceptual framing mechanism, then the peak frequency of the echo should be directly related to the optimal inter stimulus interval for segregation across subjects. Moreover, we should be able to see that this segregation or merging of information would be dependent on the phase of the constrained oscillation (see Figure 8-2). Thus, with one experiment, we could show both effects.

8.5 CONCLUSION

Let us go back to at the cross-roads, in the middle of bustling city centre. You are about to cross a busy road to reach the other side of the street where your friend awaits for you in front of the cinema. Your brain is bombarded with thousands of information at the same time: you see the oncoming traffic, your friend waving, the posters for the latest movie and that passer-by picking up his phone... Limited by its finite resources, our brain needs to make sense of all this information in a fast and efficient manner. By rhythmically sampling the incoming visual information, priorities can be set and attention directed at the most important visual information: the car and not the posters or the passer-by. These rhythms of our brains create rhythms to our perception, which are instrumental in coding incoming information and sampling the relevant information in the visual environment.

As we have seen throughout this thesis, the rhythms giving rise to perceptual cycles can be directly accessed and investigated by presenting random continuous (white noise) stimuli. This is evidenced by the presence of perceptual echoes. The perceptual echoes allow us to directly tap into the rhythmic activity giving rise to the fluctuations in perception, and, as such, are a great avenue for better understanding the role of oscillations in vision.

Chapter 9. REFERENCES

Adrian, E. D. (1928). *The Basis of Sensation: The action of the Sense Organ*. London: Christophers.

Adrian, E. D. (1941). Afferent discharges to the cerebral cortex from peripheral sense organs. *Journal of Physiology*, *100*(2), 159–191. http://doi.org/10.1097/00005053-194212000-00038

Ahumada, A. J. (2002). Classification image weights and internal noise level estimation. *Journal of Vision*, *2*(1), 121–131. http://doi.org/10.1167/2.1.8

Baldauf, D., & Desimone, R. (2014). Neural Mechanisms of Object-Based Attention. *Science*, *344*(6182), 424--427. http://doi.org/10.1126/science.1247003

Banerjee, S., Snyder, A. C., Molholm, S., & Foxe, J. J. (2011). Oscillatory alpha-band mechanisms and the deployment of spatial attention to anticipated auditory and visual target locations: supramodal or sensory-specific control mechanisms? *The Journal of Neuroscience*, *31*(27), 9923–9932. http://doi.org/10.1523/JNEUROSCI.4660-10.2011

Barlow, J. S. (1960). Rhythmic activity induced by photic stimulation in relation to intrinsic alpha activity of the brain in man. *Electroencephalography and Clinical Neurophysiology*, *12*(2), 317–326. http://doi.org/https://doi.org/10.1016/0013-4694(60)90005-5

Başar, E., & Güntekin, B. (2013). Review of delta, theta, alpha, beta and gamma response oscillations in neuropsychiatric disorders. In *Clinical Neurophysiology* (Vol. 62, pp. 19–54). Elsevier. http://doi.org/10.1016/B978-0-7020-5307-8.00002-8

Başar, E., Schmiedt-Fehr, C., Mathes, B., Femir, B., Emek-Savaş, D. D., Tülay, E., … Başar-Eroğlu, C. (2016). What does the broken brain say to the neuroscientist? Oscillations and Connectivity in Schizophrenia, Alzheimer's Disease, and Bipolar Disorder. *International Journal of Psychophysiology*, *103*(May), 135–148. http://doi.org/10.1016/j.ijpsycho.2015.02.004

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J. M., Oostenveld, R., Dowdall, J. R., … Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401.

http://doi.org/10.1016/j.neuron.2014.12.018

Baumgarten, T. J., Schnitzler, A., & Lange, J. (2015). Beta oscillations define discrete perceptual cycles in the somatosensory domain. *Proceedings of the National Academy of Sciences*, *112*(39), 12187–12192. http://doi.org/10.1073/pnas.1501438112

Bazanova, O. M., & Vernon, D. (2013). Interpreting EEG alpha activity. *Neuroscience & Biobehavioral Reviews*, *44*, 94–110. http://doi.org/10.1016/j.neubiorev.2013.05.007

Becker, R., Ritter, P., & Villringer, A. (2008). Influence of ongoing alpha rhythm on the visual evoked potential. *NeuroImage*, *39*(2), 707–716. http://doi.org/10.1016/j.neuroimage.2007.09.016

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. Arch. Psychiatr. Nervenkr. 87, 527 – 570. *Archiv Für Psychiatrie Und Nervenkrankheiten*, *81*(1), 527–570.

Bishop, G. H. (1932). Cyclic changes in excitability of the optic pathway of the rabbit. *American Journal of Physiology*, *103*(1), 213–224.

Bishop, G. H., & O'Leary, J. (1938). Potential records from the optic cortex of the cat, 391–404.

Bisley, J. W. (2010). The neural basis of visual attention. *Journal of Physiology*, *589*(1), 49–57. http://doi.org/10.1113/jphysiol.2010.192666

Bollimunta, A., Chen, Y., Schroeder, C. E., & Ding, M. (2008). Neuronal Mechanisms of Cortical Alpha Oscillations in Awake- behaving Macaques. *The Journal of Neuroscience*, *28*(40), 9976–9988. http://doi.org/10.1523/JNEUROSCI.2699-08.2008.Neuronal

Boncompte, G., Villena-Gonzalze, M., Cosmelli, D., & Lopez, V. (2016). Spontaneous alpha power lateralization predicts detection performance in an un-cued signal detection task. *PLoS ONE*, *11*(8), 1–13. http://doi.org/10.1371/journal.pone.0160347

Bonnefond, M., Kastner, S., & Jensen, O. (2017). Communication between Brain Areas Based on Nested Oscillations. *Eneuro*, *4*(April), ENEURO.0153-16.2017. http://doi.org/10.1523/ENEURO.0153-16.2017

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.

http://doi.org/10.1163/156856897X00357

Briggs, F., & Usrey, W. M. (2008). Emerging views of corticothalamic function. *Current Opinion in Neurobiology*, *18*(4), 403–407. http://doi.org/10.1016/j.conb.2008.09.002

Brüers, S., & VanRullen, R. (2017). At what latency does the phase of brain oscillations influence perception? *eNeuro*, *4*(3), ENEURO.0078-17.2017. http://doi.org/http://dx.doi.org/10.1523/ENEURO.0078-17.2017 1

