



HAL
open science

The boundaries of perception : a conceptual engineering perspective

Géraldine Carranante

► **To cite this version:**

Géraldine Carranante. The boundaries of perception : a conceptual engineering perspective. Philosophy. Université Paris sciences et lettres, 2020. English. NNT : 2020UPSLE073 . tel-03662720

HAL Id: tel-03662720

<https://theses.hal.science/tel-03662720>

Submitted on 9 May 2022

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

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

THÈSE DE DOCTORAT
DE L'UNIVERSITÉ PSL
Préparée à Ecole Normale Supérieure

**The boundaries of perception: a conceptual engineering
perspective**

Soutenue par

Géraldine Carranante

Le 1 décembre 2020

École doctorale n°ED540

**Lettres, arts, sciences hu-
maines et sociales**

Spécialité

Philosophie

Composition du jury :

Paul Egré DR, ENS/IJN	<i>Président</i>
Denis Forest Professeur, Paris 1 Panthéon Sorbonne	<i>Rapporteur</i>
Susanna Siegel Professor, Harvard University	<i>Rapporteur</i>
Sebastian Watzl Associate professor, University of Oslo	<i>Examineur</i>
Frédérique de Vignemont DR, ENS/CNRS	<i>Directrice de thèse</i>

ECOLE NORMALE SUPÉRIEURE
Institut Jean Nicod, Département d'études cognitives, ENS, EHESS,
CNRS, PSL University

THE BOUNDARIES OF PERCEPTION: A
CONCEPTUAL ENGINEERING PERSPECTIVE

PHD THESIS

submitted in partial fulfillment of the requirements for the degree of

Doctor in Philosophy

by

GÉRALDINE CARRANANTE

Under the supervision of
FRÉDÉRIQUE DE VIGNEMONT
Institut Jean Nicod, CNRS

Abstract

Scientists try to uncover the mental mechanisms that allow us to see, feel, touch, etc. They are helped by philosophers, who elaborate characterizations of perception, of its function and nature. Yet, perception as a scientific object, is quite vague. In particular, perception is hard to circumscribe in the architecture of the mind.

The aim of this dissertation is to shed light on the kind of scientific object perception is, through a philosophical work based on analytical philosophy of mind, philosophy of science and on current philosophical developments concerning perception. How could we know that a specific phenomenon is a perceptual one? How does the scientist know that she is studying perception and not another activity of the mind? What are the boundaries of perception?

These questions have direct implications in everyday scientific practice, especially in the interpretation of experimental results. They also pertain to philosophical debates such as the cognitive penetrability of perception.

In this dissertation, I first explore the usual method for determining the boundaries of perception, based on the idea that perception is a natural kind of mental states. I then elaborate another strategy called "conceptual engineering". In this strategy, I study perception as a scientific concept. I show that there are at least four different concepts of perception. I argue in favor of conceptual pluralism, i.e. that these concepts of perception are all legitimate. This pluralism is part of a classical process commonly undergone by scientific concepts - conceptual fragmentation – in which a scientific concept is fragmented into several sub-concepts.

In a third part, I explore metaphysical foundations for the boundaries of perceptions. I show that there are several alternative ways to metaphysically ground the boundaries of perception. The choice between them influence how empirical results are interpreted.

Finally, I argue that the concept of perception should be today considered an organizational concept in cognitive science, whose main function is to guide, coordinate and integrate interdisciplinary research about perception. Recognizing this specific place of the concept perception in scientific and philosophical investigations about the mind would contribute to enrich discussions, as well as to avoid ill-posed questions and fruitless debates.

Résumé

Les scientifiques tentent de découvrir les mécanismes mentaux qui nous permettent de voir, de sentir, de toucher, etc. Les théories sur la perception abondent et pourtant, la perception, en tant qu'objet scientifique, reste vague. Comment pouvons-nous savoir qu'un phénomène spécifique est un phénomène perceptuel ? Comment les scientifiques savent-ils qu'ils étudient la perception et non une autre activité de l'esprit ? Quelles sont les limites de la perception ?

Ces questions ont des implications directes dans la pratique scientifique en laboratoire, en particulier dans l'interprétation des résultats expérimentaux. Elles sont également au cœur de plusieurs débats philosophiques tels que la pénétrabilité cognitive de la perception. L'objectif de cette thèse est de contribuer au débat sur les frontières de la perception par une étude du concept de perception au sein des sciences cognitives.

J'explore en premier lieu la stratégie utilisée habituellement pour déterminer les limites de la perception, qui repose sur le présupposé que la perception constitue une espèce naturelle d'états mentaux. Dans un second temps, j'élabore une stratégie appelée « ingénierie conceptuelle ». J'y étudie la perception en tant que *concept scientifique*. Je montre qu'il existe quatre concepts différents de perception. Je défends que ces quatre concepts sont légitimes. Ce pluralisme conceptuel s'inscrit dans un processus de « fragmentation conceptuelle », processus courant lors de la maturation d'un domaine scientifique.

Dans une troisième partie, j'étudie les fondements métaphysiques des limites de la perception. Je montre qu'il existe plusieurs manières de concevoir métaphysiquement les frontières de la perception, et que ce choix a une influence sur la manière dont les résultats empiriques sont interprétés.

Enfin, dans une dernière partie, je défend que le concept de perception joue aujourd'hui un rôle de « concept organisationnel », dont la fonction principale est de guider, coordonner et intégrer la recherche interdisciplinaire sur la perception. Reconnaître la place spécifique du concept de perception en science permet d'éviter les débats infructueux et d'enrichir les discussions scientifiques en présentant certaines théories comme complémentaires plutôt que rivales.

Acknowledgements

First of all, I would like to express my gratitude towards my supervisor, Frédérique de Vignemont, for her unfailing support, enthusiasm and availability. Thank you for your guidance in writing this work, as well as your kindness towards the events of life. I am amazed to have given birth in such a short time to a dissertation and a child, and I would never be able to thank you enough for your help. You found the subtle balance between encouragement and patience for making this possible in such good conditions. I am also grateful to Uriah Kriegel, my previous supervisor, who gave me solid methodology and knowledge in philosophy of mind during our two years of work together.

I would like to thank Paul Egré and Sebastian Watzl for their help during my annual PhD committees. Our three meetings during these four years have structured my work and your comments have been highly appreciated. My appreciation also goes to Ned Block and Susanna Siegel, who invited me respectively at NYU and Harvard during my PhD. I thank both of you for your comments and discussions, they greatly improved my work.

I would like to thank the Institut Jean Nicod community : the doctoral students, the researchers and the support staff, especially Louis Rouillé, Roberto Casati, Jérôme Dokic, Nathalie Evin, Vincent Gaudefroy, Anna Giustina, Michael Murez, François Le Corre and Takuya Niikawa. You have all contributed to this work in one way or another. Thank you. The list could be much longer if I include all my colleagues and friends who gave me joy, comfort, and endless debates to feed my reflection.

Between June 2017 and April 2019, parts of this work have been presented at seminars and conferences at the Institut Jean Nicod, the University of Turin and Harvard University. I would like to thank the participants for their comments. I am especially grateful to the participants of the workshop at the Ecole Normale Supérieure, on April 1, 2019 for their insightful comments and our discussions : Jake Quilty-Dunn, Joulia Smortchkova, Fiona Macpherson and Jérôme Dokic.

I would like to say a special thank you to Alan Fincke, Jean Levaux, Anya Olson and all the participants of the Tuesday night dinner at Cambridge. Thank you for offering me a home for a few months on the other side of the world.

I gratefully acknowledge the financial support of the Institut Jean Nicod, the department of cognitive studies of the Ecole Normale Supérieure and PSL.

Pour finir, je voudrais remercier mes amis et ma famille, sans qui rien n'aurait été et ne sera jamais possible. Merci à Kevin, Elodie, Basile, Victor et Adèle pour être de si bons amis. Merci Pierre d'être ma maison. Merci maman.

Contents

Abstract	i
Résumé	iii
Acknowledgements	v
1 Introduction	1
1.1 The boundary question	2
1.2 A much-needed answer	3
1.3 The challenges of the boundary question	8
1.4 Plan	10
I Perception as a natural kind: disagreements and springboard	15
2 The natural-kind strategy	17
2.1 Perceptual as a natural kind	17
2.2 The strategy	20
2.3 The alleged hallmarks of perception	24
2.3.1 Preliminaries	24
2.3.2 The representational hallmarks	25
2.3.3 The architectural hallmark	34
2.3.4 The function of perception	37
2.4 A hard choice	38
3 Object-level oriented disagreements	43
3.1 Two kinds of empirical disagreements	44
3.1.1 Disagreement on properties	44
3.1.2 Disagreement on second-order properties	47
3.2 Disagreement on natural kinds	49
3.3 A threat for the natural-kind strategy	51
4 Representation-level oriented disagreement	53
4.1 Disagreement on extension	55
4.2 Metalinguistic negotiation	56
4.3 A Normative Approach	61
4.4 Conclusion	68

II	Conceptual Engineering	69
5	The concepts of perception	71
5.1	A <i>scientific</i> concept	71
5.1.1	Concept-as-theoretical-tool	71
5.1.2	Perception: A <i>scientific</i> concept	72
5.1.3	Objection: what about the ordinary concept of perception?	73
5.1.4	Some consequences	77
5.2	Individuation of concepts	78
5.2.1	Brigandt's framework	78
5.2.2	The challenge of sameness of topic	80
5.3	The multitude of scientific concepts of perception	83
5.3.1	The epistemic goals of PERCEPTION	83
5.3.2	Independence and synergies	91
5.4	Conclusion	97
6	Evaluation, Pluralism, Fragmentation	99
6.1	Evaluation	100
6.1.1	Concept Utility	101
6.1.2	Concept utility of versions of k-PERCEPTION	103
6.1.3	Comments and limits	113
6.2	Pluralism and Fragmentation	120
6.2.1	Pluralism and its two sources	121
6.2.2	A decent pluralism	123
6.2.3	Fragmented PERCEPTION	126
III	Joint-in-Nature	131
7	Three ways to carve nature at its joints	133
7.1	Ontological boundaries: preliminaries	133
7.1.1	Conceptual vs Ontological boundaries	133
7.1.2	Realist framework and the joint-in-nature hypothesis	136
7.2	Joint-in-nature	138
7.2.1	Literature	138
7.2.2	A metaphysical tool	140
7.3	Three joints	143
7.3.1	Taxonomical joint-in-nature	144
7.3.2	Mereological joint-in-nature	146
7.3.3	Focal joint-in-nature	149
7.4	Conclusion	155
8	Consequences on penetrability	157
8.1	Penetrability of joints-in-nature	158
8.2	Maintaining Impenetrability	161
8.3	Resistance to Eliminativism	164
8.4	Conclusion	166

IV	Against Eliminativism	169
9	The scientific role of perception	171
9.1	Organizational Concept	172
9.1.1	PERCEPTION is an organizational concept	172
9.1.2	The epistemic contribution of organizational concepts	174
9.1.3	Instrumental and intrinsic epistemic value	175
9.2	Eliminativist arguments	176
9.2.1	The risk of confusion	177
9.2.2	You (damn) conceptual luxuriance!	180
9.2.3	Natural kinds... again?	182
10	Conclusion	185
	Appendices	199
	Evaluation P_i with 10 items	201
	Evaluation P_i with 15 items	205
	Utility random partition	207

Chapter 1

Introduction

Cognitive science is currently revealing little by little how we represent the world through our senses. Scientific optimism leads us to believe that at the end of the road, we will understand how the mind works, paving the way to a better understanding of ourselves. Science would therefore become a prerequisite for any philosophy that prides itself on talking about the mind. This conclusion is too quick, though. Empirical research about the mind - including perception - is actually far from being free of philosophical questioning. It is filled with methodological, conceptual, and foundational questions, which cannot be tackled without philosophical reflection on the objects that science seeks to explain. When perception is at stake, science tries to uncover the mental mechanisms that allow us to see, feel, touch, etc. But what precisely are these phenomena?

We have an intuitive idea of what perception is (or perception in a specific modality, such as vision). We see when we open our eyes and the world presents itself in front of us at once. We hear when a sound reaches our ears. These abilities allow us to obtain information about the world: I see my son falling, I hear him screaming, I know he has been injured.

These insights are not enough to effectively feed rigorous scientific research. What does constitute the ability to see? Did I really see my son falling? Or did I see something, that I interpret as my son falling? Did I infer from my visual and auditory experience that my son has been injured? Or did I see him being injured? Our intuition is silent about the properties that we actually see versus those that we infer based on our perceptual experience. The question is therefore where the boundaries of perception are. Where do we stop perceiving and begin to interpret, judge, infer, believe?

Perception, as a scientific object, is quite vague. Yet, philosophers have given many characterizations of perception, of its function and nature. The aim of this dissertation is to shed light on the kind of scientific object perception is, through a philosophical work based on analytical philosophy of mind, philosophy of science and, more generally, on current philosophical developments concerning perception. How could we know that a specific phenomenon is a perceptual one? How does the scientist know that she is studying perception and not another activity of the mind? What are the boundaries of perception? These are the starting questions of this work.

Asking "what are the boundaries of perception?" instead of "what is perception?" is a change of perspective that reveals - or so I argue - some of the blind spots of classical

questioning. In this introduction, I first present in more details the boundary question, then I list the scientific and philosophical areas that require an answer to this specific question. In a third section, I show the difficulties the boundary question faces. I finish with a presentation of the dissertation.

1.1 The boundary question

”What are the boundaries of perception?” is the main question of this dissertation. This is actually more a starting point than a question to which I give a definitive answer at the end of this work. I do not elaborate a particular theory of perception, from which I could derive its boundaries. Instead, I question the grounds on which one may draw boundaries to perception.

The boundary question is looking for the extension of the concept PERCEPTION¹. What are the mental phenomena, studied by cognitive science, that fall under the concept PERCEPTION?

Philosophical questions about perception traditionally take the following forms: What is perception? What does constitute the faculty of perceiving? How does it work? Or, what is it like to see, to hear, etc.?. These questions point towards theories of perceptual mechanisms, and characterization of the nature of perception. Here, I take an alternative standpoint, that consists in questioning the criteria that makes a phenomenon perceptual. It is a classificatory enterprise. This work is part of the larger perspective to organize the mind into categories or parts. Is perception a legitimate category or part of the mind? If yes, what are the criteria to meet for a state, event, process or experience to be part of this category? What are the boundaries of this part/category of the mind?

The switch of perspective would not have been possible without the development of philosophical theories and without scientific progress that have brought out the pressing issue to clearly demarcate perception from other mental abilities.

The problem is not so much to find a satisfactory characterization of perception, but to determine the place of perception within the cognitive architecture. Characterizing perception is therefore essential, but with a more global goal. Here is another way to express the shift of questioning. There are two ways of asking ”what is perception?”. The first one points to an object and asks ”what is that?”. The second takes the mind as an object of scientific study, and asks what can legitimately be called ”perception” within this complex object being studied. Pointing is part of the first question, while it is part of the answer to the second question. In this work, I take the second perspective.

Let’s take an example with vision (which will be the privileged example through the whole dissertation). Visual studies inform us today on the organization of early vision, i.e. the mental mechanisms that receive retinal information. Visual mechanisms extract contrast, orientation, colors, and other low-level properties from the visual stream. It is widely accepted that this extraction is part of the whole story about perception. We also know, by psychological studies, that categories or event-types are mentally represented following some visual stimuli (Strickland and Scholl, 2015; Potter et al., 2014; Thorpe et al., 1996). A question arises: Are the processes that produce category or event-type representation part of perception? Or are they part of another mental ability? Are they perceptual or

¹I use small capital letters to denote concepts, by contrast with the objects to which the concepts refer.

post-perceptual?

I need to make a short but important comment at this point. Most of the literature in philosophy articulates the debate with a narrower scope, and targets the border between perception and *cognition*, understood as the faculty of thinking and reasoning. This border is often called the *perception/cognition divide*. The literature about the perception/cognition divide is one of my privileged source of reference in this work, and one may worry that the border at stake in this debate is different from the one I am looking for. First of all, it is not always clear that the authors writing about the perception/cognition divide are restraining cognition to the ability to believe or to think. Attention, memory, and other non perceptual abilities are usually taken as being part of cognition. Furthermore, I consider here that the perception/cognition divide is a part of the boundaries of perception, so that focusing on it partly answers the boundary question. The interface between perception and beliefs is probably the most problematic as perceptual states and beliefs interact: it is therefore a permeable boundary that has to demarcate them. In my dissertation, I consider the perception/cognition divide as a constituent of the boundaries of perception, but it does not exhaust it.

1.2 A much-needed answer

Why do we need to answer the boundary question? Couldn't we just wait for empirical findings to inform our best psychological theories, from which we will infer the scope of perception?

However, the boundary question is implicitly called in scientific practice, especially in the interpretation of scientific results. Furthermore, several philosophical debates, required for theorizing about perception, are particularly concerned with the boundary question: the cognitive penetrability of perception, the admissible content of perception and the classification of peculiar mental phenomena. Finally, the boundary question has also consequences in epistemology and even ethics. I shortly present them at the end of the section.

Experimental design and interpretation of results The identification of perception and its borders is crucial in experimental studies of perception. In order to run experiments on perceptual abilities, one has to make sure that what is measured is genuinely perception, and not something else.

Today, cognitive scientists are equipped with methods of differentiation, specific experimental designs and control techniques that allow them – to a certain extent – to correctly attribute the cause of observed effects². Their goal is to disentangle genuine perceptual effects from response-bias, decision criteria and other non-perceptual factors.

It is especially difficult to disentangle perceptual effects from judgment (post-perceptual beliefs) in experiment asking for perceptual reports (Firestone and Scholl, 2016). When subjects answer to an experimental task, do they answer based on the way they perceived the world, or do they follow their post-perceptual inferences? This distinction can be important in terms of results: some cognitive effects could be considered unsurprising if

they are part of our belief machine, but revolutionary if they are attributed to perception.

Overall, scientists are looking for a clear distinction between perceptual effects and non-perceptual ones in order to run a broad range of experiments in their everyday scientific practice.

The cognitive penetrability of perception First, the cognitive penetrability of perception is the thesis according to which non-perceptual states such as beliefs, desires or knowledge of an agent can influence how she perceives the world. Some empirical evidence point towards a direct and systematic influence. For example, a desired object could be seen closer (Balcetis and Dunning, 2010), and racial categorization could modify the perceived lightness of a face (Levin and Banaji, 2006). These pieces of evidence constitute today the main ground on which psychologists and philosophers argue in favor of the cognitive penetrability of perception.

This thesis is however controversial. In the late twentieth century, the favored hypothesis was impenetrability of perception. Fodor (1983) argues indeed that perception is a module, informationally encapsulated from the rest of the mind in such a way that cognition cannot influence perception. In the same tradition, Pylyshyn (1999) argues that the earliest stages of vision (in a hierarchical model of vision) are impenetrable by cognitive content.

Impenetrability is supported by the high speed of perception, as well as its reliability. The rapidity of perceptual processes is said to be possible because they are only taking into account a limited amount of information in order to generate outputs. Secondly, for perception to be reliable, it should provide information to the organism about "what's there, not what it wants or expects to be there" (Fodor, 1983, p.68). According to the proponents of impenetrability, the best way to ensure reliability is the informational isolation of perceptual processes. If the processes do not have access to cognitive expectations, desires, emotions, etc., the perceptual outputs will not be biased by them, and perceptual information will be reliable. These two properties of perception (high-speed and reliability) are expected by evolutionary considerations. Consider the example of distinguishing a tiger in a bush. The quicker and more reliably you see it, the sooner you run to save your life. Individuals with a quick and reliable perceptual ability would have a selective advantage.

However, empirical studies put pressure on the impenetrability thesis. Let's focus shortly on a specific example, in which the perception of colored fruits is influenced by high-level visual memory (Hansen et al., 2006)³. The perception of color is considered low-level, processed in the early stages of vision. Even if these processes are complex (visual processes disentangle the contribution of illumination from reflectance of objects in the incoming input), they are thought as being essentially local. In other words, all that is needed for seeing colors would be early visual processes. However, Hansen et al. (2006) show that the known color of objects, like the typical color of fruits, also participates in the perception of the color of these objects. Experimenters presented to participants photographs of natural fruits on a gray background (on a screen) and ask participants to manipulate the color of the fruit until it matches the color of the background. They found that participants

³For example, in psychophysics, one may cite experimental designs such as two-alternative forced-choice method and the multiple alternative perceptual search, or theoretical measures such as the point of subjective equality or the decision criterion of the signal detection theory (Finlayson et al., 2018).

adjusted the color of the fruit away from the neutral gray in the opposite direction of the typical color. They had to make the banana a little bluish for it to appear gray (if the banana is in neutral grey, it still appears yellowish). It is a case in which our knowledge of the world (the typical color of objects) influences our perception of it, which speaks in favor of the penetrability of perception.

These empirical studies in favor of the penetrability of perception have been much discussed, and some researchers cast doubt on the interpretation of empirical results in favour of cognitive penetrability (Firestone and Scholl, 2016; Machery, 2015; Mandelbaum, 2017). They highlight a systematic *locus problem* (Machery, 2015), i.e. the systematic failure of empirical evidence to secure the fact that the influence of higher-order states is *on perception* and not on post-perceptual judgments. As mentioned before, this problem is partly due to a lack of clear identification of the scope of perception.

This ambiguity has already been pointed out, and the penetrability debate has been refined into two specific thesis: (1) the cognitive penetrability of early vision and (2) the cognitive penetrability of perceptual experience (Macpherson, 2017). The latter holds that the way the world appears to us through conscious perceptual experience is influenced by cognition, but acknowledge that the earliest stages of visual processes are impenetrable. The former, however, states that perception is penetrable all the way down to the earliest stages of visual processes. Cognitive penetrability of early vision is therefore stronger than cognitive penetrability of perceptual experience.⁴

One may push the question further and ask what counts as genuine perceptual experience, or what counts as a perceptual mechanism? The identification of the boundaries of perception is at the heart of the debate, and is required to judge whether empirical findings are in favor of the penetrability thesis or not. If perception is shown to not have boundaries at all, i.e. if perception is not a distinct part of the mind, then it is not clear whether the penetrability debate can still take place at all. As Lupyan says: "I am committing to a collapse between perception and cognition. Following such a collapse, the very question of whether cognition affects perception becomes ill-posed." (Lupyan, 2015, p.38)

The admissible content of perception Another debate involving the issue of the boundaries of perception is the admissible content of perception. When I say "I see a red square", it seems legitimate to conclude that the content of my perceptual experience is "a red square". On the other hand, if I saw light on the bottom of the door of my roommate, I may say "I saw that Ahmed is at home", but it does not seem right to conclude that I literally saw Ahmed at home. The possible content of our perceptual experiences seems then to be limited to what our perceptual system can detect in the environment – visual properties for vision, sounds for audition, etc. However, the debate is still open on the kinds of properties that are perceivable in this sense. A first camp argues that only low-level properties are perceivable – for vision, properties such as colors, orientation, contrast, etc. detected by the primary visual cortex (Clark, 2000; Tye, 1995; Dretske, 1995). The

³For a quasi-exhaustive list of papers about 'top-down' effects, see the bibliography gathered by Firestone and the Yale Perception and Cognition Lab: <http://perception.yale.edu/Brian/refGuides/TopDown.html>

⁴A second lively debate is the nature of the cognitive influence on perception for it to be genuine penetration. This is an important part of the discussion, but I do not need it here to show the relation of the cognitive penetrability of perception with the boundary question.

other camp argues that higher-level properties are also perceived, such as being a dog, being drinkable, "being a pine tree" (Siegel, 2006), or "being an occlusion event"/"being a containment event" (Strickland and Scholl, 2015). Roughly, research in vision science tells us that the visual system is hierarchical in terms of property detection. In the first areas of the visual cortex, sets of neurons detect orientation, contrasts and colors. This information is integrated and higher properties are detected/recognized, such as faces. Once a face is recognized, the perceiver may react to it. As for me, if I see the face of my mother, I will smile. If I see the face of my old History teacher, I will flee. Where should we put the limit of perception? Where do we stop to perceive for us to think, decide, react?

Perceptual content can be defined either as the content presented in perceptual experiences, or as the content produced/manipulated by perceptual mechanisms. In any case, the boundary question is central to this debate. When one asks "Is 'being a chair' a possible content of perception?", one wonders whether the scope of perception includes representation of categories such as 'chair', like colors, orientations, and other low-level properties. Another way to reformulate the debate could be: given the set of perceptual states, what is the possible contents of these states? This formulation clarifies the importance of knowing which mental states are perceptual - independently of their content - in order to identify the possible content of perception.⁵

The classification of peculiar mental events Some processes, states or mechanisms in cognitive science have a perceptual dimension, but researchers wonder whether or not these states should be thought as full-fledged perceptual states. One may cite cases of mental imagery, emotions, introspection, hallucinations, synesthesia.

Synesthesia, for example, is a phenomenon where the stimulation of one sensory modality such as audition elicits an experience in a second unstimulated modality such as vision. Some synesthetes perceive colors and shapes while listening to music. It is unclear however whether this phenomenon has to be explained inside the perceptual domain only or has other cognitive dimensions. The classification depends on the mechanisms underlying these phenomena. As for synesthesia, it is still not clear which mechanism(s) give rise to these peculiar experiences. An explanation would be that these experiences result from a mechanism of hyper-association, in which there is a structure or functional connection between two perceptual brain areas typically unconnected (Ramachandran and Hubbard, 2001; Grossenbacher and Lovelace, 2001). Another hypothesis (semantic-based theories), makes use of conceptual information, that would play a role of mediation between the typical perceptual experience and the other one (Chiou and Rich, 2014).

The mechanism of hyper-association seems more perceptual than the semantic-based mechanism. But one could cast doubt on the fact that multi-sensory experiences are perceptual in the same way as unimodal experiences, leading to the idea that hyper-association is actually not perceptual. On the other hand, some conceptual information can be considered as being genuinely part of perception (Mandelbaum, 2017), which could lead to accept that semantic-based mechanisms of synesthesia are perceptual. The upshot is that it is not sufficient to find the mechanism underlying a phenomenon such as synesthesia

⁵This debate is orthogonal to the cognitive penetrability of perception, as perception of higher-level properties may or may not be included in an impenetrable perceptual system (see Mandelbaum, 2017 for higher level properties included in an impenetrable perceptual module).

in order to state whether or not this phenomenon is perceptual. Theoreticians are still confronted with the question of whether the discovered mechanism is indeed perceptual.

What is required is a characterization of perception which allows to identify mechanisms as perceptual or not. But what grounds this characterization? Several criteria seem legitimate, but deliver incompatible classifications. Let's take mental imagery as an example. Mental imagery partly uses the same neurological network as prototypical vision. However, an experience of mental imagery is said to lack a certain presence, such that one is not inclined to judge that what is experienced is really out there in the world. If perception is defined by its neurological implementation, it seems legitimate to classify mental imagery as perceptual. On the other hand, if the phenomenological presence in perceptual experience is an essential feature of perception, mental imagery should be classified as non-perceptual.

More generally, it is important to note that integrating a mechanism into perception has some epistemological consequences. Perception is considered as a privileged source of knowledge. This is shown in our tendency to assert propositions based on our perceptual experiences. When I said this morning: "Oh, it's sunny today", it is because I saw through the window that it is sunny. And nobody in the house asked me: "Really? How do you know that?". This epistemological value of perception is problematic though. We are all well acquainted with the illusions induced by our senses, like a straight stick bended in water. However, saying that a mechanism or an experience is perceptual will at least put them under epistemological scrutiny. These epistemological consequences of the boundary question are an additional pressure for the three debates evoked earlier in philosophy of perception. New questions arise indeed: If our beliefs or emotions influence perception, how can we trust perception to provide accurate information about the world? Depending on the possible content of perception, can I trust to the same degree that the tree that I see in front of me is (i) green, and (ii) a pine tree? If synesthesia is perceptual, should we expect it to justify knowledge of some sort?

The boundary question may be thought of as to be a technical question between philosophers and scientists, such that it has no broader consequences outside of these expert areas. However, the boundaries of perception may also have social and political consequences. Here is an example: we are intuitively responsible for what we believe, but not for what we perceive. There is a kind of moral responsibility associated with our beliefs (and from which perception is exempt). Therefore, if a behavior, thought to be guided by beliefs, is actually guided by perceptual experiences, it may be a reason to change the moral responsibility associated with it.

Here is an application on racial bias. Multiple studies have shown that Black men are *seen* as more threatening and dangerous than White ones. In Wilson et al. (2017), it is said that "people perceive young Black men as taller, heavier, more muscular, more physically formidable and more capable of physical harm than young White men of the same actual size" (p.60). But is it really a matter of perception or of post-perceptual interpretation? Do people perceive black men as dangerous, or do they believe that they are dangerous, based on other sources such as stereotypes and prejudices? Should we change the responsibility associated with racial biased behavior based on these research? And finally, what policy measures can change these behaviors, given that some of them are guided by automatic perceptual responses?

The resonance of perception in debates of moral responsibilities and maybe justice, as well as its role in epistemology mentioned earlier, makes the boundary question a much-needed investigation in a wide range of different areas.

1.3 The challenges of the boundary question

I showed that the boundary question comes from (and is thus justified by) the urge to determine the scope of our perceptual abilities. I should now insist on the fact that the boundary question raises fundamental questions that are still open.

I use these difficulties as challenges throughout my dissertation to guide my discussion. In this section, I present these three challenges:

- The distinctness challenge: the boundary question presupposes that perception is a *distinct* part of the mind. Is it true? On which ground?
- The methodology challenge: there is no consensus on the scientific methodology required to find the boundaries of perception. Perception is not an object like the image of an apple on a white background: we cannot merely draw an obvious line between the apple-figure and the non-apple-background. How should we proceed?
- The grounding challenge: there are actually numerous suggestions on how and where to draw the distinction between perception and the rest of the mind. On which grounds should we evaluate these proposals?

The distinctness of perception Metaphorically, the mind can be conceived as a machine, with several separable pieces or subsystems - each of them can be built separately and put together later. If we want to build a mind this way, surely we would build a subsystem for perception before assembling it with the other subsystems such as memory, the motor system, the decision-making center, etc. However, our starting point here is not an ensemble of spare parts to assemble, but on the contrary, a unified whole to cut into parts. It is, in a sense, a work of reverse-engineering.

Until now, I presupposed that perception is legitimately a distinct part of the mind, and that the boundary question focuses on where we should draw the borders of this part. But on which ground do we presuppose this? Actually, recent literature casts doubt on this distinctness of perception, arguing mostly that an organization of the mind based on clear-cut abilities such as perception does not reflect the way the mind works. Goldstone and al. say:

[The] tactic of isolating visual perceptual processes from attention and action has the negative consequence of neglecting interactions that are critically important for allowing people to perceive their world in efficient and useful ways. (Goldstone et al., 2016, p.33)

The boundary question does not merely ask where to draw the borders of perception, but also questions the ground on which we should assume that perception is a legitimate part of the mind.

I will present in this dissertation several way to "be distinct". For now, I take the term "distinctness" is a loose way. I am not necessarily looking for an actual part of the mind, or the brain - instead, I am looking for a set of mental abilities, mechanisms, states or processes that I can conceptually separate from the rest of the mind, and call "perception".

Methodology The question requires to be careful with methodology. How should we proceed to determine the boundaries of perception?

In the question in which one points to an object and asks for an explanation, clear boundaries of that object are not required. The object is defined *intuitively*. But is intuition a good guide for our investigation here? Intuition can arguably make us grasp paradigmatic cases of perception: I am currently seeing my computer, and this mental state is a perceptual one. Yet, paradigmatic cases are not sufficient for delimiting a kind. If you know paradigmatic cases of birds (like robin or pigeon), it will hardly give you the tools to judge whether pinguins are birds or not.

One may argue that it is entirely an empirical matter. It is by exploring psychological processes and neural systems that we will eventually discover the architecture of the mind and its different subsystems such as perception. Vision science already discovered an advanced description of the visual system. Should we be optimistic that one day this line of inquiry will deliver a final answer concerning the boundaries of perception? Obviously, empirical findings are necessary for building an answer to the boundary question. Yet, as I mentioned earlier, the boundaries of perception are required partly for empirically studying perception in the first place. Waiting for empirical studies to do *all* the work is not an option.

Another pitfall consists in considering the boundary question as a matter of terminology, a "verbal dispute". A verbal dispute is a debate in which people talk past each other because they use the same term to talk about different things. There is a sense in which the boundary question can be seen as a question about the meaning of the word "perception". Let's take the case of illusion to illustrate it. Whether illusions are special cases of perception (non-veridical perception), or non-perceptual cases, is an old philosophical dispute. Should we draw the borders of perception such that illusions are included or excluded from the ensemble of perceptual phenomena? Evidence to resolve this question could be found in empirical studies (i.e. do illusions and other perceptual phenomena share common neural networks?) or in philosophical investigation (i.e. do illusions and other perceptual phenomena share essential properties?), but there is also a sense in which it seems to be a verbal dispute. Whether illusions are part of perception may just be a question of what you mean by the term "perception". Let's consider that the meaning of a term is partly constituted by its extension, namely the list of things to which the term can be applied. Then, whether illusory phenomena are perceptual can be seen as the question whether the extension of "perception" does or not include illusions. This is why it can be seen as a question about the meaning of "perception", and therefore (at least partly) a "verbal dispute".

The problem of verbal disputes is their "familiar and distinctive sort of pointlessness" (Chalmers, 2011b, p.10). After all, philosophers and scientists defending positions about the boundaries of perception are interested in the first place in how the mind works, and

how it is organized, and far less interested in the meaning of words. The verbal part of a disagreement can therefore be seen as an obstacle in the path of important investigations and should be quickly resolved to allow the whole energy of research to focus on substantive questions.

Interestingly, there is a sense in which identifying verbal disputes is very important precisely because these are not the disagreements on which we want to focus. As Chalmers says:

The diagnosis of verbal disputes has the potential to serve as a sort of universal acid in philosophical discussion, either dissolving disagreements or boiling them down to the fundamental disagreements on which they turn (Chalmers, 2011b, p.3)

An investigation about the boundaries of perception has to find a safe path between the temptation to rely on intuition, the necessary but insufficient knowledge coming from empirical studies, and the pointlessness of verbal disputes. The journey will not be easy.

The grounding challenge The last challenge of the boundary question is to find a ground based on which proposals of boundaries can be evaluated. There are actually a lot of different suggestions of criteria in the literature to delimit perception: its distinctive phenomenology, the format or the content of perceptual states, the modularity of its processes, etc. (I present them in the first chapter). Through all of these suggestions, authors argue to capture what makes perception distinctive with respect to the rest of the mind. The main challenge is less to find a way to delimit perception than to evaluate the different suggestions. But how can we evaluate them? On which basis? What makes an alleged boundary better than another one? What we need is to ground the boundaries of perception into a more fundamental fact. It means that the boundaries of perception should be what they are in virtue of something else.

The basis on which the distinctness of perception is drawn should be fundamental, as the distinction is supposed to organize the mind as a scientific object. Imagine an organization that divides the mind into two sub-parts: (i) all the mental events about dogs and (ii) the rest. This is a good metaphoric way to kindly make fun of dog-lovers, but it seems inappropriate when it comes to providing a scientific description of the mind. The boundaries of perception should reflect how the mind is actually organized, or at least, should be a source of scientific progress in one way or another.

The goal of my work is neither to praise the value of my favored boundaries, nor to elaborate a new proposal that is supposedly be better than the others. Instead, my aim is to scrutinize the whole problem at its source and question its ground and legitimacy.

1.4 Plan

I will not tackle each of the challenges separately in this work. Instead I will navigate through several strategies and orientations that will shed light on different aspects of the boundary question. It is nonetheless important to keep in mind the three challenges (distinctness, methodology and grounding), as they will regularly constitute motivations to

move forward in this work.

The methodology challenge is one of the main reasons to move from the natural-kind strategy (part 1) to conceptual engineering (part 2). The grounding challenge constitutes the motivation to raise metaphysical questions that are left aside by conceptual engineering (part 3). Finally the distinctness challenge comes back at the end, in the suggestion to give up on perception and its boundaries in favor of a different organization of the mind (part 4).

My work is organized in four parts. In what follows, I highlight important concepts in *italic*.

Part 1 I open this work with a presentation of the usual method for determining the boundaries of perception (**chap.2**). I call it the "natural-kind strategy". This strategy consists in assuming that perception is a *natural kind* of mental states or processes, and then finding the property(ies) of the states/processes that underlie this natural kind.

Roughly (and arguably), natural kinds are classes of things sharing a common nature. They are supposed to reflect a kind of *objective similarity*. Chemical elements and biological species would be natural kinds. By contrast, the ensemble of plants in my garden is not. Natural kinds are taken to play a crucial role in science, in underlying induction. What is observed for a few members of the class can be generalized to the whole group. In this way, botanical observations about tulips can be generalized (in some conditions) to all tulips. However, botanical observations about the tulips of my garden cannot be generalized to the other plants in this same garden.

Once the properties thanks to which all perceptual phenomena share the same nature are found, one can use them to determine the boundaries of the class of perceptual phenomena. All the things that possess these properties are perceptual, all the things that do not possess them are not.

Attempts to fulfill the natural-kind strategy for perception are numerous in the literature (based on the format or the content of mental states, or the modularity of processes); yet, this strategy does not lead to any consensus. I propose to study the disagreements among the debaters and identify four different sources of disagreements. Two are empirical and harmless for the natural-kind strategy. However, the third one, a theoretical disagreement about natural kinds, jeopardizes it, so that it is better to explore other strategies to draw the boundaries of perception (**chapter 3**). Fortunately, the fourth disagreement is about the meaning of the concept PERCEPTION and can be used as a springboard. I identify this disagreement as a *metalinguistic negotiation*, i.e. a dispute in which people do not disagree about the world, but about the concept they should use to describe the world (**chapter 4**). But, if PERCEPTION does not refer to a natural kind, how can its extension be determined? How can we decide who is right if we cannot base the extension of PERCEPTION on the hypothetical boundaries of a natural kind? I conclude that this kind of disagreement should be tackled by the method of *conceptual engineering*.

Part 2 The second part of the dissertation is devoted to conceptual engineering about PERCEPTION. The goal is to arbitrate the metalinguistic disagreement identified in the first part. I identify several different meanings of the concept PERCEPTION in the literature

(chap.5).