Burgess, N., & O'Keefe, J. (2011). Models of place and grid cell firing and theta rhythmicity. *Current Opinion in Neurobiology*, *21*(5), 734–744. http://doi.org/10.1016/j.conb.2011.07.002

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *The Journal of Neuroscience*, *29*(24), 7869–7876. http://doi.org/10.1523/JNEUROSCI.0113-09.2009

Busch, N. A., & VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, *107*(37), 16048–16053. http://doi.org/10.1073/pnas.1004801107

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780195301069.001.0001

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. http://doi.org/10.1038/nrn3241

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929. http://doi.org/10.1126/science.1099745

Calderone, D. J., Lakatos, P., Butler, P. D., & Castellanos, F. X. (2014). Entrainment of neural oscillations as a modifiable substrate of attention. *Trends in Cognitive Sciences*, *18*(6), 300–309. http://doi.org/10.1016/j.tics.2014.02.005

Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515. http://doi.org/10.1016/j.tics.2010.09.001

Capilla, A., Pazo-Alvarez, P., Darriba, A., Campo, P., & Gross, J. (2011). Steady-state visual evoked potentials can be explained by temporal superposition of transient event-related

responses. *PloS ONE*, *6*(1), e14543. http://doi.org/10.1371/journal.pone.0014543

Carrasco, M. (2011). Visual attention : The past 25 years. *Vision Reseach*, *51*(13), 1484–1525. http://doi.org/10.1016/j.visres.2011.04.012.Visual

Carrasco, M. (2014). Spatial Attention: Perceptual modulation. *The Oxford Handbook of Attention*, 183–230.

Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*(8), 414–421. http://doi.org/10.1016/j.tics.2014.04.012.Frontal

Chakravarthi, R., & VanRullen, R. (2012). Conscious updating is a rhythmic process. *Proceedings of the National Academy of Sciences*, *109*(26), 10599–10604. http://doi.org/10.1073/pnas.1121622109

Childers, D. G., & Perry, N. W. (1971). Alpha-like activity in vision. *Brain Research*, *25*, 1–20.

Ciganek, L. (1969). Variability of the human visual evoked potential: Normative Data. *Electroencephalography and Clinical Neurophysiology*, *27*, 35–42. Retrieved from http://www.sciencedirect.com/science/article/pii/0013469469901060

Cohen, M. X. (2014). Fluctuations in Oscillation Frequency Control Spike Timing. *The Journal of Neuroscience*, *34*(27), 8988–8998. http://doi.org/10.1523/JNEUROSCI.0261-14.2014

Cohen, M. X. (2016). Rigor and replication in time-frequency analyses of cognitive electrophysiology data. *International Journal of Psychophysiology*, 1–8. http://doi.org/10.1016/j.ijpsycho.2016.02.001

Cohen, M. X., & Donner, T. H. (2013). Midfrontal conflict-related theta-band power reflects neural oscillations that predict behavior. *Journal of Neurophysiology*, *110*(12), 2752–2763. http://doi.org/10.1152/jn.00479.2013

Cosmelli, D., Lopez, V., Lachaux, J.-P., Lopez-Calderon, J., Renault, B., Martinerie, J., & Aboitiz, F. (2011). Shifting visual attention away from fixation is specifically associated with alpha band activity over ipsilateral parietal regions. *Psychophysiology*, *48*(3), 312–322. http://doi.org/10.1111/j.1469-8986.2010.01066.x

Crosse, M. J., Di Liberto, G. M., Bednar, A., & Lalor, E. C. (2016). The Multivariate Temporal Response Function (mTRF) Toolbox: A MATLAB Toolbox for Relating Neural Signals to Continuous Stimuli. *Frontiers in Human Neuroscience*, *10*(November), 1–14. http://doi.org/10.3389/fnhum.2016.00604

Cudeiro, J., & Sillito, A. M. (2006). Looking back: corticothalamic feedback and early visual processing. *Trends in Neurosciences*, *29*(6), 298–306. http://doi.org/10.1016/j.tins.2006.05.002

Dawson, G. D. (1954). A summation technique for the detection of small evoked potentials. *Electroencephalography and Clinical Neurophysiology*, *6*, 65–84. http://doi.org/http://dx.doi.org/10.1016/0013-4694(54)90007-3

de Graaf, T. A., Gross, J., Paterson, G., Rusch, T., Sack, A. T., & Thut, G. (2013). Alpha-Band Rhythms in Visual Task Performance: Phase-Locking by Rhythmic Sensory Stimulation. *PloS ONE*, *8*(3), 29–32. http://doi.org/10.1371/journal.pone.0060035

de Graaf, T. A., Koivisto, M., Jacobs, C., & Sack, A. T. (2014). The chronometry of visual perception: Review of occipital TMS masking studies. *Neuroscience and Biobehavioral Reviews*, *45*, 295–304. http://doi.org/10.1016/j.neubiorev.2014.06.017

Delorme, A., & Makeig, S. (2004). EEGLAB : an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21.

Di Russo, F., Martínez, A., Sereno, M. I., Pitzalis, S., & Hillyard, S. A. (2002). Cortical sources of the early components of the visual evoked potential. *Human Brain Mapping*, *15*(2), 95–111. http://doi.org/10.1002/hbm.10010

Di Russo, F., Pitzalis, S., Spitoni, G., Aprile, T., Patria, F., Spinelli, D., & Hillyard, S. A. (2005). Identification of the neural sources of the pattern-reversal VEP. *NeuroImage*, *24*(3), 874–886. http://doi.org/10.1016/j.neuroimage.2004.09.029

Doesburg, S. M., Roggeveen, A. B., Kitajo, K., & Ward, L. M. (2008). Large-scale gamma-band phase synchronization and selective attention. *Cerebral Cortex*, *18*(2), 386–396. http://doi.org/10.1093/cercor/bhm073

Douglas, R. J., & Martin, K. A. C. (2004). Neuronal Circuits of the Neocortex. *Annual Review of Neuroscience*, *27*(1), 419–451.

http://doi.org/10.1146/annurev.neuro.27.070203.144152

Douglas, R. J., & Martin, K. A. C. (2012). Behavioral architecture of the cortical sheet. *Current Biology*, *22*(24), R1033–R1038. http://doi.org/10.1016/j.cub.2012.11.017

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. http://doi.org/10.1523/JNEUROSCI.4795-10.2011

Dubois, J. (2011). *Perceiving the world under the strobe of attention : psychophysical and electroencephalographical investigations*.