I study PERCEPTION as a scientific concept, whose meaning can be broken down in three components: the classical intension and extension (respectively the definition and the list of things that fall under the definition), and the *epistemic goal* of the concept (Brigandt, 2010). The epistemic goal of a concept is the goal the scientists pursue by using the concept. According to my analysis, there are at least four different epistemic goals of PERCEPTION, leading to four different concepts of perception, each of which displaying some conceptual variation.

I suggest a method for evaluating these concepts, and argue in favour of conceptual pluralism (**chap.6**), i.e. that there are several concepts of perception, each as legitimate as the others. The method of evaluation is based on the notion of *concept utility* (Egré and O'Madagain, 2019), which is a measure of the value of a concept for scientific inquiry in terms of its contribution in formulating plausible and informative hypotheses.

In addition to conceptual pluralism, I identify other factors that tend to show that the conceptual pluralism of PERCEPTION is a particular stage in a classical process commonly undergone by scientific concepts - *conceptual fragmentation*. This makes the conceptual pluralism of PERCEPTION easier to accept, and even expected. A scientific concept is fragmented into several subconcepts when it is discovered to refer ambiguously to several objects, or to be used for different theoretical uses (Taylor and Vickers, 2017). Several scientific concepts such as SPECIES, ACID, ATTENTION, CONSCIOUSNESS are subject to conceptual fragmentation. PERCEPTION is one of them.

Part 3 Conceptual pluralism gives an answer to the boundary question, namely there are several boundaries to perception, determined by the extensions of the legitimate concepts of perception. However, this answer does not seem to be satisfying from a metaphysical point of view. Intuitively, the extension of a concept should not *ground* the boundaries of the object to which it refers. It seems that the grounding relation should be reversed. This is my motivation to question metaphysical grounds in a third part of my work.

Here, I explore metaphysical foundations for the boundaries of perception, apart from naturalkindness (**chap.7**). I show that there are several ways to ground realism towards perception and its distinctness, without the need to bow to the greedy demands of natural kinds. I propose to use the notion of *joint-in-nature* to analyze these different possibilities.

The notion of joint-in-nature, coming from Plato's metaphor of carving nature at its joints, refers in my work to a distinction between two things (entities, groups of entities, parts of entities, properties) that is grounded in the architecture of reality. I disentangle at least three kinds of joint-in-nature:

1. Taxonomical joint-in-nature: a distinction between two kinds of things.
2. Mereological joint-in-nature: a frontier between two parts of a whole.
3. Focal joint-in-nature: the contours of a portion of a system, described through a partial description (a *flattened* description using only a set of variables).

Applied to perception, its boundaries may correspond to (i) the extension of a kind of mental states, (ii) the border of a (proper) part of the mind or (iii) the contours of a pattern revealed by a partial description of the mind. Depending on the metaphysical object perception is taken to be, empirical results about it will be interpreted differently (**chap.8**). To show this important consequence, I use the debate about the cognitive penetrability of perception as a case study.

Part 4 After these first three parts, PERCEPTION has been shown to contain several more specific concepts and to be ambiguous from a metaphysical point of view (with important consequences). These reasons motivate some authors to advocate for the elimination of the notion of perception within cognitive science, and for its replacement by the alternative notions generated through the process of conceptual fragmentation (e.g. early-vision, top-down/bottom-up processes, etc.). This is a way to answer the boundary question: there isn't actually any boundary to perception, because it does not refer to any genuine scientific object. In this last part, I defend the epistemic value of PERCEPTION in cognitive science, even if it is a complex, ambiguous and even confusing concept.

I first argue that the general concept PERCEPTION still has a role to play in science, as an organizational concept, which provides a genuine epistemic contribution to cognitive science (**chap.9**). I then confront PERCEPTION to arguments in favor of eliminativism and show that its organizational role protects it against elimination.

I end up by giving some recommendations for future philosophical and scientific research about perception.

The main conclusions of this work are the following:

- **The failure of the natural-kind strategy:** For methodological reasons, it is not possible to find the boundaries of perception based on its alleged naturalkindness.
- **Conceptual pluralism:** There are several legitimate concepts of perception, with different extensions.
- **Conceptual fragmentation:** As a scientific concept, PERCEPTION is currently going through a process of conceptual fragmentation - i.e. a concept thought as having a single meaning, appears to be scientifically used with several distinct meanings. This process is a common scientific one, situating PERCEPTION in a classical history of scientific concepts.
- **Metaphysical ambiguity:** Beside naturalkindness, there are several alternative ways to metaphysically ground the boundaries of perception. The choice between the alternatives is crucial for the interpretation of empirical results on perception.
- **Resistance to eliminativism:** Despite a pragmatic urge to eliminate the notion of perception from cognitive science, we should keep it, but recognize its particular role in today's scientific research. It is an organizational concept, whose main function is to guide, coordinate and integrate interdisciplinary research about perception.

At the end of this journey, I have no perceptual boundaries to sell (or too many) but this may be for the best. Giving the notion of perception its rightful place in scientific and philosophical investigations about the mind would contribute, or so I hope, to avoid ill-posed questions and fruitless debates. By stoping to look for boundaries of perception,

we open the frontiers between disciplines, research traditions and models of the mind.

Part I

Perception as a natural kind: disagreements and springboard

Chapter 2

The natural-kind strategy

The purpose of this chapter is to outline a fairly common strategy for finding the boundaries of perception. This is the *natural-kind strategy*. This strategy equates perception with the class of perceptual states. This class is then considered as an *homogeneous* class of mental states. Of course, an auditory state is different from a visual states, so are two different visual experiences (like the visual experience of a sunset vs the experience of a crowded street). But the idea is that all of these states are fundamentally similar in terms of the *kind* of states they are in the mind. They share fundamental properties in virtue of which they behave in a similar way in the mind, have similar functions and thanks to which researchers may hope to find an integrated theory of perception. In other words, perception (as the class of perceptual states) is assumed to be a *natural kind*.

Because perception is a natural kind, there is at least a property shared by all (and only) perceptual states. Once this property is known, it becomes the tool needed to establish the boundaries of perception. This is roughly the proceedings of the natural-kind strategy.

This chapter is presented as follows: first, I introduce the central hypothesis of this part, namely that perception is a *natural kind* (section 2.1), then I develop the natural-kind strategy (section 2.2). In a third section, I illustrate this strategy by the properties presented in the literature as the *hallmarks* of perception, i.e. those distinctive properties of perception, which make it possible to determine the perceptual nature of a state when it possesses it (section 2.3). Finally, I introduce the main difficulties of the natural-kind strategy: the proposed hallmark properties of perception are incompatible, and none of them stands out from the crowd (section 2.4).

In this chapter, I conclude that the natural-kind strategy requires a tool to choose between the alleged hallmarks of perception.

2.1 Perceptual as a natural kind

Natural kinds are classes of things that are supposed to reflect the real structure of Nature, independent from human interests and descriptions. By saying that perception is a natural kind, one acknowledges that all the mental states can naturally be classified into kinds and perception is one of them¹.

¹The others could be beliefs, imagination, memory, etc., or completely different kinds, but this is not the goal of this work to find them.

The assumption that perception is a natural kind is widespread, but is rarely explicitly expressed. An exception is the following:

Each legitimate way of marking a border between perception and cognition invokes a notion I call 'stimulus-control'. Thus, rather than being a grab bag of unrelated kinds, the various categories of the perceptual are unified into a superordinate natural kind. (Phillips, 2019)

Phillips points out the nightmare for theoreticians of perception, i.e. that the class of perceptual states turns out to be a *grab bag of unrelated kinds*. In the perception/cognition divide debate, authors try to distinguish perception from cognition mostly in terms of intrinsic properties like the content and the format of the state (Crane, 1992; Evans, 1982; Heck, 2000; Dretske, 1969; Block, 2014; Burge, 2010). Alternatively, some researchers try to find the distinctive nature of perception in the architectural features of the processes or mechanisms underlying these states (Fodor, 1983; Mandelbaum, 2017; Firestone and Scholl, 2016). More rarely, authors suggest to demarcate perception from beliefs with their distinctive function (Beck, 2018; Phillips, 2019). All of these properties will be developed later during this chapter. What is important here is that nobody provides a demarcation in terms of conventional or ad hoc definition. People are trying to find a distinction that is already there, in the external world. If perceiving a cat is different from thinking at a cat, it is because perceiving implies a different kind of mental states than believing. This is a crucial idea for theoreticians of perception, that justifies the fact that they are using the category of perceptual states as a natural kind.

The notion of natural kind is however a technical one, with a complex history. I begin with providing a simple version of the conditions a class should meet in order to be a natural kind. These conditions will be important for the natural-kind strategy.

As quickly presented before, natural kinds are groups of things that reflect how Nature is organized, in a mind-independent way. Examples of natural kinds are chemical elements and biological species. Theories have flourished about natural kinds and especially about the necessary and sufficient conditions for a class of objects to be a natural kind. I present here three conditions that capture a classical version of natural kind. This notion is 'double-edged', because it has largely been commented, criticized and modified in the literature. I present here this version of natural kind, without much criticism. Discussion will come later, and will be a rich source of energy to pursue my inquiry further.

The three conditions a group of object should satisfy to be considered a natural kind are:

1. The membership condition: all the members of a natural kind resemble each other.
2. The mind-independence condition: the source of their similarity is mind-independent (i.e. natural and objective).
3. The fundamentality condition: they resemble each other in virtue of a common nature.

The membership condition Paradigmatic natural kinds are classes of chemical elements, or fundamental particles in Physics. In contrast, let's consider the group constituted by living cats, the number 3, the clock in my office, all bottle caps recycled in 2017 and the floor. These things have nothing in common, except the fact that I wrote them

next to each other in this page, in an arbitrary way. This is why this group cannot be a natural kind. In order to constitute a natural kind, all members of a group have to share something, most of the time a property, or a cluster of properties, in virtue of which their grouping is non-arbitrary. For example, all atoms of hydrogen share the same number of protons, in virtue of what they are classified into the class of hydrogen. It makes 'hydrogen' a good candidate for being a natural kind. Being a natural kind is therefore being a class of things resembling each other. This is what I call the *membership condition*. The membership condition is the minimal characterization of a natural kind.

If perception is a natural kind, it means that all perceptual states resemble each other in some sense. The easiest way is to share at least one property. This property should be exclusive and exhaustive of all perceptual states, i.e. all and only perceptual states should possess it. This (or these) property(ies) will be the one(s) targeted by the natural-kind strategy.

The mind-independence condition All pieces of luxury jewelry is an example of a *non-natural* kind because their source of resemblance is mind-dependent. The members of this group probably share some similarities, and consequently, they satisfy the membership condition. However, these similarities are due to human activities and thoughts like historical source, monetary or social value, etc. These similarities are therefore mind-dependent. In contrast, members of natural kinds are grouped together in virtue of *mind-independent* commonalities, and that is why they represent the architecture of reality, and not interests of humans. For example, the number of protons is (arguably) a mind-independent property of classes of chemical elements. It does not depend on human activity and thoughts.

It may seem weird here to talk about *mind-independence* of perceptual states (as they are mental states). This is true that if there were no minds, there would be no psychological phenomena and therefore no kinds of mental states. But there is a sense in which mental states can be considered as mind-independent, because their existence (arguably) is not constituted or determined by any beliefs, theories or conventions about them.

The fundamentality condition Finally, for a kind to be a natural kind, the similarity between its members have to be tied to their nature, the resemblance has to be fundamental. For example, red things have something in common, namely being red, based on which we could consider them as a kind. However, this property is not tied enough to the fundamental nature of things – in the group of red things, we would find red chairs and red birds, red flowers and red carpet, and they do not resemble each other enough in nature to be grouped together. This vague notion of fundamental resemblance is the feature of natural kinds that is particularly difficult to grasp. Quine captures it this way:

Things are similar to the extent that they are interchangeable parts of the cosmic machine (Quine, 1969, p.20).

Imagine that you have the super-power of moving things through the universe. If you switch two protons in the universe, it will arguably not change anything. If you switch a plant with another of the same species, it may not change a lot. But if you switch a plant with another one of another species, even if they are both blue - you may have changed

the balance of a whole eco-system (of course, this metaphor has its limits).

In order to understand this idea of fundamentality, it is interesting to look at the purported role of natural kinds in induction. The problem of induction is how we can extrapolate conclusions from a limited number of known cases to unknown cases. How do I know that all instances of salt would dissolve in water when I can only test it on a sample of cases? How do I know that I cannot do the same with the following case: I check the colors of a sample of flowers in a garden and they are all red, why cannot I infer with confidence that all the flowers in that garden are red? The difference is that the dispositional property ‘dissolve in water’ is entrenched in the chemical compound ‘salt’, whereas ‘being red’ is not a fundamental property of flowers. These fundamental properties, which all members must share to constitute a natural kind, are also those on which generalizations can be made (Goodman, 1955; Quine, 1969).²

Other conditions have been discussed in the literature for a class of objects to be a natural kind, e.g. whether the properties shared by all members of a natural kind have to be intrinsic or not (Ellis, 2001). For the sake of clarity, I let these conditions aside and only take into account membership, mind-independence and fundamentality as conditions. Let’s sum up now. A natural kind is a non-arbitrary grouping of things that resemble each other in a fundamental and mind-independent way, such that some true facts about one member of this group are also non-accidentally true for the others.

The assumption that perception is a natural kind means therefore that the class of perceptual states is constituted by mental states that resemble each other in a fundamental and mind-independent way, such that some true facts about some perceptual states are also non-accidentally true for the whole class.

Because of the tumultuous history of the notion of natural kind in philosophy, this characterization could already be largely attacked. First, this is a realist characterization of natural kinds. Some philosophers build natural kinds as partly depending on some conceptual schemes or interests of the categorizer (Kitcher, 2007). Furthermore, the membership condition follows the essentialist tradition, in which all (and only) members of a kind share the same property(ies), considered as their essence. Non-essentialist alternatives have been offered, like Boyd’s homeostatic property cluster theory. According to Boyd, the members of a natural kind share a clustering of properties, but no property is necessary for membership (Boyd, 1999). These alternatives will be discussed later in this work (chapter 3, section 3.2).

For now, this characterization will help us meet the three challenges I mentioned in the introduction: the distinctness challenge, the methodology challenge, and the grounding one (see introduction, section 1.3).

2.2 The strategy

Based on the assumption that perception is a natural kind, the natural-kind strategy allows in a three-step method to find the boundaries of perception. These boundaries are drawn in virtue of the naturalkindness of perception, so that they are grounded into a more

²More specifically, natural kinds are supposed to underlie law-like generalizations, i.e. universal and necessary generalizations.

fundamental fact. This strategy meets therefore the *grounding challenge* of the boundary question.

Furthermore, the fact that perception is a natural kind secures that the class of perceptual states is a distinct ensemble of mental states. Natural kinds have indeed *sharp* boundaries. The common nature shared by perceptual states as members of a natural kind cannot be partially shared by other mental states. In other words, there is no intermediary cases between two natural kinds. If it is accepted that the mental realm is organized into natural kinds, then it is expected that there is no gap and no overlap between them. Natural kinds are mutually exclusive and globally exhaustive. Consequently, perception as a natural kind fulfills the *distinctness* challenge of the boundary question, i.e. the distinctness of perception is grounded in its naturalkindness.

Finally, the natural-kind strategy works partly thanks to the contribution of natural kinds to valid generalizations. The assumption that perception is a natural kind provides therefore a methodology for finding the boundaries of perception. I elaborate on this point during the presentation of the method.

The strategy roughly consists in looking for the property (or properties) related to the common nature of all perceptual states. Once we have found it (and verified that it does indeed generate a natural kind that satisfies all the conditions), we can use it as a criterion for determining the boundaries of perception. Here are the three steps of the natural kind strategy:

1. **Starting assumption:** the class of perceptual states is a natural kind
2. **Discovery of commonalities:** this ensemble has a common nature (because it is a natural kind), which is manifested by the sharing of one (or more) property that all and only members of the set of perceptual states possess. Let's discover this property. It is usually called the "hallmark" of perception.
3. **Determination of boundaries:** The hallmark can be used as a criterion of individuation of the natural kind, as all and only members of the ensemble of perceptual state possess it. The boundaries of perception can therefore be established between the mental states that possess the hallmark of perception and those that do not possess it.

One may worry that the weight of the starting assumption is too high given its contribution to the overall conclusion. This can be nuanced. The assumption opens the research of hallmarks and will cease to be an assumption when the hallmark will be found. By finding a common property that generates a class of things satisfying the membership, the mind-independence and the fundamentality conditions, one secures the fact that this class is a natural kind. Therefore, at the end of the second step, we do not need the starting assumption anymore. But then why do we need it in the first place? Can't we merely begin with a research of the common properties across perceptual states? The problem is that we do not know the extension of the ensemble of perceptual states, as it is what we are looking for. How can we affirm that a property is common to all perceptual states if we do not know what counts as a perceptual state?

Imagine that I want to find a property common to all swans. If I base my answer on all the swans I met in my life, I can end up easily thinking that all the swans are white, so that

whiteness is a property shared by all swans. However, there are (very few) black swans (and I never saw one in my life). In order to know that whiteness is not a property shared by all swans, I need to know that the class of swans includes these particular cases of swans, namely black ones. My judgment on the common property of the class is therefore influenced by the extension I take the class to have.

Therefore, without the starting assumption of naturalkindness (step 1), the two following steps are circular: we need to know the scope of perception to find the common properties of perceptual states (we need step 3 to answer step 2), and we need the common property to find the boundaries and the scope of perception (we need step 2 to answer step 3). The assumption of naturalkindness breaks the circularity thanks to the possibility to generalize fundamental properties found in a sample of cases to the whole natural kind. In assuming that perception is a natural kind, the research of hallmark can be made on a sample of perceptual states (paradigmatic ones), and the result can be generalized to the whole natural kind, without knowing the extension of the natural kind.

It is useful here to introduce the notions of paradigmatic cases and controversial cases. *Paradigmatic cases* of perception are cases on which everybody agrees that they are indeed cases of perceptual states. There is no doubt that they are part of the class of perceptual states (if this class exists). An example of paradigmatic case of perception would be a visual state whose content is "a red square".

On the contrary, *controversial cases* of perception are cases in the vicinity of the perceptual border whose perceptual nature is being debated. They may be states whose content is controversially perceptual, like seeing a pine tree, or seeing a causal event. It may also be complex cases like illusions created by some environmental elements and mental imagery (e.g. a trompe l'oeil).

It is not recommended to draw conclusions on the class of perceptual states based on controversial cases. But the fact that perception is taken to be a natural kind allows to generalize the fundamental properties found in a sample of paradigmatic cases to the whole class of perceptual states. This would be similar to observe the property 'dissolve in water' for a limited number of salt cases and generalize it to all instances of salt. Then, I would be able to apply these properties as criteria of membership for controversial cases. Similarly, if I want to know whether a white powder is salt, I can perform a test of water solubility. Of course, other white powders dissolve in water, but this would give me a first idea (e.g. calcium sulfate, calcium carbonate, cornstarch, and potassium bitartrate are four insoluble white powders).

The second and third steps of the natural kind strategy can therefore be refined with the new notions of paradigmatic and controversial cases:

1. **Starting assumption:** the ensemble of perceptual states is a natural kind
2. **Discovery of commonalities:** find the hallmark of perception (the common and fundamental property of perceptual states) on a sample of paradigmatic cases of perceptual states, then generalize to the whole natural kind.
3. **Determination of boundaries:** use the hallmark as a criterion to establish whether controversial cases belong to the extension of perceptual states or not. This will establish the border of perception.