Dugué, L., Marque, P., & VanRullen, R. (2011). The Phase of Ongoing Oscillations Mediates the Causal Relation between Brain Excitation and Visual Perception. *The Journal of Neuroscience*, *31*(33), 11889–11893. http://doi.org/10.1523/JNEUROSCI.1161-11.2011

Dugué, L., Marque, P., & VanRullen, R. (2015). Theta oscillations modulate attentional search performance periodically. *Journal of Cognitive Neuroscience*, *27*(5), 945–958. http://doi.org/10.1162/jocn

Dugué, L., & VanRullen, R. (2017). Transcranial magnetic stimulation reveals intrinsic perceptual and attentional rhythms. *Frontiers in Neuroscience*, *11*(154), 1–7. http://doi.org/10.3389/fnins.2017.00154

Engel, A. K., & Singer, W. (2001). Temporal binding and the neural correlates of sensory awareness. *Trends in Cognitive Sciences*, *5*(1), 16–25. http://doi.org/10.1016/S1364-6613(00)01568-0

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. http://doi.org/10.1016/j.cogbrainres.2004.03.009

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierachical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*(1), 1–47. http://doi.org/10.1093/cercor/1.1.1

Fellner, M. C., Volberg, G., Mullinger, K. J., Goldhacker, M., Wimber, M., Greenlee, M. W., & Hanslmayr, S. (2016). Spurious correlations in simultaneous EEG-fMRI driven by in-scanner movement. *NeuroImage*, *133*, 354–366. http://doi.org/10.1016/j.neuroimage.2016.03.031

Foxe, J. J., & Snyder, A. C. (2011). The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention. *Frontiers in Psychology*, *2*(JUL), 1–13. http://doi.org/10.3389/fpsyg.2011.00154

Frey, H.-P., Kelly, S. P., Lalor, E. C., & Foxe, J. J. (2010). Early Spatial Attentional Modulation of Inputs to the Fovea. *The Journal of Neuroscience*, *30*(13), 4547–4551. http://doi.org/10.1523/JNEUROSCI.5217-09.2010

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480. http://doi.org/10.1016/j.tics.2005.08.011

Fries, P. (2015). Rhythms for Cognition: Communication through Coherence. *Neuron*, *88*(1), 220–235. http://doi.org/10.1016/j.neuron.2015.09.034

Fries, P., Nikolić, D., & Singer, W. (2007). The gamma cycle. *Trends in Neurosciences*, *30*(7), 309–316. http://doi.org/10.1016/j.tins.2007.05.005

Gonçalves, N. R., Whelan, R., Foxe, J. J., & Lalor, E. C. (2014). Towards obtaining spatiotemporally precise responses to continuous sensory stimuli in humans: A general linear modeling approach to EEG. *NeuroImage*, *97*, 196–205. http://doi.org/10.1016/j.neuroimage.2014.04.012

Gulbinaite, R., İlhan, B., & VanRullen, R. (2017). The Triple-Flash Illusion Reveals a Driving Role of Alpha-Band Reverberations in Visual Perception. *The Journal of Neuroscience*, *37*(30), 7219–7230. http://doi.org/10.1523/JNEUROSCI.3929-16.2017

Gulbinaite, R., van Rijn, H., & Cohen, M. X. (2014). Fronto-parietal network oscillations reveal relationship between working memory capacity and cognitive control. *Frontiers in Human Neuroscience*, *8*(September), 761. http://doi.org/10.3389/fnhum.2014.00761

Haegens, S., Barczak, A., Musacchia, G., Lipton, M. L., Mehta, A. D., Lakatos, P., & Schroeder, C. E. (2015). Laminar Profile and Physiology of the alpha Rhythm in Primary Visual , Auditory , and Somatosensory Regions of Neocortex. *The Journal of Neuroscience*, *35*(42), 14341–14352. http://doi.org/10.1523/JNEUROSCI.0600-15.2015

Haegens, S., Cousijn, H., Wallis, G., Harrison, P. J., & Nobre, A. C. (2014). Inter- and intra-individual variability in alpha peak frequency. *NeuroImage*, *92*, 46–55. http://doi.org/10.1016/j.neuroimage.2014.01.049

Haegens, S., Nácher, V., Luna, R., Romo, R., & Jensen, O. (2011). α - Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceedings of the National Academy of Sciences*, *108*(48), 19377–19382. http://doi.org/10.1073/pnas.1117190108

Halbleib, A., Gratkowski, M., Schwab, K., Ligges, C., Witte, H., & Haueisen, J. (2012). Topographic Analysis of Engagement and Disengagement of Neural Oscillators in Photic Driving. *Journal of Clinical Neurophysiology*, *29*(1), 33–41. http://doi.org/10.1097/WNP.0b013e318246ad6e

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage*, *37*(4), 1465–1473. http://doi.org/10.1016/j.neuroimage.2007.07.011

Hanslmayr, S., Gross, J., Klimesch, W., & Shapiro, K. (2011). The role of α oscillations in temporal attention. *Brain Research Reviews*, *67*(1–2), 331–343. http://doi.org/10.1016/j.brainresrev.2011.04.002

Hanslmayr, S., Klimesch, W., Sauseng, P., Gruber, W. R., Doppelmayr, M., Freunberger, R., & Pecherstorfer, T. (2005). Visual discrimination performance is related to decreased alpha amplitude but increased phase locking. *Neuroscience Letters*, *375*(1), 64–68. http://doi.org/10.1016/j.neulet.2004.10.092

Hanslmayr, S., Volberg, G., Wimber, M., Dalal, S. S., & Greenlee, M. W. (2013). Prestimulus oscillatory phase at 7 Hz gates cortical information flow and visual perception. *Current Biology*, *23*(22), 2273–2278. http://doi.org/10.1016/j.cub.2013.09.020

Harter, M. R. (1967). Excitability cycles and cortical scanning: a review of two hypotheses of central intermittency in perception. *Psychological Bulletin*, *68*(1), 47–58. http://doi.org/10.1037/h0024725

Helfrich, R. F., Schneider, T. R., Rach, S., Trautmann-Lengsfeld, S. A., Engel, A. K., & Herrmann, C. S. (2014). Entrainment of brain oscillations by transcranial alternating current stimulation. *Current Biology*, *24*(3), 333–339. http://doi.org/10.1016/j.cub.2013.12.041

Herring, J. D., Thut, G., Jensen, O., & Bergmann, T. O. (2015). Attention Modulates TMS-Locked Alpha Oscillations in the Visual Cortex. *Journal of Neuroscience*, *35*(43),