This whole strategy is rarely explicitly advocated in the literature (the exception is Block, ms.). However, the generalization of the possession of a property from a sample of paradigmatic cases to the whole class of perceptual states is common in the quest for the hallmark of perception.

For example, Crane (1992) argues that perceptual states have non-conceptual content (I elaborate on this suggestion in the next section). He does not explicitly defend this property as being the hallmark of perception, but he makes it the distinctive property between perception and cognition (understood as beliefs). He says:

The (causal) relation between perception and belief takes place at the level of the whole contents. But on the perception side of this transaction, the contents are not composed of concepts: concepts come later when thinkers employ the beliefs they thus formed, and the desires they have, in reasoning-belief conceptualises the content of perception. (Crane, 1992, p.20)

According to him, the border consists of a transformation that takes non-conceptual content and conceptualises it. It is therefore reasonable to understand non-conceptuality of content as the criterion of what it takes to be a perceptual state (at least in the ensemble of perception and belief states). He argues that perception has non-conceptual content based on the fact that perceptual states are not revisable like beliefs and he justifies it by taking the example of the Müller-Lyer illusion in which "P cannot help but see L1 as longer as L2" (Crane, 1992, p.15). He takes therefore a sample of perceptual states (some specific visual illusions), and generalizes from it to the whole category.

More explicitly, Block defends a similar method to find the border of perception. Block's methodology consists in finding scientific indicators "that converge on the cases we are most sure of, classifying some cases as perceptual, others as cognitive and none as both perceptual and cognitive" (Block, ms.)³. In other words, Block looks for properties shared by paradigmatic cases of perception, that individuate a group of cases sharing a common nature in an exclusive way.

Even if he does not use the term of "natural kind"⁴, he nonetheless acknowledges a common nature across all instances of perceptual states. Here, the formulation of the strategy makes explicit this crucial starting assumption of naturalkindness.

Now that I justified the assumption, the second step consists in finding the hallmark of perception (based on paradigmatic cases). These are the properties constitutive of perceptual states, the properties in virtue of which a perceptual state is a member of the natural kind perception. The object of the next section is to explore the different suggestions offered in the literature.

Caveat: It does not matter whether the hallmarks presented in the next section has been explicitly proposed in the perspective of finding the borders of perception. The boundary question is only one question among many concerning perception and its investigation. However, they may all serve as criteria of identification of the scope of perception if they

³He specifies that these indicators may not be the constitutive features of perception, but should point towards them. These indicators are symptoms of the possession of more fundamental properties. (Block, ms.)

⁴Block actually rejects the term of *natural kind* in favor of the more liberal notion of *joint-in-nature*. I explore the difference in the third part of this dissertation. For this part, the distinction is unnecessary.

are indeed properties common to all and only perceptual states.

2.3 The alleged hallmarks of perception

I present here several candidates of hallmarks of perception. I am well aware that the list is not exhaustive of the literature, but it shows at least the main propositions in contemporary philosophy of perception. I first begin to present some assumptions I chose to reduce my field of investigation.

2.3.1 Preliminaries

Our intuition tells us that perceiving an object is different from imagining it, or thinking of it. Our intuition is mostly based on phenomenological differences between perceptual experiences and other mental events: perceiving a piano, either by seeing it or hearing it, 'feels' different from thinking or imagining a piano, or performing mental arithmetic. Capturing the distinctive features of this phenomenology with words is complicated.

One suggestion would be to say that perceptual experiences are transparent. The transparency of perceptual experiences refers to the idea that through introspection, the subject of a perceptual experience can only reach the mind-independent objects she is perceiving, and not the properties of the experience itself. If blueness is one of the properties perceived through my experience, I experience it as the property of something in the world. I do not experience my own experience *as blue* (Moore, 1903; Tye, 2000).

Another characteristic of perceptual phenomenology is its perceptual presence (Dokic and Martin, 2017; Dorsch and Macpherson, 2018): the objects perceived are presented as existing, as being present in the environment of the perceiver.

The specific phenomenology of perception could be a candidate for being the hallmark of perception, however I do not follow this path here. The main reason is that phenomenology is sometimes misleading. For example, I was convinced a few months ago that I have heard some mouses in the attic of my mother's house, but after checking, the attic was clear of any rodent. I then apologized, I really thought I heard some, but I must have imagined them instead. The phenomenology of imagination may be close enough to the phenomenology of perception, at least in some occasions, for me to be mistaken between them. This story puts some doubt on the specificity of perceptual phenomenology, that may not be distinctive enough to ground the distinction between perception and the rest of the mind. Kriegel (2019) tries to find a "first-personally manifest difference between perception and thought" (Kriegel, 2019, p.22), and concludes that none of the options he explores in his work are satisfactory. I take it as an encouragement to look elsewhere.

Furthermore, phenomenology is a phenomenon of consciousness. Using phenomenology as a way to distinguish perception forces me to remain in the sphere of consciousness. For a long time, the realm of the mental was considered as co-extensional with the realm of consciousness (Brentano, 1874), therefore perception was also co-extensional with the class of conscious perceptual experiences. However, works on unconscious cognition are now legion, making the study of consciousness and conscious states a sub-object of the philos-

ophy of mind in general. When it comes to perception, paradigmatic cases are conscious - I am conscious of what I see on the TV, I am conscious of the music I listen on my phone, etc. However, psychological studies, neuro-psychology and psycho-pathological cases have paved the way to a possible unconscious perception that is today hard to dismiss from the study of perception (Prinz, 2010; Merikle et al., 2001; Dretske, 2006; Block, 2016; Burge, 2010).⁵

Once phenomenology has been discarded as potential hallmark of perception, three other routes are available. The first one - the representational approach - is focusing on the particularity of perception in its way to *represent the world*, either in terms of *content* or *format*. I present the representational hallmarks in the next section. The second route consists in finding the distinctive feature of perception in the way its underlying processes are organized in the architecture of the mind. These are the architectural hallmarks of perception. Finally, a third approach distinguishes perception by its distinctive function.

2.3.2 The representational hallmarks

In this section, I focus on views in which a perceptual state represents in some sense the actual state of the world. Most of the time, "world" is taken as the external environment of our bodies, but a perceptual state can arguably also represent internal information, such as the location of body parts through proprioception. The main point is that perceptual states are about the world, in a way that its content – or what it represents – can be correct or incorrect with respect to the world. In other words, contents of perceptual states are assessable for accuracy. For example, if I open my eyes and see a camel in front of me, my visual experience represents in some sense the external world. What is represented in my visual experience - the camel – may correspond to something true in the world (there is a camel in my visual field) or not (there is no camel in my visual field). In the former case, my experience is said to be correct, whereas in the latter, it is incorrect. The correctness conditions of my perceptual state, i.e. facts in the world that makes my experience correct, define the *content* of a perceptual state.

Other mental states have content. The paradigmatic mental state with content is belief. The content of belief is also assessable for accuracy. When I believe that there is a camel in front of me, my belief has also correctness conditions. What makes this belief true, namely that there is a camel in front of me, is the content of my belief.

The content of a mental state has to be differentiated from its *vehicle*. In representationalism, perceptual states are representations in the same sense than maps and pictures are. These representations involve two entities : (1) what is represented and (2) what represents. In the case of a map, what is represented is a spatial location with usually some properties such as distance and topography, and the colors and lines on the paper or electronic map are what represents it. For a picture, the piece of paper on which the picture is printed is the vehicle. The objects depicted on the paper constitute the content. Two identical reproductions of the same picture will have the same content, but not the same vehicle. Likewise, a perceptual state has a content and a vehicle. In perception,

⁵Some authors maintain that phenomenology can still be a legitimate marker for perception, despite discoveries on unconscious perception (Phillips, 2018). I will show indeed in the fifth chapter that consciousness may still have a role to play for a specific kind of philosophical and scientific questions for which perception is recruited.

what is represented is the world, what represents is the perceptual state as vehicle. The vehicle is the medium on which the representation stands, while the content is what is represented in the representation.

The distinction between content and vehicle is different from the classical distinction in philosophy of mind between the content and the attitude of a mental state. "Believing", "perceiving", "doubting", "fearing" are different attitudes that someone can hold toward the same content such as "it will rain today". For short, the expression of a mental state is often written with a "that"-sentence: I fear that the cat is dead, I see that my dad is home, I think that Villejuif is the capital of France. What is before the "that" is the attitude of the mental state, what is after is the content. Arguably, all contents can be targeted by all attitudes. I will not use the distinction content/attitude in my dissertation as it seems to me that the idea that perception is constituted by the attitude "perceiving that x" does not help to grasp the specificity of perception.

The fact that a perceptual state represents 'how the world is' is not specific enough for it to be the criterion of individuation of the class of perceptual states. Beliefs indeed also represent how the world is. This common trait between perceptual states and beliefs is captured by the notion of direction of fit. Direction of fit is roughly the direction of correspondence, or influence, between the mind and the world. Perception and beliefs have a mind-to-world direction of fit, meaning that mental states are supposed to fit the world, while desires and intentions have a world-to-mind direction of fit. In these states, the world is called to change to fit the mental state. The common direction of fit between perception and beliefs may explain partly why representational hallmarks of perception (specificity of content and format) are usually discussed in order to mark the distinction between perception and cognition (understood as the class of belief-like states). In the representational strategy, if one finds how to distinguish perception from belief, one finds how to distinguish perception from the rest of the mind.

There are two kinds of representational hallmarks of perception: (1) perception has a specific kind of content and, (2) perception has a specific kind of vehicle. On the former approach, perception is said to have a richer content than belief. On the latter, a perceptual state is said to be more like an image, while a belief is more like a sentence.

A rich and non-conceptual content

Perceptual states are said to be a specific kind of mental states, in virtue of having a specific kind of content. This kind of content is characterized by its *richness* and *non-conceptuality*.

A rich content. Through perceptual experience, we are aware of a lot of details in our environment. It probably means that perceptual states represent a lot of very fine-grained properties of the world – a very particular shade of red, every single books in the shelves in front of me, all the shades of grey that the light draws on the walls. In contrast, a belief is less detailed, and uses general terms to express the same state of the world, such as "there are a lot of books in various colors on the shelves". The richness of perceptual content is somehow inexpressible or inaccessible to the rest of the mind, such as beliefs, but maybe also memory, or some specific kind of consciousness (access consciousness). As Heck writes:

Consider your current perceptual state – and now imagine what a complete description of the way the world appears to you at this moment might be like. Surely a thousand words would hardly begin to do the job. (Heck, 2000, p.487)

Prima facie, this richness of content is a good candidate for being the hallmark of perception, as it is specific to the class of perceptual states. Furthermore, this specific trait would explain why perception is considered as a privileged source of knowledge: it represents the world in a lot more details than other states of the mind. However, the notion of richness is relative: the perceptual content is said to be richer than the content of other mental states, and especially beliefs. This could be problematic as it suggests that the difference between contents is a difference in degree, and not in kind. This would prevent richness of content to be a good criterion of individuation of perception as a natural kind, as this difference would not be fundamental enough - it would not mark a difference in kind.

However, there is a sense of richness that may be taken to not be relative. Let's distinguish two senses of richness. The first notion, *quantitative richness*, is roughly the idea that a perceptual state has a content that contains *much more information* than other mental states, and especially beliefs. The second one, *qualitative richness*, expresses the idea that perceptual contents are *more fine-grained*, more detailed than the content of other mental states such as beliefs. A perceptual state can be said to contain *more shades of red* than a belief (quantitatively) or a perceptual state can be said to contained *a more specific shade of red* than a belief (qualitatively). The difference in qualitative richness is a difference in terms of determinacy. Beliefs have determinable properties in content (red) while perception has determinate properties (red⁵⁸ or crimson red). Determinate and determinable properties are relative: 'colored' is a determinable for 'red', and 'red' is a determinable for 'crimson red'. Perception would be qualitatively rich by having a content with *maximally determinate* properties (i.e. properties that are not determinable ones)⁶.

Most of the authors that have discussed the richness of perception do not specify which sense of richness they use and both senses are often entangled. I take here the example of richness in Dretske (1981b).

Dretske defines our perceptual states as “informationally profuse and specific” (Dretske, 1981b, p.145). This rich character of perceptual states is given, according to him, through our perceptual experiences:

Our own perceptual experience testifies to the fact that there is more information *getting in* than we can manage to *get out* (Dretske, 1981b, p.145, Dretske's emphasis).

There are two ways of reading this characterization of perceptual states. In the first reading - the quantitative one - the content of perceptual states is more abundant than the content of cognitive states. In other words, there is more information in perceptual states than in cognitive states. This reading is supported by other parts of Dretske's book:

the visual system is processing and making available a quantity of information far in excess of what the subject's cognitive mechanisms can absorb [...].

⁶The distinction between determinate and determinable is different from the distinction genus/species. You can describe a species as a genus with specific properties. Describing a species requires therefore *more* information. However, you cannot add a property to 'red' to find 'crimson red'. This is not a question of quantity of information, but of (in)determinacy of it.

Our sensory experience is informationally rich and profuse in a way that our cognitive utilization of it is not. (Dretske, 1981b, p.146).

The second reading of Dretske focuses on the idea that the content of perceptual states is more “specific” (Dretske, 1981b, p.145) than the content of beliefs. For Dretske, the information contained in a perceptual state is translated into the content of beliefs through a process of conceptualization, which consists in ignoring differences and generalizing. For example, when I am conceptualizing my visual experience of a cup of coffee, I am ignoring the differences between this particular cup of coffee and all the other cups of coffee I ever saw in order to classify what I see as a cup of coffee. Therefore, the information in cognitive states is more general than in perceptual states, and this generality prevents us to grab in our cognitive states the fineness of grain of perceptual states. This is the qualitative reading of richness of content.

Perceptual states may represent the same type of things as beliefs, but represent a lot more of them. This large number of properties represented explains the impression of richness. On the other hand, it is possible to accept that the richness comes from the determinacy of properties represented, without them being more numerous. They are more determinate and precise in their nature. These two senses of richness are independent in principle ⁷.

Quantitative richness meets a problem for delimiting perception, as it is a difference in degree. What is the threshold of the number of things represented in the content of a mental state for it to be a perceptual state? All answers to this question seem arbitrary. Furthermore, as it is not a difference in kind, quantitative richness is not fundamental enough for being the criterion of individuation of a natural kind. Whatever the number of things represented in the content of a perceptual state, subtracting one thing will never change the nature of the state from perceptual to something else. On the contrary, qualitative richness can be considered as a difference in kind, as the difference is that properties represented in non-perceptual states are determinable, while properties represented in perceptual states are determinate. Generalizing the determinate content of a perceptual state to make it determinable is typically taken to be the process that transforms a perceptual state into a belief. Some equates this process with the process of *conceptualization* (Dretske, 1981b). This is why perceptual content has been said to be non-conceptual, while contents of beliefs are typically conceptual. This is the object of the next subsection.

A non-conceptual content. The content of perceptual state has been said to be non-conceptual, by contrast with the conceptual content of beliefs. The change from non-conceptual to conceptual content would therefore mark the limit between perception and beliefs, and *non-conceptuality of content* would be a good criterion of individuation of perception as a natural kind. What does it mean to have a non-conceptual content? Let’s begin by the conceptual content of belief.

The content of beliefs is conceptual, because it has a certain structure (propositional structure), with *concepts as components*. Intuitively, when I believe that my computer is grey, I need to have concepts like COMPUTER and GREY. These concepts constitute partly the content of my belief. By contrast, the content of perceptual states is not composed by concepts (or at least not only), as someone could see a grey computer without possessing

⁷It is not the case in Dretske’s work - in his general system, quantitative and qualitative richness are linked together through the concept of format: the change of format from perception to cognition explains both quantitative and qualitative richness (see the next section on format).

neither COMPUTER nor GREY. That is how I understand the notion of non-conceptual content here.

Two short remarks about this characterization of non-conceptual content:

- I do not need a specific theory on what concepts are here. I assume that they are mental representations and that having a concept means being able to build some reasonings about it and recognize/identify its objects. This is not a formal definition, but it is sufficient for my purpose here.
- My way of understanding non-conceptual content is close to the one first introduced by Evans (1982). Another distinction has been introduced later by Heck: the distinction between non-conceptual *content* versus non-conceptual *state*. The content view consists in claiming that perceptual states have a different kind of content than beliefs. This claim is about the kind of elements, if any, that composed the content of these states. On the contrary, the state view is the claim that (1) perceptual states and beliefs have the same kind of content, but (2) perceptual states are concept-independent, when beliefs are concept-dependent. It means that one cannot have a belief about P if she does not possess the concept P, whereas she may have a perceptual state of P without possessing the concept P. As Heck writes:

A perceiver can be in perceptual states a faithful report of whose content would necessarily employ concepts she did not possess. (Heck, 2000, p.484).

Evans' way of thinking of non-conceptual content - as well as mine - falls under the definition of the content view, even if Heck claims that some of Evans' claims can be interpreted as a support to the state view (Heck, 2000, p.485). I am not interested here in the state view, because it is incompatible with the claim that concepts are elements of beliefs (which is one of my assumption). Heck acknowledges this point:

I suspect that the state view is indefensible-even incoherent, if coupled with the claim that the contents of beliefs are conceptual (Heck, 2000, p.486, note 6).

This is why I focus on the following understanding of the non-conceptual content of perceptual state: the content of a perceptual states has components and some of them are not concepts.

As I evoked earlier, the intuitive reason for positing non-conceptual content is that when I see a computer in front of me, even if I do not know what is a computer, if I never saw computers before, I nevertheless see 'something' in front of me. My perceptual state is still representing the world in some sense, so it still has a content. However, this content cannot have COMPUTER as component, since I don't have the concept COMPUTER. One easy answer is that the content may be composed of other concepts such as ELECTRONIC OBJECT, or even GREY THING. This is where qualitative richness plays a fundamental role in the justification of non-conceptual content of perceptual states: qualitative richness is best explain by non-conceptual content. The level of determinacy in the content of a perceptual state is a sign that elements of content of a perceptual state are not concepts, because none of our concepts have this level of specificity. A good example is the contrast between our color concepts and our ability to discriminate much more precise shades of colors though perception. Here is how Tye expresses it:

Color experiences, to take one obvious case, subjectively vary in ways that far outstrip our color concepts. For example, the experience or sensation of the determinate shade, red²⁹, is phenomenally different from that of the shade, red³². But I have no such concept as red²⁹. So, I cannot see something as red²⁹, or recognize that specific shade as such. For example, if I go into a paint store and look at a chart of reds, I cannot pick out red²⁹. My ordinary color judgment are, of necessity, far less discriminating than my experiences of color. (Tye, 2000, p.61).

Beliefs or judgments that we can have about colors involve our concepts of colors. But these concepts of colors are too general, not specific enough to grasp what is present in the content of a perceptual state. That is why elements of perceptual states, such as a specific shade of red like red²⁹, have to be non-conceptual.

There are other arguments in favor of non-conceptual content of perceptual states, beside qualitative richness. The most famous are (1) the argument about animals and children, and (2) the argument of conceptual learning.

The argument about animals is the following: (i) some non-human animals do not have concepts; (ii) non-human animals have the same kind of perceptual states than human (at least to some extent); (c) therefore, human perceptual states cannot be conceptual. This argument is defended by Peacocke (2001):

While being reluctant to attribute concepts to the lower animals, many of us would also want to insist that the property of (say) representing a flat brown surface as being at a certain distance from one can be common to the perceptions of humans and of lower animals. [...] If the lower animals do not have states with conceptual content, but some of their states have contents in common with human perceptions, it follows that some perceptual representational content is nonconceptual. (Peacocke, 2001, p.614), see in (Speaks, 2005, p.6).

The argument about children is analogous to the argument about animal. We want to accept that children have perceptual abilities, that are to some extent similar to ours. However, children do not have a full-fledged conceptual apparatus, like human adults. Therefore, the perceptual states of children, that represent the world in some sense, must have a non-conceptual content.

A last consideration comes from the idea that concepts are learnt. Some can be learnt from the combination of other concepts, but some must be learnt 'from scratch', based on something that is not conceptual. The best candidates for being the precursors of concept are the elements in the content of perceptual states. For example, it seems possible that children learn what a dog is by perceiving dogs that adults show them. Therefore, the content of perceptual states (or at least some elements of content) are not concepts, as they are pre-conceptual. Being pre-conceptual is one way of being non-conceptual.

So far, the potential good hallmarks of perception related to the content of perceptual states are the *qualitative richness* of content, and the *non-conceptuality* of content. Both are usually related (but not necessarily). Let's move on now on specificities of the vehicle of perceptual representations.

An analog and iconic format

The format of a representation corresponds to the type of vehicle of the representation. For example, the same picture of my cat can be printed with an ID-photo format, or as a poster to hang it up on my wall. These two representations of my cat would have the same content, but not the same vehicle (as they are printed on two different pieces of paper). More importantly, they do not have the same *type* of vehicle, as one is an ID-photo, and the other one a poster. Therefore, they do not have the same format.

Saying that perceptual states have a specific format means that perceptual states have a specific *kind of vehicle*, compared to the other mental states. Intuitively, the difference between a perceptual state -let's say a visual state - and a belief, is that a perceptual state is more like a picture, when a belief is more like a sentence. When I see my cat in front of me, I can best express my mental state by a painting, than a sentence. But if I believe that my cat is at home when I'm not, the best way to express it is with the sentence "I believe that my cat is at home". Two technical terms have been used to grasp this intuition: (1) analog format and (2) iconic format.

An analog format. A perceptual state is said to have an analog format of encryption, while beliefs have a digital format. The two notions of analog and digital formats have been elaborated mostly by Dretske (1981a) and have to be understood in relation with his general theory about perception and knowledge. Roughly, Dretske uses a theory of information to explain perception and the constitution of knowledge, reducing the associated mental states to information carried by signals. There is however a difference between the information coded in perceptual signals and post-perceptual signals: their format of encryption. An information is coded in *digital* form when it is *the most precise information carried by a signal*. On the contrary, an information is coded in an *analog* form when *there is a more precise information carried by the signal*.

Dretske uses the example of the information F: "the cup has coffee in it" (Dretske, 1981b, p.136). If this information F is conveyed through the sentence S "the cup has coffee in it", there is no more information than F in S. F is then coded in digital form. However, if F is conveyed through a picture P of a cup full of coffee, F will not be the only information conveyed by P. P also conveys information about the size, the color and the shape of the cup, as well as the amount of coffee in it. F is then coded in analog form in P. One important consequence of these definitions of analog and digital is that a signal always carries some information in analog form and some information in digital form. For example, the picture P carries information in digital form. However, a sentence expressing the digital information of a picture would need to express all the details of the picture and would probably be very long and complex, as it need to be the most precise information carried by the picture.

In Dretske's work, the translation from analog to digital format is done through a process of conceptualization, which implies a loss of information:

To describe a process in which a piece of information is converted from analog to digital form is to describe a process that necessarily involves the loss of information. Information is lost because we pass from a structure [...] of greater informational content to one of lesser information content. (Dretske, 1981b, p.139)

This loss of information is due to a limitation of our cognitive capacities:

The sensory systems overload the information-handling capacity of our cognitive mechanisms so that not all that is given to us in perception can be digested (Dretske, 1981b, p.144).

According to Dretske, the change of format comes with a change of kind of content, from a rich content to a parser one. The change of format of encoding comes therefore with a change of richness of content, and a change from non-conceptual to conceptual content. Dretske does not explicitly prioritize one of these differences (format of informational encryption, richness of content or conceptuality of content) as a more important property than the others to mark the specificity of perception. However, it is possible to interpret that it is the *change of structure* (analog to digital) that causes the loss of information (a difference in content). In this reading, the format of the vehicle (how the information is carried by the signal) is more fundamental than the content (what information is carried by the signal). Therefore, Dretske can be understood as a proponent of format of encryption as the hallmark of perception (even if richness and non-conceptuality of content are also properties of perceptual states in his view).

An iconic format. The other notion used to describe this specificity of format is iconicity. Perceptual representations are *iconic* (Fodor, 2007; Burge, 2010; Block, 2014). By contrast, cognitive representations are *discursive*.

The distinction between discursive and iconic representations lies in their different kinds of compositionality. The compositionality of a representation is the way the structure and content of its parts contribute to the structure and content of the whole representation. Let's take an example. The paradigmatic example of a discursive representation is a sentence of a natural language, as S: "My house is red". This is in virtue of the meaning and grammatical role of each of the words that compose the sentence S that S means that my house is red.

Now the distinction between the compositionality of discursive representation and iconic one is the following:

- Discursive representations have some parts (but not all) that are constituents with intrinsic syntactic and semantic properties. These constituents contribute exhaustively to the compositionality of the discursive representation. The decomposition of the representation into these constituents is called the canonical decomposition of the representation.
- Iconic representations have no canonical decomposition, in the sense that all the parts of the representations contribute to the compositionality of the whole representations. None of the parts of an iconic representation has a privileged position in terms of contribution to the structure and content of the whole representation.

The paradigmatic case of an iconic representation is a picture. Contrary to the sentence "my house is red", a picture of my red house has no canonical decomposition: all parts of the picture, whatever its size or shape, contribute to the representation of my red house in terms both of structure and content. This specificity of iconic representations is also expressed by Fodor's picture principle:

The picture principle: if P is a picture of X, then parts of P are pictures of parts of X. (Fodor, 2007, p.7).

A picture of my red house follows the picture principle: parts of the picture are pictures of parts of my house.

Another characteristic of iconic representations is their *holistic* format of representation. A part of an iconic representation represents many properties at once (Quilty-Dunn, 2019). For example, a small portion of the picture of my house represents at the same time the color and the texture of my house. On the contrary, a constituent of a discursive representation stands for a single property or individual. Quilty-Dunn sums up:

The core difference between iconic and discursive representations is that the latter do, and the former do not, break down into recombining parts each of which can stand uniquely for a particular property, location, or individual. (Quilty-Dunn, 2019, p.8).

The holistic characteristic of iconic representations refers to the same phenomenon than Dretske's analog format. A property is coded in an analog format when it comes within a bundle of other properties, and conveying more than one property at a time is precisely the holistic characteristic of iconic representation.

Iconicity, like analog format, goes hand in hand with non-conceptual content. For Fodor (2007), "it's in the nature of iconic representation not to be conceptualized" (p.5). The core idea is that concepts have intrinsic syntactic and semantic properties (at least in Fodor's Language of Thought theory), as well as a 'power of individuation' (by applying a concept, we put conditions of individuation on what we subsume under the concept). Concepts are therefore constituents of a discursive representation. If a representation is iconic, it does not have a canonical decomposition into constituents, and therefore it cannot have concepts as constituents (in Fodor's sense).

There is good empirical evidence that some perceptual representations are iconic representations (Fodor, 2007; Quilty-Dunn, 2019).

First, iconic representations are expected to be *item-insensitive*. As they do not individuate features or individuals, they can represent as many properties/individuals as possible without asking for more resources. An intuitive example is that a photograph of a herd of giraffes does not require more space or ink than a photograph of a single giraffe (Fodor, 2007, p.15). There is good evidence that we have a high-capacity short-term perceptual memory. Its high-capacity in terms of number of items supports the idea that stored representations are iconic. For example, in the famous Sperling experiments, subjects were briefly exposed to a visual stimulus of 3 rows of 3 letters. Subjects were able to report correctly only 3 to 4 letters. However, when asked to report one of the three lines cued shortly *after* the disappearance of the stimulus, they were as good as reporting any of the row. It suggests that the visual stimulus was shortly stored *as a whole* (and therefore in iconic format), for the subject to be able to report any row of letters. (Sperling, 1960).

Another source of evidence of the iconic format of perceptual representations is *ensemble perception* (Quilty-Dunn, 2019), which consists in the perception of a great number of items - such as 16 circles of varying diameters (Ariely, 2001), that gives rise to an accurate representation of an average feature of the group of items (in Ariely's experiment, the mean of diameters). However, subjects are unable to recognize a single item. It suggests that all the 16 circles were represented, but only as a whole, in an iconic format, and not

individually.

Block (ms.) and Burge (2010) both argue in favor of iconicity as (at least partial) hallmark of perception. For example, Burge says that the structure of a mental state - its format - is part of its nature:

The structures of mental states include the structures of their representational contents. These structural features are not only constitutive. They are aspects, ‘parts’, of the states’ natures. The representational content of a belief and its structural elements are aspects of the nature of the belief [...] They are part of what it is to be that mental state. Psychological explanation makes explanatory use of these structures. (Burge, 2010, p.67)

Even if he do not use ‘iconicity’ explicitly, he later acknowledges:

Perceptual representation has a structure relevantly like that of pictorial representation [such that] just as one cannot draw a line without drawing its length, shape and orientation, one cannot visually represent an environmental edge as such without representing its length, shape, and orientation, as such. (Burge, 2014b, p.493), see in (Quilty-Dunn, 2019, p.8).

Both Burge and Block advocate the representational approach of the hallmarks of perception, in which the distinctive feature of perception is its content and format, e.g. its non-conceptuality (of content) and iconicity (of format).

2.3.3 The architectural hallmark

I presented the representational strategy for finding the limits of perception, consisting in finding a specificity of perception in its representational content or vehicle. By contrast, the second approach focuses on the architectural specificity of perceptual processes. It has been argued that perceptual processes are *modular*, such that they are architecturally distinct from other mental processes. This thesis has been defended in the eighties by Fodor (1983), who characterizes modularity of a mental subsystem with nine properties:

A modular system is **(1) domain specific**, which means that its processes are triggered by a specific range of inputs (transducer’s outputs) and its information is proprietary; A modular system generates **(2) shallow outputs** (relatively simple outputs). The best candidates for perceptual outputs, according to Fodor, are representations of basic level categories. These are the categorizations “that you would want the input systems to deliver assuming that you are interested in maximizing the information per unit of perceptual integration” (Fodor, 1983, p.97). The computational architecture of a module is such that its processes exhibit **(3) limited central accessibility** and are **(4) informationally encapsulated** from the rest of the mind. The processes are **(5) mandatory** and **(6) fast**. Finally, the neural architecture of a module is **(7) fixed**, and exhibits **(8) characteristic and specific breakdown patterns**, as well as **(9) characteristic ontogenetic pace and sequencing**.

Among this quite long list, Fodor specifies that a system can be considered as modular when it manifests most of these features in a reasonable extent. He indicates, however, that the informational encapsulation of the processes is essential to the modularity of a system:

The informational encapsulation of the input systems is, or so I shall argue, the essence of their modularity. (Fodor, 1983, p.71)

It is worth noting that modularity is not a proprietary property of perceptual processes as Fodor attributes modularity to the language analyzer too. Furthermore, in *the massive modularity theory* - the theory that the whole mind is organized exhaustively into a multitude of modules - perception is only a module among many. Modularity of perception is therefore not in itself a good criterion of individuation of perception as a natural kind. However, by adding the functional information of types of inputs and outputs of perception, the characterization becomes specific to perception. Perceptual processes would therefore be the mental processes that constitute the module with transducer's output as input and basic-level categories as output.

In this characterization of perception, the core characteristic for distinctiveness of perception is its informational encapsulation, resulting in a clear border between the class of perceptual processes and the other mental processes. It is because the perceptual module is informationally encapsulated that all processes constituting this module are proprietarily perceptual.

There are three main motivations in Fodor's defense of the informational encapsulation of perception. First, perceptual processes are too fast in generating an output for them to use information from other parts of the mind. If perceptual processes have access to the background knowledge and values of an organism, these would be much slower. Therefore, perceptual processes must only have access to a limited body of information. Fodor says:

Speed is purchased for input systems by permitting them to ignore lots of facts. (Fodor, 1983, p.70)

This lack of accessibility to background knowledge is best explained by the informational encapsulation of perceptual processes. Therefore, the best explanation for the high speed of perceptual processes is their informational encapsulation.

The second motivation is evolutionary. Perception is/has to be reliable about what is in the world, and this reliability can only be preserved if perception is encapsulated. This reliability is supported by evolutionary considerations, as Fodor points out:

A condition for the reliability of perception, at least for fallible organism, is that it generally sees what's there, not what it wants or expects to be there. Organisms that don't do so become deceased. (Fodor, 1983, p.68)

Reliability of perception is therefore also best explained by its informational encapsulation.

Finally, the last motivation is the persistence of (visual) illusions. Some visual illusions persist even if one is aware of their illusory nature. For example, most of us can't help but seeing the length of Müller-Lyer's arrows as different, even though we know they are identical. However, if perception has access to one's knowledge, the phenomenon of the persistence of visual illusion would not exist. Therefore, in these cases, perception has no access to one's knowledge. The best explanation for the fact that perception has no access to one's knowledge in these cases is that it never has access to it (namely, that perception is informationally encapsulated).

The architectural strategy is today defended by numerous philosophers and psychologists, even if empirical evidence points toward cognitive penetrability of perception (Quilty-Dunn, 2019; Siegel and Byrne, 2017). The post-fodorian proponents of the modularity of perception can be divided into two categories:

1. The first category is composed by the encapsulationnalsists, or in other words, people who maintain the encapsulation of information as the mark of perceptual modularity, and consequently of perception. Pylyshyn (1999), Raftopoulos (2001), Firestone and Scholl (2016) and Mandelbaum (2017) are all encapsulationnalsists, with variable positions.
2. The second category is composed by the revisionists. They preserve the modularity of perception, but at the cost of its informational encapsulation (Ogilvie and Carruthers, 2016; Barrett and Kurzban, 2006).

The encapsulationnalsist The novelty of post-fodorian encapsulationnism, compared to Fodor's position, is the need for its proponents to defend their position against the increasing number of evidence of cognitive influence on perception. Pylyshyn (1999), and more recently Firestone and Scholl (2016), argue that all cases of purported cognitive penetration can be explained by alternative accounts preserving informational encapsulation of perception. For example, cognitive penetration is often deduced from results that can be explained by intra-modular top-down effects, or by influence of cognitive states on post-perceptual states such as perceptual judgment.

Firestone and Scholl (2016) also highlight some general and pervasive pitfalls among the empirical works on cognitive penetration, which invalidate – or at least considerably weaken – the results of the experiments, such as a confirmatory research strategy or the lack of presentation of alternative explanation.

It is worth noting that this strategy of defense is negative. The general argument is that there is no good evidence for denying the fodorian architectural and functional characterization of perception as modular. Recently, the fodorian modularity of perception has been positively defended by Mandelbaum (2017). He accepts some results of empirical studies – e.g. that one's categories and concepts influence perception - but argues that these influences are intra-modular. In a sense, he widens the fodorian module of perception to incorporate some processes of categorization. The main argument relies on the high speed of visual categorisation, which precludes the interpretation of data as cognitive influences. He writes:

Why should we see such categorization as being underwritten by perceptual processes? Why not just take this datum as evidence for top-down penetration? The main reason not to is based in neurological wiring: 13ms is just too short a time to allow for top-down connections to take hold. (Mandelbaum, 2017, p.12)

Mandelbaum's position preserves the informational encapsulation of the perceptual module but entails a slight modification on the architectural characteristics of a module. As concepts and categories are part of the perceptual module, but are also sensitive to perceptual learning, the neural and computational architecture of the perceptual module is

not fixed anymore.

The encapsulationists maintain that the informational encapsulation of perceptual processes is what best characterizes the distinction between perception and the rest of the mind, as it is what characterizes the limits of the perceptual module, and explains properties of perception, such as visual illusion, reliability and even rapid visual categorization.

The revisionists The second strategy consists in rejecting informational encapsulation of the perceptual module. Contrary to the fixed architecture discarded by Mandelbaum (2017), the informational encapsulation of a module is considered by Fodor as its essential feature. That is why modularity without encapsulation can be called *non-fodorian modularity* and its proponents are revisionists with respect to Fodor's traditional picture of modularity.

Revisionists reject encapsulation based mostly on the empirical evidence of cognitive penetrability of perception. However, they maintain that the notion of modularity is still crucial in the characterization of perception. To do so, Carruthers (2006) disentangles two notions of encapsulation. The first one, called *narrow-scope encapsulation*, characterizes systems that are not affected by information coming from elsewhere in the mind. This is the encapsulation of fodorian modules. The second notion of encapsulation, *wide-scope encapsulation*, describes systems with the ability to reach and use information from outside during the course of their processing, but cannot use them all at once. According to Carruthers, the visual system is modular, but without narrow-scope encapsulation. His notion of modularity is then much weaker than Fodor's one. It keeps however the following features: mandatoriness, conscious inaccessibility of its internal operations, domain specificity and dissociability (Carruthers, 2006). This last feature has to be understood at the functional level. The visual system is indeed individuated functionally "as the set of brain-mechanisms specialized for the analysis of signals originating from the retina" (Ogilvie and Carruthers, 2016, p.722).

According to revisionists, modularity with wide-scope encapsulation is therefore sufficient to distinguish the class of perceptual processes from other mental processes. This is why modularity is still a good hallmark of perception.

2.3.4 The function of perception

Another approach to find what distinguishes perception from the rest of the mind (and particularly from cognition) is to focus on its distinctive *function*. Recently, Phillips (2019) and Beck (2018) argue that this strategy allows us to draw the boundaries of perception. They both advocate that what makes a state or process perceptual is their function to represent the environment through the proximal stimuli that this environment produces on the subject transducers.

Perception is said to be stimulus-dependent (Beck, 2018) or stimulus-controlled (Phillips, 2019), which means that a perceptual state is *caused* by a distal stimulus via a proximal stimulus. The former is the object perceived in the environment, while the latter is the array of light that directly stimulates the retina (for the case of vision). The hallmark of perception, according to this approach, is that perceptual mechanisms have the *function* to

represent the environment based on information extracted from current proximal stimuli. One sometimes says that perception represents the environment 'online' or through a direct contact with it, contrary to beliefs or memories which may represent the world 'offline'. Phillips says:

A process is perceptual just in case it has the function of producing representations of environment entities by being causally controlled by those proximal stimuli that these entities produce. (Phillips, 2019, p.11)

It is far from being the first time that perception is characterized by its function (Graham, 2014; Armstrong, 2002). The particularity of these two approaches is that the function of perception is the distinctive element that allows us to draw the boundary between perception and the rest of the mind. Beck mentions it at the beginning of his article:

My primary aim will be to flesh out one proposal to demarcate perception from cognition: perceptual states are dependent on a stimulus, or are stimulus-dependent, in a way that cognitive states are not. (Beck, 2018, pp.1-2)

The function is also explicitly acknowledged as being the criterion of individuation of a natural kind. Phillips (2019) considers "the class of stimulus-controlled states as a superordinate Natural Kind" (Phillips, 2019, p.40).