14435–14447. http://doi.org/10.1523/JNEUROSCI.1833-15.2015

Herrmann, C. S. (2001). Human EEG responses to 1-100 Hz flicker: resonance phenomena in visual cortex and their potential correlation to cognitive phenomena. *Experimental Brain Research*, *137*(3–4), 346–353. http://doi.org/10.1007/s002210100682

Herrmann, C. S., Strüber, D., Helfrich, R. F., & Engel, A. K. (2016). EEG oscillations: From correlation to causality. *International Journal of Psychophysiology*, *103*(May), 12–21. http://doi.org/https://doi.org/10.1016/j.ijpsycho.2015.02.003

Hillyard, S. A., Teder-Sälejärvi, W. A., & Münte, T. F. (1998). Temporal dynamics of early perceptual processing. *Current Opinion in Neurobiology*, *8*(2), 202–210. http://doi.org/10.1016/S0959-4388(98)80141-4

Hubel, D. G. H., & Wiesel, T. T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, *148*(3), 574. Retrieved from http://jp.physoc.org/content/148/3/574.full.pdf

Huster, R. J., Debener, S., Eichele, T., & Herrmann, C. S. (2012). Methods for Simultaneous EEG-fMRI: An Introductory Review. *Journal of Neuroscience*, *32*(18), 6053–6060. http://doi.org/10.1523/JNEUROSCI.0447-12.2012

Iemi, L., Chaumon, M., Crouzet, S. M., & Busch, N. A. (2017). Spontaneous Neural Oscillations Bias Perception by Modulating Baseline Excitability. *The Journal of Neuroscience*, *37*(4), 807–819. http://doi.org/10.1523/JNEUROSCI.1432-16.2017

Ilhan, B., & VanRullen, R. (2012). No Counterpart of Visual Perceptual Echoes in the Auditory System. *PLoS ONE*, *7*(11). http://doi.org/10.1371/journal.pone.0049287

Jacobs, J., Kahana, M. J., Ekstrom, A. D., & Fried, I. (2007). Brain Oscillations Control Timing of Single-Neuron Activity in Humans. *Journal of Neuroscience*, *27*(14), 3839–3844. http://doi.org/10.1523/JNEUROSCI.4636-06.2007

Jensen, O., Bonnefond, M., & VanRullen, R. (2012). An oscillatory mechanism for prioritizing salient unattended stimuli. *Trends in Cognitive Sciences*, *16*(4), 200–205. http://doi.org/10.1016/j.tics.2012.03.002

Jensen, O., & Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences*, *11*(7), 267–269. http://doi.org/10.1016/j.tics.2007.05.003

Jensen, O., Gips, B., Bergmann, T. O., & Bonnefond, M. (2014). Temporal coding organized by coupled alpha and gamma oscillations prioritize visual processing. *Trends in Neurosciences*, *37*(7), 357–369. http://doi.org/10.1016/j.tins.2014.04.001

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity : gating by inhibition. *Frontiers in Human Neuroscience*, *4*(November), 1–8. http://doi.org/10.3389/fnhum.2010.00186

Jia, J., Liu, L., Fang, F., & Luo, H. (2017). Sequential sampling of visual objects during sustained attention. *PLOS Biology*, *15*(6), e2001903. http://doi.org/10.1371/journal.pbio.2001903

Kanai, R., Chaieb, L., Antal, A., Walsh, V., & Paulus, W. (2008). Frequency-Dependent Electrical Stimulation of the Visual Cortex. *Current Biology*, *18*(23), 1839–1843. http://doi.org/10.1016/j.cub.2008.10.027

Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. 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. http://doi.org/10.1152/jn.01234.2005

Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Research Reviews*, *29*(2–3), 169–95. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10209231

Klimesch, W., Fellinger, R., & Freunberger, R. (2011). Alpha oscillations and early stages of visual encoding. *Frontiers in Psychology*, *2*(118), 1–11. http://doi.org/10.3389/fpsyg.2011.00118

Klimesch, W., Hanslmayr, S., Sauseng, P., Gruber, W. R., & Doppelmayr, M. (2007). P1 and Traveling Alpha Waves: Evidence for Evoked Oscillations. *Journal of Neurophysiology*, *97*(2), 1311–1318. http://doi.org/10.1152/jn.00876.2006

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition-timing hypothesis. *Brain Research Reviews*, *3*(1), 63–88. http://doi.org/10.1016/j.brainresrev.2006.06.003

Krauzlis, R. J., Lovejoy, L. P., & Zénon, A. (2013). Superior Colliculus and Visual Spatial Attention. *Annual Review of Neuroscience*, *36*(1), 165–182.

http://doi.org/10.1146/annurev-neuro-062012-170249

Kuffler, S. W. (1953). Discharge Patterns and Functional organisation of Mammalian Retina. *Journal of Neurophysiology*, *16*(1), 37–68. Retrieved from http://jn.physiology.org/content/16/1/37.article-info

Lakatos, P., Shah, A. S., Knuth, K. H., Ulbert, I., Karmos, G., & Schroeder, C. E. (2005). An Oscillatory Hierarchy Controlling Neuronal Excitability and Stimulus Processing in the Auditory Cortex. *Journal of Neurophysiology*, *94*, 1904–1911. http://doi.org/10.1152/jn.00263.2005.

Lalor, E. C. (2009). Modeling the human visual system using the white-noise approach. In *Proceeding of the 4th International IEEE/EMBS Conference on Neural Engineering* (pp. 589–592). http://doi.org/10.1109/NER.2009.5109365

Lalor, E. C., & Foxe, J. J. (2009). Visual evoked spread spectrum analysis (VESPA) responses to stimuli biased towards magnocellular and parvocellular pathways. *Vision Research*, *49*(1), 127–133. http://doi.org/10.1016/j.visres.2008.09.032

Lalor, E. C., Kelly, S. P., & Foxe, J. J. (2012). Generation of the VESPA response to rapid contrast fluctuations is dominated by striate cortex: Evidence from retinotopic mapping. *Neuroscience*, *218*, 226–234. http://doi.org/10.1016/j.neuroscience.2012.05.067

Lalor, E. C., Kelly, S. P., Pearlmutter, B. A., Reilly, R. B., & Foxe, J. J. (2007). Isolating endogenous visuo-spatial attentional effects using the novel visual-evoked spread spectrum analysis (VESPA) technique. *European Journal of Neuroscience*, *26*, 3536–3542. http://doi.org/10.1111/j.1460-9568.2007.05968.x