The notion of *function* can be understood in different ways. For biological organisms, a function F of a trait (or organ) is the ability of this trait to contribute to the fitness of the organism, i.e. its ability to survive and reproduce (Graham, 2014). It can also be the past ability of the trait to contribute to the survival and reproductive success of the ancestors of the organism (Godfrey-Smith, 1994). Both Phillips and Beck are neutral concerning the specific notion of function they need to recruit. What matters is the ability of a trait to fail to perform its function in some cases, while remaining what it is. As Beck says: "A heart that stops pumping blood doesn't cease to be a heart; it just ceases to be a good heart" (Beck, 2018, p.8). Likewise, perceptual states may fail to represent environmental entities through proximal stimuli, e.g. illusions generated by noise in stimuli, but the output experiences and states would still be genuinely perceptual.

Furthermore, other states may represent the world in a way that is caused by the proximal stimulus, but if it is not part of their function to do so, they will not be perceptual. For example, I may form the belief that my socks are red based on my vision of my red socks. This belief is causally controlled by the proximal stimulus produced by my red socks, but this is not part of the function of this belief. This causal origin is not constitutive of my belief, but it is constitutive of my perceptual state.

In addition to the possibility of failure, the functional approach has the advantage to be compatible with the cognitive penetrability of perception, contrary to the architectural hallmark. A perceptual process can be influenced by a cognitive state/process and be causally controlled by a proximal stimulus. Nothing prevents multiple causal factors.

2.4 A hard choice

So far, I have set out several candidates for the prized position of hallmark of perception. These different alleged properties of perceptual states or processes are not necessarily in-

compatible, and it is possible to defend the idea that they are all conjunctively hallmarks of perception. This solution is close to Block's position (ms.). He argues that the difference in kind between perception and cognition is constituted by their difference in terms of contents and formats of representation, perception being non-conceptual and iconic. Furthermore, he accepts that perception is a module (even if he is closer to revisionists than encapsulationists by accepting some forms of cognitive penetration). Finally, he highlights the difference of function between perception and cognition. Perception is a provider of information about the environment, whereas cognition uses these information to reason and decide.

Even if it is in principle possible to accept all the previous candidates of hallmarks together, some of them are often considered as incompatible. In particular, the representational strategy is usually put in competition with the architectural strategy (Quilty-Dunn, 2019). The main reason is that they define classes of perceptual states with different extensions. In other words, they do not classify some mental states and processes the same way.

Hallucinations are good ambiguous cases in the literature, as they are sometimes considered as full-fledged perceptual states, and sometimes are excluded from perception. *Hallucinations* refers to situations in which it seems to the subject that something is perceived, but there is actually nothing to be perceived. In contrast, an *illusion* occurs when there is something to be perceived but the subject misperceives it, attributing properties that do not exist to an object that exists. For example, if I see a pirate ship on the pond of Meudon when there is nothing on the water, I'm hallucinating. However, I may be under the illusion that the group of paddle boats on the pond are a fleet of pirate ships. In this latter case, I saw correctly that there is something on the water, but I misperceived its attributes. This is an illusion. Both hallucinations and illusions are cases of non-veridical perception, and are opposed to veridical perception, in which a subject perceives correctly the object and properties present in her perceptual field. In some cases of hallucinations, the subject is not aware that she is hallucinating. The best explanation is that, in these cases, an hallucination is indiscriminable from a veridical perceptual experience. It tends to show that these two kinds of states have the same kind of content and format, namely non-conceptual content and iconic format. This is why the proponents of the representational strategy classify hallucinatory states as perceptual states.

However, contrary to veridical perceptual states, hallucinations cannot be (only) driven by environmental inputs and are arguably caused or at least influenced by cognitive states. Yet, according to the architectural strategy, the perceptual system is encapsulated from the rest of cognition. The upshot is that hallucinations are not perceptual for an encapsulationist, as it is not informationally encapsulated from cognition.

As the different approaches are incompatible, a criterion should be found to differentiate between them.

A first idea is to see if the proposed criteria allow us to clearly classify the controversial cases, without leaving borderline, hybrid or intermediate cases. However, each hallmark encounters difficulties when faced with certain controversial cases.

The functional hallmark Some non-perceptual states may partially have the function to be stimulus-driven. Beck discusses the case of perceptually grounded demonstrative

thoughts (Beck, 2018). If I look at my dog and think "this dog is brown", my belief is causally controlled by my current visual stimulus of my dog. Yet, the state I encounter is supposed to be a belief. Beck's solution is to acknowledge that perceptually grounded demonstrative thoughts have a perceptual element (the demonstrative part of the belief), and a cognitive element (their attributive part). The upshot is that the functional hallmark of perception does not classify states and processes in natural kinds with no hybrid or intermediate cases. I remind you that a natural kind cannot *partially* share its nature with non-members of its kind, so that it is impossible to be *partially perceptual*. If the function of perception is its hallmark, then the controversial case of demonstrative thoughts is unclassifiable.

The architectural hallmark The architectural approach also struggles with some potential hybrid cases. A first example is the existence of some illusions such as size distortion illusions, which globally affect visual experience (everything is perceived smaller or bigger than in reality). If it is acknowledged that the illusion is cognitively mediated, the proponents of the architectural approach would have to deny the perceptual nature of the experience. However, they cannot deny that the individual is indeed perceiving the world.

Other mental systems and capacities may prove to be difficult to classify as perceptual or cognitive, because they share some of the properties of the perceptual module, but also of cognitive capacities. This is the case of core cognition systems (Carey, 2009; Jenkin, 2020; Spelke, 2000). Core cognition systems are innate domain and task specific systems that takes perceptual information and generate basic conceptual representation, such as basic concepts of object, agent or numerosity. They share some features usually associated with perception, like encapsulated modularity. Furthermore, they seem to participate fully to perceptual phenomenology (we tend to see individuated objects and not indistinct flow of sensory information). On the other side, their inputs are not sensory information like retinal image, but already-processed representations such as 3D surface arrangements. Their inputs are therefore the output of some early perceptual systems, which makes core cognition closer to cognition than perception. Furthermore, core cognition is thought to be at the origin of the development of complex conceptual representations in human adults, which also places it on the side of cognition.

The representational hallmark The iconicity of format is actually quite restrictive for classifying perceptual states. First, philosophical and psychological literature shed light on mental representations such as maps and diagrams, that are genuinely participating in human cognition (Camp, 2007). These are neither iconic nor sentential, and these alternative formats of representation make the ensemble of states to classify more diverse, and therefore put pressure on the neat divide between two mono-format natural kinds.

Another controversial case for the representational hallmark is the representation of amodal contour. Amodal completion is the alleged representation of hidden parts of an object (either occluded, in the shadow or invisible from the point of view of the perceiver), so that the object is perceived as a unit with hidden parts (and not as unrelated distinct objects). Some cases of amodal completion are represented in early stage of visual processing, helping especially in the parsing of sensory information into perceptual units. However, they do not give rise to perceptual phenomenology (contrary to modal completion) (Briscoe,

2011). Amodal contour is arguably non-conceptual, but it is difficult to attribute it a format of representation as it is not a visible part of the icon. The proponent of the representational approach would therefore have difficulties to arbitrate on the perceptual or cognitive nature of amodal contour.

All the hallmarks generate borderline cases. That is why we need another way to arbitrate between them.

To recapitulate, we face a lack of consensus on the criterion of individuation of perception as natural kind and the suggested hallmarks are incompatible. Furthermore, none of them stands out from the crowd in terms of providing a clear demarcation without borderline cases.

What I suggest now is to explore the disagreements between the advocates of the different hallmarks. The goal is to identify a method or tools to arbitrate between them. I will show, however, that a kind of disagreement jeopardizes the natural-kind strategy itself.

Chapter 3

Object-level oriented disagreements

If people disagree about the hallmark of perception, it is partly because they disagree about perception and its characterization. This is the kind of disagreement on which I focus in this chapter: object-level oriented disagreements. By contrast, in the next chapter, I will focus on representation-level oriented ones. The former are disputes about the world, about the characteristics of our mental apparatus, or about grounding or supervenience relations. They are the ones that can be investigated, at least partly, by empirical methods. The latter raise questions about the theoretical tools used in scientific investigations.

I first focus on object-level oriented disagreements because, at first sight, they will inform us on disagreements that can be arbitrated by empirical studies (past, present or future). It would then be possible to use empirical results to find the hallmark of perception and feed the natural-kind strategy. I will show, however, that the literature on the hallmark of perception contains also a more theoretical disagreement on natural kinds that jeopardizes the natural-kind strategy itself.

Listing all the specific disagreements, or making an exhaustive list of all the positions, arguments and answers in favor of one candidate or the other would be very long. That's why I decided to organize the following two chapters into *four kinds* of disagreements. These disagreements are not specific to one hallmark, but are instead formulated with an indeterminate alleged hallmark H. That is why this organization allows me to travel across most of the literature on the subject with only a few examples. I let the reader classify whatever specific debate about hallmarks she has in mind into one of the following categories of disagreements. Finally, this organization gives me the ability to disentangle object-level oriented issues (the first three disagreements) from representation-level oriented one (the last disagreement).

The four kinds of disagreements are individuated by four objections that one can build against the generic claim (C): "H is the hallmark of perception". According to the assumptions presented in the preceding chapter, (C) presupposes other claims, that are:

- (C_1) H is a property of mental states or mental processes.
- (C_2) H meets the membership, mind-independence and fundamentality conditions, in virtue of which H is the criterion of individuation of a natural kind
- (C_3) H individuates perception

The rejection of (C_1) provides a first object-level oriented issue, namely whether or not H is a property of mental states or mental processes. This is what I call a *disagreement on properties*.

There are two main objections against (C_2) , each of them produces a different kind of disagreement. The first objection consists in arguing that H does not meet the membership, mind-independence or fundamentality conditions. Objections along this line generate what I call *disagreements on 2nd-order properties*.

The second objection against (C_2) states that the membership, mind-independence and fundamentality conditions are not the conditions in virtue of which a property constitutes a criterion of individuation of a natural kind. I call it the *disagreement on natural kinds*.

Finally, the last disagreement consists in putting into question (C_3) and arguing that H does not individuate perception (but something else). I come back to this objection in the next chapter, as I identify it as a representation-level oriented disagreement.

These two first disagreements are harmless for the natural-kind strategy. Their respective resolution would help find the hallmark and therefore will feed the second step of the natural-kind strategy. The third one, however, jeopardizes the natural-kind strategy in putting in doubt its assumptions.

In what follow, I first present the disagreements on properties and 2nd-order properties (section 3.1), then I focus on the disagreement on natural kinds (section 3.2). I conclude that it would be better to explore other strategies to draw the boundaries of perception. In the next chapter, I will use the fourth disagreement as a springboard to elaborate a new strategy.

3.1 Two kinds of empirical disagreements

The two first kinds of disagreements are the following:

1. **Disagreement on properties** : H is not the hallmark of perception because H is not a property of mental states/processes.
2. **Disagreement on 2nd-order properties** : H is not the hallmark of perception because H does not meet the membership, mind-independence or fundamentality conditions.

In this section, I present in more details these two disagreements and give several examples from the literature on perception.

3.1.1 Disagreement on properties

The disagreement on properties is a disagreement about whether a property x exists in the mental realm as a property of any mental states or mental processes. Here is a non-mental example. In his fictional novel *His Dark Materials*, Pullman (2012) offers a criterion for distinguishing adults from children: adults, contrary to children, attract some particles called 'dust' on them. If we suppose that the class of children and the class of adults are natural kinds, one can therefore adopt the property of 'attracting dust' as a good criterion

of individuation of the natural kind adulthood. However, *His Dark Materials* is a fiction, and there is no such thing as dust. Therefore, 'attracting dust' cannot be the hallmark of adulthood in reality.¹

Let's go back to the debate about the hallmarks of perception. I present here two examples (among many) of disagreements on properties: (1) McDowell's position against perceptual content being non-conceptual and (2) Prinz's position against perception being encapsulated.

McDowell's position In *Mind and World*, McDowell (1996) not only denies the existence of non-conceptual content in perception, but he denies the existence of non-conceptual content tout court. This is why his position is typically what I called a disagreement on properties about non-conceptualism.

McDowell argues that perceptual experiences have conceptual content. His defense is complex, but one of the argument can be reconstructed this way:

- The notion of 'non-conceptual content' is intrinsically inconsistent
- An inconsistent notion cannot refer to an existing property
- Therefore 'having non-conceptual content' is not an existing property (of any mental states).

The first premise needs explanation. Here is a brief tentative clarification of McDowell's position. Non-conceptual content, according to McDowell, grew in the literature in a reaction to a philosophical tension between two pressures that leads to an antinomy. The first pressure is the idea that our beliefs and reasoning, in which we use our conceptual capacities, must be in some sense connected to the world. When I believe that the TV is off, I must have a way to confront/justify this belief with respect to how things are in the world - the actual state of the TV - or I just think into the void. Our way to confront with the actual state of the world is through our senses and perceptual experiences. This is why perceptual experiences should be the place where beliefs find justification (a belief can of course be justified by another belief, but at some point, the justification process must reach perceptual experiences). In other words, it is through perception that beliefs can be connected to the world, and do not constitute a close system in the void. Beliefs must therefore be justified by perceptual experience.

The second pressure is the idea that perceptual experience and beliefs are part of different *logical spaces*. Beliefs are part of a normative domain, while perceptual experiences are descriptive. It means that our system of beliefs "includes relations such as one thing's being warranted or [...] correct, in the light of another" (McDowell, 1996, xv). On the contrary, perceptual experiences are empirical descriptions of the world, and do not constitute a system with such relations. It is therefore a fallacy to use perceptual experience as justification for beliefs. This is the antinomy: perceptual experiences both must and cannot be justification for our beliefs.

This tension generates two opposed positions :

¹One might also say that even if it exists, it is not fundamental, but this will be *another kind* of disagreement, namely a disagreement on 2nd-order properties.

1. A coherentism a la Davidson (1984), in which it is acknowledged that perceptual experiences cannot play the role of justifiers for any beliefs. Only a belief can justify another belief. The impressions that the world produces to our senses are outside epistemological concerns, they do not constitute reasons for our thoughts, they are only causally related to our system of reasoning. As an upshot, our powers of understanding are floating (almost) free from empirical considerations.
2. A non-conceptualism, which accepts the following move:

[W]hen we have exhausted all the available moves within the space of concepts, all the available moves from one conceptually organized item to another, there is still one more step we can take: namely, pointing to something that is simply received in experience (McDowell, 1996, p.6)

Therefore perceptual experience could justify beliefs. The idea is that we are able, through an operation of abstraction, to form observational concepts based on our rich perceptual experience. And from them, we can move along the inferential links of our conceptual capacities.

McDowell recognizes the advantage of non-conceptualism. Empirical reality is then infused into our reasoning system thanks to observational concepts that results directly from our received experience. However, according to him, the 'operation of abstraction', which is supposed to link non-conceptual content with observational concepts, cannot be part of a justificatory process. It is just a "brut impact from the exterior" (McDowell, 1996, 8). Therefore, non-conceptual content cannot play the role non-conceptualism would like it to play, namely being the contact with the world *that can constraint our judgments* in a way that it can constitute reasons for thinking (and not a brut force). Because experiences are just a 'brut force', and is not truly part of our justificatory processes, they do not actually have *representational* content. Therefore having a (representational) content, and being non-conceptual is impossible. That is McDowell conclusion:

If experience is pictured as input to spontaneity from outside, then it is another case of fraudulent labelling to use the word "content" for something we can even so take experience to have, in such a way that reason-constituting relations can intelligibly hold between experiences and judgements (McDowell, 1996, 53).

McDowell's solution consists in accepting that the realm of the concept runs all the way down to perceptual experience, that there is no 'Given' on which our conceptual capacities applies. Our conceptual capacities are present at the same time we receive information from the world. One of the non-conceptualist errors, according to him, consists in giving independent existence to things that are only conceptually separable, namely the receptivity of perception and its intellegibility, the fact that perceptual states receive information from the world, and the fact that it is understandable by our judgment as being about the world.

Prinz's position A second example of disagreement on properties can be found in Prinz's 2006 paper titled *Is the Mind really modular?*. Prinz rejects the fodorian picture of perception, in which perception is a module following the nine fodorian marks of modularity (see chapter 2, section 2.3.3). The objection raised is the following:

(P1.) A fodorian module is characterized by nine properties.

(P2.) These nine properties are not jointly present in the mind.

(C.) Therefore, perception, as a part of the mind, cannot be a fodorian module.

As in McDowell's example, the second premise attacks the existence of modularity in the mind – and not specifically as properties attributed to perception. That is why Prinz's objection is a disagreement on properties. Here is a summary of his argumentation in favor of premise 2:

- Localized domain specific area with characteristic breakdown are largely overestimated in the mind in general. First, anatomical area taking to be the center of some specific functions (such as Broca areas, or visual center of color vision) are often merely hotspots of activation during tasks, while experimental tasks involve large-scale overlapping networks. The moderately active locations are often disregarded, which leads to the misleading assumption that some tasks are performed by clear-cut brain areas. Secondly, brain lesions tend to produce multiple deficits that are often difficult to clearly identified, and are not stable across population of patients.
- Nobody denies the existence of mandatory, fast or shallow processes in the mind. For example, the system underlying circadian rhythm, or semantic priming are mandatory. There is little doubt that a ton of unconscious processes are very fast. However, there is no reason that automaticity calls for either rapidity of processing or shallow output. For example, circadian rhythm is a quite slow process, and semantic priming involves conceptual knowledge.
- The existence of fixed schedules of development of alleged 'innate' capacities of the mind has been exaggerated, and developmental studies tend to show that a large number of environmental variables may influence the development of perception, language, or core knowledge like mindreading. Furthermore, specialization is not incompatible with learning. A specialized function is therefore not necessarily innate. There is no sound argument in favor of any inaccessible processes or subsystems in the mind. The fact that we do not consciously access the processes/information does not entail that other subsystem do not have access to them.
- They may be some encapsulated subsystem in the mind, but they must be very few. Empirical literature tends to show more cross-talk between mental functions than encapsulation.

He concludes that there is little chance that any subsystem in the mind gather the nine fodorian features of a module, therefore perception has little chance to be modular. Actually, he even goes further: "My claim is not just that Fodor's criteria are not jointly satisfied by subsystems within the mind, but they are rarely satisfied individually" (Prinz, 2006, p.1).

3.1.2 Disagreement on second-order properties

The second kind of disagreement is the disagreement on second-order properties of H, the alleged hallmark of perception. Now, the existence of H is not put in doubt. The question is whether the hallmark-candidate possesses the second-order properties required for being a criterion of individuation of a natural kind. These second-order properties of the candidate must ensure that its associated class of mental states satisfies the membership, the mind-independence and the fundamentality conditions. It follows that the hallmark of

a natural kind must be **exhaustive** and **exclusive** for the kind (membership condition). It also must be **mind-independent** (mind-independence condition) and **fundamental**, i.e. linked to the nature of the kind (fundamentality condition).