Lalor, E. C., Lucan, J. N., & Foxe, J. J. (2009). Estimation of the impulse response of the visual system using stochastic modulation of stimulus spatial frequency. *2009 4th International IEEE/EMBS Conference on Neural Engineering, NER '09*, 593–596. http://doi.org/10.1109/NER.2009.5109366

Lalor, E. C., Pearlmutter, B. A., Reilly, R. B., McDarby, G., & Foxe, J. J. (2006). The VESPA: a method for the rapid estimation of a visual evoked potential. *NeuroImage*, *32*(4), 1549–1561. http://doi.org/10.1016/j.neuroimage.2006.05.054

Lalor, E. C., Power, A. J., Reilly, R. B., & Foxe, J. J. (2009). Resolving Precise Temporal Processing Properties of the Auditory System Using Continuous Stimuli. *Journal of*

Neurophysiology, *102*(1), 349–359. http://doi.org/10.1152/jn.90896.2008

Lalor, E. C., Reilly, R. B., Member, S., Pearlmutter, B. A., & Foxe, J. J. (2006). A Spectrum of Colors : Investigating the Temporal Frequency Characteristics of the Human Visual System Using a System Identi fi cation Approach. *EMBS Annual International Conference*, 3720–3723.

Lange, J., Keil, J., Schnitzler, A., van Dijk, H., & Weisz, N. (2014). The role of alpha oscillations for illusory perception. *Behavioural Brain Research*, *271*, 294–301. http://doi.org/10.1016/j.bbr.2014.06.015

Lansing, R. W. (1957). Relation of brain and tremor rhythms to visual reaction time. *Electroencephalography and Clinical Neurophysiology*, *9*(3), 497–504. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/13447855

Lansing, R. W., & Barlow, J. S. (1972). Rhythmic after-activity to flashes in relation to the background alpha which precedes and follows the photic stimuli. *Electroencephalography and Clinical Neurophysiology*, *32*(2), 149–160. http://doi.org/10.1016/0013-4694(72)90137-X

Laufs, H., Kleinschmidt, A., Beyerle, A., Eger, E., Salek-Haddadi, A., Preibisch, C., & Krakow, K. (2003). EEG-correlated fMRI of human alpha activity. *NeuroImage*, *19*(4), 1463–1476. http://doi.org/10.1016/S1053-8119(03)00286-6

Lindsley, D. B. (1952). Psychological phenomena and the electroencephalogram. *Electroencephalography and Clinical Neurophysiology*, *4*, 443–456. http://doi.org/10.1016/0013-4694(52)90075-8

Lisman, J. E., & Idiart, M. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science*, *267*(5203), 1512–1515. http://doi.org/10.1126/science.7878473

Lisman, J. E., & Jensen, O. (2013). The θ-γ neural code. *Neuron*, *77*(6), 1002–1016. http://doi.org/10.1016/j.neuron.2013.03.007

Livingstone, M., & Hubel, D. (1988). Segregation of Depth: Form, Anatomy, Color, Physiology, and Movement, and Perception. *Science*, *240*(4853), 740–749. http://doi.org/10.1126/science.3283936

Lopes da Silva, F. H. (1991). Neural mechanisms underlying brain waves: from neural membranes to networks. *Electroencephalography and Clinical Neurophysiology*, *79*(2),

81–93. http://doi.org/10.1016/0013-4694(91)90044-5

Lopes da Silva, F. H. (2013). EEG and MEG: Relevance to neuroscience. *Neuron*, *80*(5), 1112–1128. http://doi.org/10.1016/j.neuron.2013.10.017

Lopes da Silva, F. H., Vos, J. E., Mooibroek, J., & van Rotterdam, A. (1980). Relative contributions of intracortical and thalamo-cortical processes in the generation of alpha rhythms, revealed by partial coherence analysis. *Electroencephalography and Clinical Neurophysiology*, *50*(5–6), 449–456. http://doi.org/10.1016/0013-4694(80)90011-5

Lorincz, M. L., Kékesi, K. A., Juhász, G., Crunelli, V., & Hughes, S. W. (2009). Temporal Framing of Thalamic Relay-Mode Firing by Phasic Inhibition during the Alpha Rhythm. *Neuron*, *63*(5), 683–696. http://doi.org/10.1016/j.neuron.2009.08.012

Lozano-Soldevilla, D., & Vanrullen, R. (2016). The hidden spatial dimension of alpha occipital EEG channels encode contralateral and ipsilateral visual space at distinct phases of the alpha cycle. *Journal of Vision*, *16*(12), 1226.

Luck, S. J. (2005a). An Introduction to Event-Related Potentials and Their Neural Origins. In *An introduction to the ERP technique* (pp. 1–50).

Luck, S. J. (2005b). *An Introduction to the Event-Related Potential Technique*.

Luck, S. J., Woodman, G. F., & Vogel, E. K. (2000). Event-related potential studies of attention. *Trends in Cognitive Sciences*, *4*(11), 432–440. http://doi.org/10.1016/S1364-6613(00)01545-X

Makeig, S., Westerfield, M., Jung, T. P., Enghoff, S., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2002). Dynamic Brain Sources of Visual Evoked Responses. *Science*, *295*, 690–694. http://doi.org/10.1126/science.1066168

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–90. http://doi.org/10.1016/j.jneumeth.2007.03.024

Marmarelis, P. Z., & Marmarelis, V. Z. (1978). The White Noise Method of System Identification. In *Analysis of physiological systems: the white-noise approach*. New York: Plenum Press.

Marmarelis, V. Z. (1980). Identification of nonlinear systems by use of nonstationary white-

noise inputs. *Appl. Math. Modelling*, *4*, 117–124.

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*, *29*(9), 2725–2732. http://doi.org/10.1523/JNEUROSCI.3963-08.2009

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*, 1–15. http://doi.org/10.3389/fpsyg.2011.00099

Mazaheri, A., & Jensen, O. (2010). Rhythmic pulsing: linking ongoing brain activity with evoked responses. *Frontiers in Human Neuroscience*, *4*(October), 177. http://doi.org/10.3389/fnhum.2010.00177

McCormick, D. A., & Bal, T. (1994). Sensory gating mechanisms of the thalamus. *Current Opinion in Neurobiology*, *4*, 550–556. http://doi.org/10.1016/0959-4388(94)90056-6

McLelland, D., Lavergne, L., & VanRullen, R. (2016). The phase of ongoing EEG oscillations predicts the amplitude of peri-saccadic mislocalization. *Scientific Reports*, *6*, 29335. http://doi.org/10.1038/srep29335

Mehta, M. R., Lee, a K., & Wilson, M. a. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, *417*(6890), 741–6. http://doi.org/10.1038/nature00807

Michalareas, G., Vezoli, J., Van Pelt, S., Schoffelen, J., Kennedy, H., & Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, *89*, 1–14. http://doi.org/10.1016/j.neuron.2015.12.018

Millett, D. (2001). Hans Berger: From Psychic Energy to the EEG. *Perspectives in Biology and Medicine*, *44*(4), 522–542. http://doi.org/10.1353/pbm.2001.0070

Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, *46*(3), 774–785. http://doi.org/10.1016/j.neuropsychologia.2007.10.005

Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., & Panzeri, S. (2008). Phase-of-Firing Coding of Natural Visual Stimuli in Primary Visual Cortex. *Current Biology*, *18*(5), 375–380. http://doi.org/10.1016/j.cub.2008.02.023

Moosmann, M., Ritter, P., Krastel, I., Brink, A., Thees, S., Blankenburg, F., … Villringer, A. (2003). Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. *NeuroImage*, *20*(1), 145–158. http://doi.org/10.1016/S1053-8119(03)00344-6

Murphy, J. W., Kelly, S. P., Foxe, J. J., & Lalor, E. C. (2012). Isolating early cortical generators of visual-evoked activity: A systems identification approach. *Experimental Brain Research*, *220*(2), 191–199. http://doi.org/10.1007/s00221-012-3129-1

Neri, P., & Heeger, D. J. (2002). Spatiotemporal mechanisms for detecting and identifying image features in human vision. *Nature Neuroscience*, *5*(8), 812–816. http://doi.org/10.1038/nn886

Neuenschwander, S., Castelo-Branco, M., & Singer, W. (1999). Synchronous oscillations in the cat retina. *Vision Research*, *39*(15), 2485–2497. http://doi.org/10.1016/S0042-6989(99)00042-5

Neuper, C., & Pfurtscheller, G. (2001). Event-related dynamics of cortical rhythms: Frequency-specific features and functional correlates. *International Journal of Psychophysiology*, *43*, 41–58. http://doi.org/10.1016/S0167-8760(01)00178-7

Nunez, P. L., & Srinivasan, R. (2006). *Electric Fields of the Brain: The Neurophysics of EEG*. *Oxford University Press* (2nd ed.). London, England: Oxford University Press. http://doi.org/10.1001/jama.1982.03320380071046

Nunn, C. M., & Osselton, J. W. (1974). The influence of the EEG alpha rhythm on the perception of visual stimuli. *Psychophysiology*, *11*(3), 294–303.

Osipova, D., Hermes, D., & Jensen, O. (2008). Gamma power is phase-locked to posterior alpha activity. *PLoS ONE*, *3*(12), 1–7. http://doi.org/10.1371/journal.pone.0003990

Panzeri, S., Brunel, N., Logothetis, N. K., & Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends in Neurosciences*, *33*(3), 111–120. http://doi.org/10.1016/j.tins.2009.12.001

Patrick, G. T. W. (1890). The psychology of prejudice. *The Popular Science Monthly*, *36*(March), 634.

Penttonen, M., & Buzsáki, G. (2003). Natural logarithmic relationship between brain oscillators. *Thalamus and Related Systems*, *2*(2), 145–152. http://doi.org/10.1016/S1472-

9288(03)00007-4

Petersen, S. E., & Posner, M. I. (2012). The Attention System of the Human Brain: 20 Years After. *Annual Reviews of Neuroscience*, *35*(1), 73–89. http://doi.org/10.1146/annurev-neuro-062111-150525

Pfurtscheller, G., & Aranibar, A. (1977). Event related cortical desynchronization detected by power measuremnets of scalp EEG. *Electroencephalography and Clinical Neurophysiology*, *42*(5), 817–826.

Pfurtscheller, G., & Lopes Da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, *110*, 1842–1857. http://doi.org/10.1016/S1388-2457(99)00141-8

Pfurtscheller, G., Stancák, A., & Neuper, C. (1996). Event-related synchronization (ERS) in the alpha band - an electrophysiological correlate of cortical idling: a review. *International Journal of Psychophysiology*, *24*(1–2), 39–46. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8978434

Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson, R., … Taylor, M. J. (2000). Guidelines for using human event-related potentials to study cognition: Recording standards and publication criteria. *Psychophysiology*, *37*(2), 127–152. http://doi.org/10.1111/1469-8986.3720127

Pitts, W., & McCulloch, W. S. (1947). How we know universals: the perception of auditory and visual forms. *The Bulletin of Mathematical Biophysics*, *9*, 127–147. http://doi.org/10.1007/BF02478291

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. http://doi.org/10.1080/00335558008248231

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, *109*(2), 160–174. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7381367

Power, A. J., Lalor, E. C., & Reilly, R. B. (2006). Can visual evoked potentials be used in biometric identification? *Conference Proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, *1*, 5575–5578.

http://doi.org/10.1109/IEMBS.2006.259493

Priebe, N. J., & Ferster, D. (2012). Mechanisms of Neuronal Computation in Mammalian Visual Cortex. *Neuron*, *75*(2), 194–208. http://doi.org/10.1016/j.neuron.2012.06.011

Regan, D. (1966, March). Some characteristics of average steady-state and transient responses evoked by modulated light. *Electroencephalography and Clinical Neurophysiology*. http://doi.org/10.1016/0013-4694(66)90088-5

Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *European Journal of Neuroscience*, *25*(2), 603–610. http://doi.org/10.1111/j.1460-9568.2007.05278.x

Rihs, T. A., Michel, C. M., & Thut, G. (2009). A bias for posterior alpha-band power suppression versus enhancement during shifting versus maintenance of spatial attention. *NeuroImage*, *44*(1), 190–199. http://doi.org/10.1016/j.neuroimage.2008.08.022

Ringach, D., & Shapley, R. (2004). Reverse correlation in neurophysiology. *Cognitive Science*, *28*(2), 147–166. http://doi.org/10.1016/j.cogsci.2003.11.003

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. *Cerebral Cortex*, *18*(9), 2010–2018. http://doi.org/10.1093/cercor/bhm229

Romei, V., Gross, J., & Thut, G. (2010). On the role of prestimulus alpha rhythms over occipito-parietal areas in visual input regulation: correlation or causation? *Journal of Neuroscience*, *30*(25), 8692–7. http://doi.org/10.1523/JNEUROSCI.0160-10.2010