I present here three examples in which the second-order properties of a hallmark of perception are discussed.

First, Quilty-Dunn (2019) argues that *iconic format* is not a good hallmark of perception, because it is not *exhaustive*. Contrary to a disagreement on properties, he accepts that some mental states - and even some perceptual states - have an iconic format. What he puts into question is the role of iconic format as the hallmark of perception, because, according to him, some perceptual states - *perceptual object representations (PORs)*- are not iconic, but discursive. According to Quilty-Dunn:

PORs are representations of objects that select individuals through the deployment of a visual index and store information about the objects they represent (Quilty-Dunn, 2019, p.14).

There are pieces of evidence that PORs are not iconic, in particular they do not represent holistically an individual object and its features. However, Quilty-Dunn argues that PORs are genuinely perceptual because they are integrated to perceptual processes and are encapsulated from cognition. What is important here is that Quilty-Dunn argues that the existence of PORs, as *non-iconic perceptual representations*, is an argument to dismiss iconic format as a good hallmark of perception. He says:

[D]emonstrating the presence of both iconic and discursive formats in perception provides abductive evidence in favor of perceptual pluralism and against the representational strategy (Quilty-Dunn, 2019, 4).

In this line of argumentation, Quilty-Dunn attacks the representational strategy on the fact that iconicity is not a property that describes *exhaustively* perceptual phenomena. This is why it is a disagreement on 2nd-order properties.

Another example concerns Beck's function of perception to be stimulus-dependent. Beck argues that "being stimulus-dependent" cannot be the distinctive characteristic of perception, because it is not specific to perception. As mentioned in the previous chapter, perceptually grounded demonstrative thoughts- like "that is blue" - are also stimulus-dependent. In other words, *stimulus-dependence* is not *exclusive* to perception. That is why Beck refines his proposition to offer a more specific characteristic for perception: perceptual states are states whose *all elements* have the function of being stimulus-dependent (Beck, 2018). Therefore, this latter proposition does not apply to perceptually grounded demonstrative thoughts.

Prinz (2006) provides a third example on failure of fundamentality. According to him, wide-scope encapsulation a la Carruthers (2006) does not give any information about the distinct nature of perception. It is no more than a functional decoupage of the mind, meaning that the mind is a whole composed of some parts, and that perception is one of this part. According to Prinz, this is not informative enough about perception for being a distinctive feature. In his words:

Some defenders of modularity are committed to nothing more than functional decomposition. They reject Fodor's list and adopt the simple view that the

mind is a machine with component parts. That view is uncontroversial. (Prinz, 2006, p.12)

Disagreements on properties and on 2nd-order properties are important for the natural-kind strategy, and their resolution is necessary to move forward. However, they are harmless for the strategy itself, in the sense that (i) there is good hope to resolve them with empirical studies and conceptual analysis and (ii) their resolution will provide a hallmark suited for being used in the natural-kind strategy. The disagreement on natural kinds, however, is more problematic, as it challenges the assumptions on naturalkindness. This is the subject of the next section.

3.2 Disagreement on natural kinds

In the case of a **disagreement on natural kinds**, the objection raised against (C) "H is the hallmark of perception" has the following form: H is not the hallmark of perception because, even if H meets the membership, mind-independence and fundamentality conditions, these are not the conditions in virtue of which a property is a criterion of individuation of a natural kind.

Disagreements on natural kinds are disagreements on the conditions a class of thing should satisfy for being a natural kind. As an upshot, they also are disagreements on the second-order properties that candidate-hallmark *must* have. Contrary to the previous disagreements on 2nd-order properties, it is not the possession of second-order qualities that is discussed, but their *legitimacy* as good qualities for making their bearer a hallmark of perception as a natural kind.

Following an essentialist characterization of natural kinds, I considered that the hallmark of perception should meet the membership, mind-independence and fundamentality conditions to be a criterion of individuation of a natural kind (chapter 2, section 2.1). This characterization has been criticized, though. I present here two of these criticisms among others.

Exhaustivity or exclusivity are not required for a hallmark This claim follows the idea that the membership condition is too demanding. The membership condition requires that (1) all and only the members of a natural kind share a property - or a cluster of properties and that (2) it is in virtue of this exhaustive and exclusive sharing that their grouping is non-arbitrary. One can reject the membership condition in a strong or a weak sense. The strong rejection claims that the members of a natural kind do not share some common property exhaustively and exclusively. The weak rejection claims that, even if the members of a natural kind share some microstructural properties, this is not in virtue of them that they are grouped together (but, instead, by some superficial or partial resemblance that can be neither exclusive, nor exhaustive). The two types of rejection imply that exhaustivity or exclusivity are not required for a property of perceptual states or processes to be the hallmark of perception.

Block is a proponent of the weak rejection of the membership condition for the hallmark of perception. He accepts that a fundamental property of a kind can also be shared by other kinds. In other words, the main condition for a class of things to be a natural kind for Block is what I called the fundamentality condition, i.e. all members of the kind

have to be *fundamentally* - in nature - similar. This is what he expresses by the use of the metaphoric notion of "joint-in-nature". According to him, there is a joint in nature between perception and cognition, namely a "fundamental and explanatorily significant difference in kind between perception and cognition" (Block, ms). However, this joint is "compatible with intermediate or indeterminate cases" (Block, ms), and in particular the fundamental properties of a kind can be found in members of other kinds. The difference lies in the fact that these properties are not constitutive, not in the nature, of the members of other kinds. Concerning the perception/cognition joint in nature, Block argues:

Perception is constitutively iconic, non-conceptual and non-propositional, but these properties are not sufficient for perception. For example, hallucination has all these properties but is not perception (Block, ms)

However, iconic format, non-conceptual and non-propositional content are the hallmarks of perception. The difference in kind is found in the fact that these properties are fundamental to perception, while they are not fundamental to non-perceptual states, such as hallucinations or other cognitive states that may use perceptual material non-constitutively. As an upshot, Block rejects that the hallmarks of perception have to be exclusive properties of perceptual states or processes.

In a different line of thoughts, Smortchkova and Murez (forthcoming) argue that what grounds a natural kind, especially in the mental realm, is a common underlying mechanism - instead of some essential properties. Following Boyd (1991) and his homeostatic property cluster theory, this underlying common mechanism is responsible for the clustering of common superficial properties. As an upshot, it is not the sharing of any specific property that constitutes a natural kind, but it is the naturalkindness of a group, in virtue of being produced by a common mechanism, that leads to the sharing of some properties by the members of the natural kind. However, these properties, called *syndrome* properties, are not systematic, as they may depend on the conditions and environment in which a mechanism takes place. That is why no property is *essential* to a natural kind. However, these syndrome properties may nonetheless be a good guide in the research of grounding mechanisms, as their clustering in an ensemble of mental states or processes may be the sign of an underlying common mechanism. As a consequence, the set of perceptual states may be a natural kind, whose limits cannot be determined by an exhaustive or exclusive property, as the sharing of a common property is not what grounds the naturalkindness of perception.

The epistemic role of a hallmark Another important way to disagree on natural kinds is to consider the epistemic role of the natural kind as constitutive of its naturalkindness. The epistemic role of natural kinds is present and even important in the classical view about natural kinds, as it is acknowledged that natural kinds allow scientists to perform law-like generalisations. However, this epistemic role of natural kinds is a consequence of their other characteristics such as their fundamentality (see chapter 2, section 2.1). On the contrary, some philosophers argue that the epistemic role of categories - their usefulness and productivity in science - is the ground on which they have to be considered as natural kinds (Quine, 1969).

In the debate about the boundaries of perception, this is not always clear whether some philosophers use the epistemic role of the class of perceptual state as a ground for its

naturalkindness - or the consequence of it. For example, Phillips (2019) explicitly says that:

[T]he most important requirement [for a class to constitute a natural kind] is that its members must share certain distinctive properties the appeal to which gives us scientific explanations and inductive inferences that we wouldn't otherwise have at our disposal (Phillips, 2019, p.2).

This quote is ambiguous between a classical use of the epistemic role of a natural kind, and a grounding use. However, he presents later his method of investigation, which is "to uncover those notions of perception and cognition that earn their explanatory keep in cognitive science" (Phillips, 2019, p.3). This quote suggests a use of the epistemic role of a natural kind as its individuating principle.

The use of the epistemic role of a natural kind as a grounding principle for naturalkindness, can be transmitted to its hallmark. A good hallmark would therefore be a characteristic of a natural kind with a privileged epistemic role. A good hallmark is in a sense a property that explains, or at least illustrates the epistemic value of the natural kind. This is how Phillips (2019) justifies that "stimulus-dependence" of mental states is a good property to delimit perception:

A perception–cognition boundary that is drawn in terms of stimulus-dependence will prove to be of theoretical interest (Phillips, 2019, p.3)

Both Phillips and Beck are proponents of a distinction between perception and cognition in terms of stimulus-dependence, but do not actually reject the representational or architectural candidates. They are both open to a *pluralism* of borders, i.e. the idea that the limits of perception may be multiple, and all legitimate. However, even if they do not openly disagree against specific candidates, they disagree about the conditions under which the limits of perception have to be drawn. This is a disagreement on natural kinds.

3.3 A threat for the natural-kind strategy

Disagreements on natural kinds put into question the conditions under which a property can individuate a natural kind. It reveals that the notion of natural kind I presented in the previous chapter (section 2.1) is actually controversial. The problem is that it threatens the natural-kind strategy in making it rests on controversial -yet, essential- assumptions.

I remind you that the natural-kind strategy consists in three steps: (1) assuming that perception is a natural kind, (2) finding the hallmark of this natural kind: the common and fundamental property of perceptual states, and (3) using the hallmark to establish the boundaries of perception.

If natural kinds are characterized with different conditions than the membership, mind-independence and fundamentality ones, the second and third steps of the strategy are threatened. First, the generalization of the hallmark from paradigmatic cases to controversial cases becomes dubious: the hallmark can indeed be present in paradigmatic cases and absent in controversial cases, without preventing them to be perceptual states. Secondly, if the hallmark is neither exclusive nor exhaustive to perceptual states, it becomes a bad tool for drawing the boundaries of perception. It will generate false positive and

false negative, i.e. states that display the hallmark without being perceptual, and states that do not possess the hallmark but are perceptual.

Disagreements on natural kinds are not discussing whether or not perception is a natural kind, nor whether there is one or several properties that can be considered as hallmarks of perception. However, in questioning the requirements for a property to be a hallmark, they weaken the appeal of the natural-kind strategy.

This difficulty justifies the idea to leave aside the natural-kind strategy and explore alternative strategies to draw the boundaries of perception. I use the fourth kind of disagreement, i.e. disagreements about the extension of the concept PERCEPTION, to draw a new strategy: conceptual engineering.

Chapter 4

Representation-level oriented disagreement

In this chapter, I show that a disagreement about the boundaries of perception consists in a metalinguistic negotiation about the concept of perception. A metalinguistic negotiation is a disagreement in which people do not disagree about the world, but about the representations they should use to describe the world (Plunkett, 2015). In other words, it is a representation-level oriented disagreement. Plunkett defines a metalinguistic negotiation this way:

A metalinguistic negotiation (or, equivalently, a normative metalinguistic dispute) is a dispute in which speakers each use (rather than mention) a term to advocate for a normative view about how that term should be used (Plunkett, 2015, p.832)

Here is an example of a metalinguistic negotiation in a non-philosophical context. Last winter, I went to Quebec to visit my sister-in-law. The temperature was around -2 Celsius degree and my little sister was visibly cold. My sister-in-law said about my little sister that she is 'frileuse', a negative french term to say that she feels the cold too easily. I answered: "No, she is not, it is just very cold here". Clearly, in this exchange, we were not disagreeing about the outside temperature, or the fact that my little sister was cold. We were disagreeing about the meaning of the term "frileuse", and whether the case of my little sister were falling under the extension of it. In other words, we held fixed our views about the world, and we used the word "frileuse" to talk about whether or not this term was adapted to describe it.

We could have explicitly talked about what 'frileuse' should mean, and my sister-in-law could have said: "In Quebec, she is frileuse". And I could have answered then: "Oh ok, in France, she is not". If our dialog went in this direction, we would have reached an agreement by pointed out regional differences in the meaning of the word 'frileuse'. As winter in Quebec is colder than the French one, the norm of cold sensitivity may be lower in terms of temperature, and consequently the meaning of a term such as 'frileuse' may be different. I could have add "frileuse is a French term, and therefore we should use the French meaning". This would have been chauvinistic, but it shows an important feature of a metalinguistic negotiation: this is a normative disagreement. The discussion is not about what the meaning of a term is, universally or in specific conditions, but about the meaning that a term *should* have.

A famous example of a metalinguistic negotiation in a scientific context is the recent debate about the status of Pluto as a planet. This debate may seem to tackle a metaphysical question about the nature of planets. What is a planet? What are the essential properties of a planet? However, whether Pluto is a planet or not is not an object-level oriented issue, as most scientists agree about the physical characteristics of both Pluto and the other planets. On the contrary, the object of the dispute is the concept *PLANET*, and whether or not this concept should include Pluto and similar celestial bodies. On 2006, August 26th, the International Astronomical Union in Prague has decided that Pluto was now a dwarf planet instead of a planet. Their decision, however, was more a revision of the meaning of the concept *PLANET*, than an evaluation of Pluto. The status change of Pluto was a consequence of a metalinguistic negotiation.

According to Plunkett (2015), there are two crucial features to recognize a metalinguistic negotiation from an object-level oriented disagreement: (1) At least one of the term is used with a different meaning by both parties in the disagreement and (2) the disagreement persists even when the two parties agree about the current meaning of the term (Plunkett, 2015, p.850).

I claim in this chapter that a kind of disagreement about the boundaries of perception is a metalinguistic negotiation about the concept *PERCEPTION*, which has consequences on how this debate can be, and should be (at least partly) resolved. I do not claim that resolving the metalinguistic negotiation will give a definitive answer to the boundary question (as they are some object-level oriented issues too), but at least it will clarify the whole debate.

Concepts like *PERCEPTION* (as well as *MEMORY*, *COGNITION* or *IMAGINATION*) are part of the theoretical background of cognitive science and are partly infused with preconceptions. Perception is intuitively passive, involuntary, and irrepressible, and these intuitions are probably partly grounded in a model of the mind in which perception is a mere receptive part of the mind. These concepts have an influence on the way we think about the mind, and disagreements about them are usually based on deeper disagreements about which model of the world is best to work with. Plunkett says:

The background reasons why speakers engage in [such disputes] are standardly bound up with their views about object-level issues (and not representational level ones about our words and concepts) (Plunkett, 2015, p.860)

Empirical apparatuses are influenced by the concepts we use, which are themselves influenced by our preconceived views about the world (and – fortunately – by empirical discoveries). However, this is an infelicitous circle of influence, in which our investigation about the world are tied to our preconceptions. Is it possible to break this circle, or at least control and/or soften the influence of preconceptions?

A radical reaction would be to eradicate all preconceptions from our investigating tools. Let's forge new theoretical tools, new concepts from scratch and let's do science with these clean, unbiased representations.¹

I do not think this solution can succeed. Preconceptions have a role to play in science and inexorably plays this role of guidance in scientific research. Building unbiased concepts,

¹This is the kind of eliminativism advocated by Churchland (1981).

metaphysically neutral, is a kind of misleading publicity, in which one shows a clean tool in the front window, but uses a dirty one in the back of the store. A moderate reaction would be to investigate our representations of perception, to inquire which intuitive models are in action in our theories, and find a way to evaluate them. We may be able to modify them, or at least choose the best one to work with. This will be the purpose of the second part of the dissertation.

In this chapter, I first present the fourth kind of disagreement I identify in the literature about the boundaries of perception (section 4.1), then I show that this disagreement is a metalinguistic negotiation (section 4.2). I end up in sketching an approach for resolving this metalinguistic negotiation (section 4.3).

4.1 Disagreement on extension

The fourth disagreement about the hallmark of perception is an objection against (C_3) "H individuates perception". It objects that H is not the hallmark of perception because H does not individuate perception as such. H may individuate a class of mental states. This class may even be a natural kind. However, this class is not *the class of perceptual states*. I call it the **disagreement on extension**.

The disagreement on extension is the disagreement on whether a property marks the *real* borders of perception, and not the border of another class of mental states. Contrary to the three other disagreements, it is acknowledged (or acknowledgeable), that the property exists, and is the hallmark of a natural kind - whatever the chosen conditions for being a natural kind are. The question is whether the class individuated by the property is perception.

I present two examples of such a disagreement in the literature on perception. The first one pertains to the perceptual nature of perceptual object representations (PORs). The second focuses on the controversial assimilation of visual perception to early vision.

First, Quilty-Dunn (2019) argues in favor of pluralism of format in perception with the following argument: (1) some perceptual states are iconic, (2) some perceptual states, namely perceptual object representations (PORs), are discursive, (3) Therefore, perceptual states can be either iconic or discursive (pluralism). The main objection against his paper is against (2), and affirms that PORs may not be *genuinely* perceptual. In his words:

One might object that PORs are not genuinely perceptual, and hence don't bear on the format of perception. (Quilty-Dunn, 2019, 21)

This objection is a disagreement on extension: discursive format may exist, may even be the hallmark of a natural kind, but mental states with discursive formats fall outside the limits of perception.

Secondly, some philosophers (Pylyshyn, 1999; Raftopoulos, 2001) defend the informational encapsulation of early vision. One could object that informational encapsulation is however not the hallmark of perception, because perception is not early vision. Ironically, Fodor himself, the 'Father of Encapsulation' refuses

to identify input analysis with perception. The point of perception is the fixation of belief, and the fixation of belief is a conservative process – one that

is sensitive in a variety of ways, to what the perceiver already knows. Input analysis may be informationally encapsulated, but perception surely is not. (Fodor, 1983, 73)

This disagreement could be misinterpreted as a disagreement on encapsulation as a criterion of individuation. In this interpretation, encapsulation would not be the hallmark of perception because it would not be exhaustive. As it leaves out a part of the class of states, it cannot be the distinctive characteristic of a natural kind. This would be a disagreement on 2nd-order properties. However, here, it is accepted (or could be accepted) that encapsulation is the distinctive characteristic of a natural kind, namely early vision. It is therefore not a disagreement on 2nd-order properties. Instead, what is rejected is that this natural kind is *perception*, that its boundaries correspond to the boundaries of perception. This is a disagreement on extension.

Disagreements on extension, contrary to the three other disagreements presented in the previous chapter, is not an object-level oriented disagreement. It may seem to be, as the location of the boundaries of perception could be seen as part of the *nature* of perception. But I will show instead that parties in disagreements on extension are actually disagreeing on the *concept of perception*. It means that disagreements on extension are representation-level oriented disagreements. More precisely, they are metalinguistic negotiations, i.e. *normative* representation-level oriented disagreements.

Here is one way to formulate the disagreement on extension: H is not the hallmark of perception because H do not individuate perception as such. Prima facie, the disagreement seems to concern a property of H, namely the property of individuating perception and could be thought as an object-level oriented issue, concerning a second-order property of H. However, this disagreement arises (or could arise) in the absence of disagreement about neither properties of H, nor the grouping of mental states into natural kinds. In other words, while agreement about physical and metaphysical properties of the world is reached, parties still continue to disagree about whether a class of states really is perception or not. This is evidence that the disagreement on extension is actually a disagreement about the meaning of "perception", instead of perception per se.