Romei, V., Gross, J., & Thut, G. (2012). Sounds reset rhythms of visual cortex and corresponding human visual perception. *Current Biology*, *22*(9), 807–813. http://doi.org/10.1016/j.cub.2012.03.025

Romei, V., Rihs, T., Brodbeck, V., & Thut, G. (2008). Resting electroencephalogram alpha-power over posterior sites indexes baseline visual cortex excitability. *NeuroReport*, *19*(2), 203–208. http://doi.org/10.1097/WNR.0b013e3282f454c4

Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R., … Whittington, M. A. (2008). Temporal interactions between cortical rhythms. *Frontiers in Neuroscience*, *2*(2), 145–154. http://doi.org/10.3389/neuro.01.034.2008

Sakai, H. M., Naka, K., & Korenberg, M. J. (1988). White-noise analysis in visual neuroscience. *Visual Neuroscience*, *1*(3), 287–296. http://doi.org/10.1017/S0952523800001942

Samaha, J., Gosseries, O., & Postle, B. R. (2017). Distinct oscillatory frequencies underlie excitability of human occipital and parietal cortex. *The Journal of Neuroscience*, *37*(11), 3413–16. http://doi.org/10.1523/JNEUROSCI.3413-16.2017

Samaha, J., & Postle, B. R. (2015). The Speed of Alpha-Band Oscillations Predicts the Temporal Resolution of Visual Perception. *Current Biology*, *25*(22), 2985–2990. http://doi.org/10.1016/j.cub.2015.10.007

Samaha, J., Sprague, T. C., & Postle, B. R. (2016). Decoding and Reconstructing the focus of spatial attention from the topography of alpha band oscillations. *Journal of Cognitive Neuroscience*, *28*(8), 1090–1097. http://doi.org/10.1162/jocn_a_00955

Sauseng, P., & Klimesch, W. (2008). What does phase information of oscillatory brain activity tell us about cognitive processes? *Neuroscience & Biobehavioral Reviews*, *32*(5), 1001–13. http://doi.org/10.1016/j.neubiorev.2008.03.014

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., … Birbaumer, N. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, *22*(11), 2917–2926. http://doi.org/10.1111/j.1460-9568.2005.04482.x

Scheeringa, R., Koopmans, P. J., van Mourik, T., Jensen, O., & Norris, D. G. (2016). The relationship between oscillatory EEG activity and the laminar-specific BOLD signal. *Proceedings of the National Academy of Sciences*, *113*(24), 6761–6766. http://doi.org/10.1073/pnas.1522577113

Scheeringa, R., Petersson, K. M., Oostenveld, R., Norris, D. G., Hagoort, P., & Bastiaansen, M. C. M. (2009). Trial-by-trial coupling between EEG and BOLD identifies networks related to alpha and theta EEG power increases during working memory maintenance. *NeuroImage*, *44*(3), 1224–1238. http://doi.org/10.1016/j.neuroimage.2008.08.041

Schroeder, C. E., & Lakatos, P. (2008). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, *32*(1), 9–18. http://doi.org/10.1016/j.tins.2008.09.012

Schyns, P. G., Thut, G., & Gross, J. (2011). Cracking the code of oscillatory activity. *PLoS Biology*, *9*(5), e1001064. http://doi.org/10.1371/journal.pbio.1001064

Singer, W. (1999). Time as coding space? *Current Opinion in Neurobiology*, *9*(2), 189–194. http://doi.org/10.1016/S0959-4388(99)80026-9

Spaak, E., Bonnefond, M., Maier, A., Leopold, D. A., & Jensen, O. (2012). Layer-Specific Entrainment of Gamma-Band Neural Activity by the Alpha Rhythm in Monkey Visual Cortex. *Current Biology*, *22*(24), 2313–2318. http://doi.org/10.1016/j.cub.2012.10.020

Spaak, E., de Lange, F. P., & Jensen, O. (2014). Local entrainment of α oscillations by visual stimuli causes cyclic modulation of perception. *The Journal of Neuroscience*, *34*(10), 3536–44. http://doi.org/10.1523/JNEUROSCI.4385-13.2014

Spillmann, L. (2014). Receptive fields of visual neurons: The early years. *Perception*, *43*(11), 1145–1176. http://doi.org/10.1068/p7721

Steriade, M., Gloor, P., Llinás, R. R., Lopes da Silva, F. H., & Mesulam, M. M. (1990). Report of IFCN Committee on Basic Mechanisms: Basic mechanisms of cerebral rhythmic activities. *Electroencephalography and Clinical Neurophysiology*. http://doi.org/10.1016/0013-4694(90)90001-Z

Steriade, M., & Llinas, R. R. (1988). The Functional States of the Thalamus and the Associated Neuronal Interplay. *Physiological Reviews*, *68*(3), 331–370.

Steriade, M., McCormick, D. A., & Sejnowski, T. J. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science*, *262*(5134), 679–685. http://doi.org/10.1126/science.8235588

Stroud, J. M. (1955). The Fine Structure of Psychological Time. In *Information Theory in Psychology* (Vol. 138, pp. 175–207).

Survillo, W. W. (1961). Frequency of the "Alpha" Rhythm, Reaction Time and Age. *Nature*, *191*, 823–824.

Tallon-Baudry, C., & Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, *3*(4), 151–162. http://doi.org/10.1016/S1364-6613(99)01299-1

Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Pernier, J. (1996). Stimulus specificity of

phase-locked and non-phase-locked 40 Hz visual responses in human. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(13), 4240–9. http://doi.org/10.1016/j.neuropsychologia.2011.02.038

Thorpe, S. J. (1990). Spike arrival times: A highly efficient coding scheme for neural networks. In R. Eckmiller, G. Hartmann, & G. Hauske (Eds.), *Parallel processing in neural systems and computers* (pp. 91–94). North-Holland: Elsevier.

Thut, G., & Miniussi, C. (2009). New insights into rhythmic brain activity from TMS-EEG studies. *Trends in Cognitive Sciences*, *13*(4), 182–9. http://doi.org/10.1016/j.tics.2009.01.004

Thut, G., Miniussi, C., & Gross, J. (2012). The functional importance of rhythmic activity in the brain. *Current Biology*, *22*(16), R658-63. http://doi.org/10.1016/j.cub.2012.06.061

Thut, G., Nietzel, A., Brandt, S. a, & Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *The Journal of Neuroscience*, *26*(37), 9494–9502. http://doi.org/10.1523/JNEUROSCI.0875-06.2006

Thut, G., Schyns, P. G., & Gross, J. (2011). Entrainment of perceptually relevant brain oscillations by non-invasive rhythmic stimulation of the human brain. *Frontiers in Psychology*, *2*(JUL), 1–10. http://doi.org/10.3389/fpsyg.2011.00170

Thut, G., Veniero, D., Romei, V., Miniussi, C., Schyns, P., & Gross, J. (2011). Rhythmic TMS causes Local Entrainment of Natural Oscillatory Signatures. *Current Biology*, *21*(14), 1176–1185. http://doi.org/10.1016/j.cub.2011.05.049

Tiesinga, P., & Sejnowski, T. J. (2009). Cortical Enlightenment: Are Attentional Gamma Oscillations Driven by ING or PING? *Neuron*. http://doi.org/10.1016/j.neuron.2009.09.009

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. *Analysis of Visual Behavior*. http://doi.org/10.2139/ssrn.1353746

van Diepen, R. M., Cohen, M. X., Denys, D., & Mazaheri, A. (2015). Attention and Temporal Expectations Modulate Power, Not Phase, of Ongoing Alpha Oscillations. *Journal of Cognitive Neuroscience*, *27*(8), 1573–1586. http://doi.org/10.1162/jocn

van Dijk, H., Schoffelen, J.-M. J., Oostenveld, R., & Jensen, O. (2008). Prestimulus

oscillatory activity in the alpha band predicts visual discrimination ability. *The Journal of Neuroscience*, *28*(8), 1816–1823. http://doi.org/10.1523/JNEUROSCI.1853-07.2008

van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences*, *111*(40), 14332–41. http://doi.org/10.1073/pnas.1402773111

VanRullen, R. (2016a). How to evaluate phase differences between trial groups in ongoing electrophysiological signals. *Frontiers in Neuroscience*, *10*, 426. http://doi.org/10.3389/fnins.2016.00426-

VanRullen, R. (2016b). Perceptual Cycles. *Trends in Cognitive Sciences*, *20*(10), 723–735. http://doi.org/10.1016/j.tics.2016.07.006

VanRullen, R. (2016c). Perceptual Rhythms. In *Stevens' Handbook of Experimental Psychology* (pp. 1–51).

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*(April), article 60. http://doi.org/10.3389/fpsyg.2011.00060

VanRullen, R., Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *The Journal of Neuroscience*, *28*(1), 1–4. http://doi.org/10.1016/j.tins.2004.10.010

VanRullen, R., & Koch, C. (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences*, *7*(5), 207–213. http://doi.org/10.1016/S1364-6613(03)00095-0

VanRullen, R., & Macdonald, J. S. P. (2012). Perceptual echoes at 10 Hz in the human brain. *Current Biology*, *22*(11), 995–999. http://doi.org/10.1016/j.cub.2012.03.050

VanRullen, R., Zoefel, B., & Ilhan, B. (2014). On the cyclic nature of perception in vision versus audition. *Philosophical Transactions of the Royal Society B*, *369*, 20130214. http://doi.org/10.1098/rstb.2013.0214

Varela, F. J., Toro, A., John, E. R., & Schwartz, E. L. (1981). Perceptual framing and cortical alpha rhythm. *Neuropsychologia*, *19*(5), 675–686. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7312152

Wallisch, P., & Movshon, J. A. (2008). Structure and Function Come Unglued in the Visual

Cortex. *Neuron*, *60*(2), 195–197. http://doi.org/10.1016/j.neuron.2008.10.008

Wang, X. (2010). Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. *Physiological Reviews*, *90*(3), 1195–1268. http://doi.org/10.1152/physrev.00035.2008.Neurophysiological

Wässle, H. (2004). Parallel processing in the mammalian retina. *Nature Reviews Neuroscience*, *5*(10), 747–757. http://doi.org/10.1038/nrn1497

Watrous, A. J., Fell, J., Ekstrom, A. D., & Axmacher, N. (2015). More than spikes: Common oscillatory mechanisms for content specific neural representations during perception and memory. *Current Opinion in Neurobiology*, *31*, 33–39. http://doi.org/10.1016/j.conb.2014.07.024

Watson, A. B., & Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*(2), 113–120. http://doi.org/10.3758/BF03202828

Whittingstall, K., & Logothetis, N. K. (2009). Frequency-Band Coupling in Surface EEG Reflects Spiking Activity in Monkey Visual Cortex. *Neuron*, *64*(2), 281–289. http://doi.org/10.1016/j.neuron.2009.08.016

Woodman, G. F. (2010). A brief introduction to the use of event-related potentials in studies of perception and attention. *Attention, Perception & Psychophysics*, *72*(8), 2031–2046. http://doi.org/10.3758/APP

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*, *20*(RC63), 1–6.

Wyart, V., & Tallon-baudry, C. (2008). Neural Dissociation between Visual Awareness and Spatial Attention. *Journal of Neuroscience*, *28*(10), 2667–2679. http://doi.org/10.1523/JNEUROSCI.4748-07.2008

Wyart, V., & Tallon-Baudry, C. (2009). How Ongoing Fluctuations in Human Visual Cortex Predict Perceptual Awareness: Baseline Shift versus Decision Bias. *Journal of Neuroscience*, *29*(27), 8715–8725. http://doi.org/10.1523/JNEUROSCI.0962-09.2009

Yamagishi, N., Goda, N., Callan, D. E., Anderson, S. J., & Kawato, M. (2005). Attentional shifts towards an expected visual target alter the level of alpha-band oscillatory activity

in the human calcarine cortex. *Cognitive Brain Research*, *25*(3), 799–809. http://doi.org/10.1016/j.cogbrainres.2005.09.006

Zaehle, T., Lenz, D., Ohl, F. W., & Herrmann, C. S. (2010). Resonance phenomena in the human auditory cortex: Individual resonance frequencies of the cerebral cortex determine electrophysiological responses. *Experimental Brain Research*, *203*(3), 629–635. http://doi.org/10.1007/s00221-010-2265-8

Zoefel, B., & VanRullen, R. (2017). Oscillatory Mechanisms of Stimulus Processing and Selection in the Visual and Auditory Systems : State-of-the-Art, Speculations and Suggestions. *Frontiers in Neuroscience*, *11*(296), 1–13. http://doi.org/10.3389/fnins.2017.00296

Zumer, J. M., Scheeringa, R., Schoffelen, J. M., Norris, D. G., & Jensen, O. (2014). Occipital Alpha Activity during Stimulus Processing Gates the Information Flow to Object-Selective Cortex. *PLoS Biology*, *12*(10). http://doi.org/10.1371/journal.pbio.1001965