4.2 Metalinguistic negotiation

I claim that the disagreement on extension is actually a metalinguistic negotiation. Plunkett (2015) gives a list of four crucial features in order to identify metalinguistic negotiations:

1. The discussion is a **dispute**. People do not agree.
2. The discussion is a **disagreement**. It means that people are not just talking past each other, but are genuinely disagreeing about something in a rational way. A disagreement involves that someone holds the belief that P, while the other holds the belief that non-P. By contrast, a dispute may not. For example, a few years ago, my cousin told me a discussion she heard between her two little girls. One said: "You're a big liar" and the second answered "I'm not big!". In this example, the two girls quarreled with each other, but in a sense, the second was not disagreeing with

the first. They just did not talk about the same thing. Here is another example of a discussion:

- Julia: "I was at the bank at 3p.m. today"
- Joanna: "No, you weren't at the bank. I was there and I did not see you."

If Julia means she went at the river, while Joanna withdraw some money, they are not really disagreeing, they are just talking past each other. This is because Julia utters the proposition P: "I was at the bank₁", while Joanna utters the proposition non-Q: "You were not at the bank₂", in which bank₁ means the river side, while bank₂ is a financial business. In this condition, the verbal exchange between Julia and Joanna is a dispute, but not a disagreement. If, on the contrary, both Julia and Joanna are talking about the river side, then Julia says P while Joanna affirms non-P, this is a real disagreement.

3. There is a term used with a **different meaning** by both parties in the disagreement. This may seem to be very similar to the example of Julia and Joanna at the bank, but the difference is that they were not talking about the same thing, while here, people are referring to the same thing but with a different meaning. For example, both Joanna and Julia may have in mind the financial establishment, but Joanna was inside, talking to her adviser, while Julia stayed outside at the withdrawal automaton. Joanna may consider that the front of the bank is not part of the meaning of bank₂. In this sense, Julia and Joanna use the term "bank₂" in a different way, but are talking about the same thing.
4. The disagreement **persists** even when the two parties specify and agree about the current meaning of the terms. Let's continue with Julia and Joanna. If Julia specifies that she remained outside the bank₂, there are good chances that Joanna then agrees with Julia that she was at the bank₂, even if she did not see her. There is no reason for the disagreement to pursue after the specification of each meaning of bank₂. On the contrary, some disagreements may persist after the specification of meanings. For example, in a discussion about whether transgender women are women, one can affirm that they are, while the other claims the contrary. The first one may specifies that she defines a woman as a person with XX chromosomes, when the second one uses the term woman to talk about a constructed social role. There are little chances that one of the parties agrees with the other on her use of the term 'woman'. This is the sign that what is at stake here is the concept WOMAN.

The last feature shows that the disagreement is a matter of concept, and not only about the meaning of a term. If people disagree about the meaning of the term 'woman', it is because they disagree about the meaning the concept WOMAN *should* have. How WOMAN is used has indeed consequences on social justice and how transgender people are treated in our society. For now, I consider concepts as (1) tools used in thoughts and theories that (2) can be described (at least) by their intension or their extension. I will be more specific in the following chapters. The *intension* of a concept is its definition in terms of predicates. In my example about women, the intension of the concept WOMAN for the first party is "person with XX chromosomes". The *extension* of a concept is the ensemble of things or individuals to which the concept applies. The extension of WOMAN is therefore the list of all people that are considered as women following a specific intension of WOMAN. For the first party, transgender women are therefore not part of the extension of WOMAN,

while they are for the second one.

Disagreements on extension about perception show these four crucial features, and that is why they are metalinguistic negotiations about the concept PERCEPTION.

The existence of disputes among philosophers about the boundaries of perception is what I showed in the previous chapter and the beginning of this one. The number of papers on the subject, workshops and conferences is another sign that different positions are debated. A consensus does not generate such a level of activity among philosophers. For specific examples, I take you back to the previous section, in which the examples of early-vision and PORs are two clear disputes.

What is more important is to show that in the case of disagreements on extension, some philosophers argue in favor of a proposition P, while the others argue in favor of its negation non-P, that they use the term 'perception' differently and that the disagreement persists after explicitly formulating these meanings. I begin with a special case, involving natural kinds, and turn after to the general case.

The case of natural kinds Here is a form of disagreement on extension involving a natural kind X:

- Speaker 1 utters P: "The natural kind X is perception."
- Speaker 2 utters non-P: "It is not the case that the natural kind X is perception."

This is special form of disagreement on extension, as this kind of disagreement must not involve necessarily a natural kind. I begin by this special case because it disentangles more clearly the disagreement on extension from other disagreements involving natural kinds. A good example of this form is the disagreement about whether early-vision is (visual) perception. In this case, it is broadly acknowledged that the ensemble of states and processes that constitute early-vision is a natural kind. It is clearer then to see that the disagreement lies in the matching of what we know about early-vision and what we expect perception is, according to our concept PERCEPTION. In this disagreement, let's say that both parties agree on all facts about early-vision. So, their view about the world is hold fixed. Their disagreement lies in *what they mean by 'perception'*. In this form of disagreements on extension, the term 'perception' is thus used differently.

Finally, I need to show that the disagreement persists, even after people explicit their meaning of 'perception'. For this purpose, let's introduce two new terms to replace the contentious term of 'perception': perception₁ and perception₂. Now, we can reformulate the previous pair of propositions:

- Speaker 1 utters P1: "The natural kind X is perception₁."
- Speaker 2 utters non-P2: "It is not the case that the natural kind X is perception₂."

Now we see that nothing prevents logically speaker 1 to agree with speaker 2 and vice versa. However, a new form of disagreement arises. Speaker 2 may answer to speaker 1: "It does not matter that the natural kind X is perception₁. What matters is that natural kind X is not perception₂. The reason is that perception₂ is perception. Perception₁ is

not really perception". The existence of this new form of disagreement is the sign that the disagreement is actually a metalinguistic negotiation. The disagreement is about the concept of perception people *should* use.

Let's illustrate this on the early-vision example. I introduce two incompatible positions over the meaning of perception. According to Maar, perception is "defined as the process of information extraction" (Maar, 1982, p.1). Suppose this is the definition of perception₁. I take Fodor's definition of perception for perception₂, i.e. perception is the ensemble of processes that lead to the formation of perceptual belief, i.e. belief based on sensory input (Fodor, 1983).

Now, we have the following two paraphrases:

1. (PA1): "early-vision is (visual) perception₁" means "early-vision is the process of (visual) information extraction"
2. (PA2): "it is not the case that early-vision is (visual) perception₂" means "it is not the case that early-vision is the ensemble of processes that lead to the formation of beliefs based on visual inputs "

I suspect that Fodor would agree on (PA1) but would answer that this is not what defines perception, because "the point of perception is the fixation of belief, and the fixation of belief is a conservative process – one that is sensitive in a variety of ways, to what the perceiver already knows." (Fodor, 1983, p.73). Therefore, perception does not consist merely in informational extraction. So, even if an agreement is reached on the content of the paraphrase, there is still a disagreement on the validity of the paraphrase. What it means is that Fodor (arguably) would think that Maar *should not* use 'perception' the way he uses it.

The general case Here is the general form of a disagreement on extension:

- Speaker 1 utters P: "The kind of states/processes Y is perceptual"
- Speaker 2 utters non-P: "It is not the case that the kind of states/processes Y is perceptual"

In this disagreement, Y does not need to be a natural kind. It is still a disagreement on extension as the disagreement is about whether Y lies within the boundaries of perception or not. As we are in a disagreement on extension, speakers 1 and 2 agree on the existence and properties of Y, which means that the disagreement arises because 'perceptual' means something different for both of them.

Concerning the question whether the disagreement persists when the different meanings of 'perception' are clarified, the same point as before can be made. Speaker 1 may actually mean that Y is perceptual₁ and speaker 2 that Y is not perceptual₂. However, both speakers will defend that their meaning of perception is the *right meaning* of perception.

I present two examples to illustrate the general case. I hope the core of the disagreement - namely the disagreement about which concept of 'perception' should be used - will emerge clearly from them.

Object representations As presented in the previous section, Quilty-Dunn (2019) argues that some object representations are perceptual (what he calls "perceptual object representations" (PORs)). The particularity of these representations is their discursive format, which allows him to defend a format pluralism in perception (perceptual representations may be either iconic or discursive). I choose Spelke (1988) to embody the other party here. She argues that object representations cannot be perceptual because "perceptual systems do not package the world into units" (Spelke, 1988, p.229). Both Quilty-Dunn and Spelke acknowledge the existence of object representations, and agree - at least to some extent - on their properties such as their discursive format. However, they do not use 'perception' in the same way. For Spelke (1988), "human perceptual systems appears to analyse arrays of physical energy so as to bring knowledge of a continuous layout of surfaces in a state of continuous change." (1988, p.229), so that a perceptual process is a process that leads to the detection of no more than unbounded surfaces. Let's use this as our first meaning of perceptual. On the contrary, Quilty-Dunn (2019) says:

Segmenting the world into coherent, enduring units that can gain, lose, or change features while retaining their identity allows visual processing to make sense of retinal input. Without object representations, vision would be as William James imagined it to be for infants, a 'blooming, buzzing confusion' Quilty-Dunn (2019, p.22)

Following this quote, something is perceptual if it participates in *making sense* of sensory inputs.

I introduce the terms perceptual_1 and perceptual_2 with the following meanings:

1. perceptual_1 : what leads to or constitutes the detection of unbounded surfaces.
2. perceptual_2 : what participates in making sense of sensory inputs.

Caveat: I do not argue that these two characterizations of perception are respectively the exhaustive definitions of perception given by Spelke and Quilty-Dunn. They are illustration of the differences between their conceptions of perception.

Now we can paraphrase Quilty-Dunn and Spelke:

1. (PA Quilty-Dunn): "Object representations are perceptual" means "Object representations participate in making sense of sensory inputs".
2. (PA Spelke): "Object representations are not perceptual" means "Object representations do not lead to or constitute the detection of unbounded surfaces".

It is probable that both Quilty-Dunn and Spelke agree with each other about the paraphrases. However, they will continue to disagree on whether perceptual_1 or perceptual_2 are the best way to capture what perception is. In other words, they would disagree on the fact that the paraphrase of the other really means "Object representations are (not) perceptual".

To summarize, Quilty-Dunn and Spelke disagree on whether object representations are perceptual, they use different meaning of "perceptual" in the disagreement, and the disagreement persists after the clarification of the meanings. The dispute about object representations presents therefore all the signs for being a metalinguistic negotiation.

Hallucinations As I presented at the end of chapter 2, *Hallucinations* are cases of non-veridical perception, in which it seems to the subject that something in the world is perceived by her in absence of anything to be perceived. "Visual perception" traditionally encompasses veridical vision, illusions and hallucinations (Fish, 2010), in virtue of the common factor principle which says: "phenomenologically indiscriminable perceptions, hallucinations and illusions have an underlying mental states in common" (Fish, 2010, p.4). In this traditional view, "visual perception" refers to these common mental states. Disjunctive theories of hallucination (Johnston, 2004; Fish, 2009; Martin, 2004) however individuate the phenomenon of perception by incorporating the object perceived, what leaves hallucinations out of the class of perceptual states. Johnston says:

It does seem that once we adopt the act/object treatment of visual experience it is more natural to individuate an act of awareness occurring at a time in terms of an object that includes all that one is aware of in the relevant time (Johnston, 2004, p.171).

It is possible to consider that, at least in some sense, both theories agree on what are hallucinations, apart from their perceptual nature. They both agree on their existence and potential indiscriminability from veridical perception. However, what is different is what (visual) perception means for these theories, and this is in virtue of this disagreement that hallucinations are treated differently. They disagree on the concept of perception, on what phenomena fall under this concept. This is why we can consider that the disagreement about whether hallucinations are perceptual - when parties agree on what hallucinations are - is a metalinguistic negotiation.

4.3 A Normative Approach

I have shown that a part of the disagreements about the boundaries of perception is actually a metalinguistic negotiation about the concept PERCEPTION. How can I tackle this problem? In this last section of the chapter, I introduce the approach I will develop in the second part of the dissertation, namely *conceptual engineering*. In a recent talk given at NYU, Chalmers (2018) defines conceptual engineering as "the process of designing, implementing and evaluating concepts". This is what I will do in the second part of this work, I will elaborate a method to evaluate the different concepts of perception used in the literature.

Conceptual engineering is a normative approach for resolving the metalinguistic negotiation. I will evaluate the actual (and possible) concepts of PERCEPTION with respect to a norm, in order to find the best concept for research. In this last section, I discard two other competing approaches to resolve the metalinguistic negotiation: (i) the descriptive approach and (ii) the discovery approach.

The descriptive approach is based on tools philosophers have to investigate meanings of words and uses of concepts (linguistics, conceptual analysis and experimental philosophy). I will show that these tools cannot help me to resolve the metalinguistic negotiation at stake. Plunkett explains:

The facts about how [a speaker] is using her words can't themselves establish how she should be using them. Similarly, unless she has good antecedent reason to defer to the usage of other outside of the dispute, settling on the

facts about how others outside of the dispute use their words will also be of limited relevance to settling the normative question on the table (Plunkett, 2015, p.867)

The second alternative approach suggests to rely on the discovery of the true meaning of perception by scientific studies. According to this strategy, the concept PERCEPTION has a meaning fixed by the world, and should be discovered. I will discard this second approach because it is based on the assumption that perception is a natural kind – which is the assumption I refused to accept in chapter 3.

The descriptive approach Some could argue that we need descriptive tools to find the concept of perception we should use. By descriptive tools, I mean methods that have the goal to describe the actual meaning of words and concepts used by speakers. Psychologists and philosophers use such descriptive methods as experimental philosophy, conceptual analysis or descriptive linguistic.

Experimental philosophy is a field that tackles traditional philosophical inquiries by using systematic empirical studies, more often associated with scientific research such as psychology and neuroscience. Typically, experimental philosophy gathers experimental data and analyses them statistically. For example, in epistemology, attribution of knowledge has been studied by presenting vignettes (or cases) to participants, and by asking them whether they take it to be a case of knowledge or not (Myers-Schulz and Schwitzgebel, 2013; Colaço et al., 2014; Friedman and Turri, 2015). This could be seen as empirical research on what the ordinary concept KNOWLEDGE is. This kind of research can reveal striking results for philosophers. For example, Myers-Schulz and Schwitzgebel (2013) have shown that knowledge can be attributed without ascription of belief. At the question: "What year did Queen Elizabeth die?", a student – who has reviewed this information several times - can't recall it during a history test. When she guessed the answer, she wrote the correct one. When given this vignette, participants in the experimental study tend to ascribe to the student knowledge about the fact that Queen Elizabeth died in 1603, but they do not attribute to her the belief that Queen Elizabeth died in 1603 (Myers-Schulz and Schwitzgebel, 2013) see on (Knobe and Nichols, 2017).

Conceptual analysis is one of the traditional method in philosophy, and can be traced back to Plato's early dialogues. The idea is to analyze philosophical concepts, such as KNOWLEDGE or TRUTH, and find necessary and sufficient conditions for something to fall under the extension of this concept. This is a descriptive method as the goal is to find the conditions "hidden behind" the use of a concept. The method is often iterative: (1) use an example of knowledge (2) extract what the necessary and sufficient conditions are from this example (3) find counterexamples to refine the conditions. For example, I know that Paris is the capital of France. I guess that this is a case of knowledge because (a) I believe that Paris is the capital of France and (b) it is true that Paris is the capital of France. The hypothesis at this stage is that knowledge is true belief. But numerous counterexamples have been developed against this simple analysis, in particular when I believe something true by accident. The analysis has therefore to be refined in the light of counterexamples. Because of accidental cases, knowledge is better analysed as justified true beliefs, instead of mere true belief. But this new analysis also raises some counterexamples, such as Gettier cases. It is time now to refine the analysis again, etc. Once you cannot find counterexamples anymore, it means that you reached the conditions that generate the

right extension of the concept *as it is used in the cases you choose as data*.

Finally, descriptive linguistics aims to describe accurately real speech behavior in a specific language or community. This goal – and the methods associated – can be contrasted with prescriptive approaches that seek to give rules and grammatical structure to a language. Meaning of words are therefore studied by their use in everyday situations and environment.

Why not use these methods to discover what is the concept of perception? The main problem is that they would lead to the discovery and description of a *multitude* of meanings and concepts, without any resources to evaluate them.

I will present and develop the diversity of the concept PERCEPTION in the next chapter. However, the few examples I gave when I presented the disagreements on extension are enough for now to suggest that a descriptive approach will reveal more than one concept of perception.

Once the descriptive approach gave us a list of concepts, three solutions may be offered to deal with it: (1) be democratic, (2) accept a disjunction or (3) focus research relatively to a community in which the concept is stable.

The first solution would be to 'vote' for the concept of perception. We can imagine an experiment asking a significant panel of the population - or maybe just philosophers – what the concept of perception is. The conception with a majority of votes would then be elected as the concept of perception. We can think of a survey inspired by the Philpapers Survey, conducted in 2009 by Bourget and Chalmers (2014), in which more than 3000 philosophers (faculty members, PhDs and graduate students) presented their philosophical views about 30 philosophical issues. However, this solution seems inadequate here. The multitude of concepts will not disappear with a vote. Saying the contrary would be similar to say that the diversity of political ideas and parties disappears when a population elects a leader among several candidates. This solution is therefore not a good one. At best, it will inform us on the size of the different communities using different concepts.

The second solution is to accept a disjunction of all the different concepts of perception. Let's say that experimental philosophy discovers n different concepts of perception among philosophers. Then, let's call *super-concept* the concept constituted by the following disjunction : concept_1 or concept_2 or ... or concept_{n-1} or concept_n . This super-concept is unique and does not fall into the problem of a vote. However, two main problems arise: (1) it is not clear that a disjunction of concepts is still a concept and (2) this super-concept seems super useless. I quickly said that a concept is characterized by its intension and extension. If the notion of intension may accept a disjunctive intension, it is not clear that a disjunctive extension is still an extension. Remember that the extension of the concept is the list of things to which the concept applies. If one of the disjunct of the super-concept perception contains the mental states m in its extension, but another disjunct excludes it, is m part of the extension of the super-concept? It seems undetermined, and therefore problematic. Secondly, I remind you that the metalinguistic negotiation about the concept PERCEPTION arises in the broader debate of finding the borders of perception. This super-concept, because it has no definite extension, does not mark any boundaries at all. Therefore it is a useless concept for our actual inquiry.

Finally, the third solution is to use descriptive tools to define communities of concept, and narrow research to each community. Therefore, given a community of concept, there is a single concept of perception. The principle seems in itself a good idea. It is a way to dissolve the metalinguistic disagreement without denying the multitude of conceptions in general. As an upshot, in each community of concept, only object-level oriented questions will be debated, and representation-level oriented questions (such as our disagreements on extension) will be dissolved by descriptive tools. However, even if it is true that *within a conceptual community*, only object-level oriented issues will be debated, the metalinguistic negotiation will still exist *across communities*.

Recognizing that your conceptual community is not just constituted by yourself may be an argument against an idiosyncratic use of a term, but is not enough to resolve a metalinguistic negotiation.

As an upshot, it shows that descriptive tools are not enough, that a disagreement persists. Descriptive tools are useful for some work in representation-level oriented disagreements - finding the different concepts used in the literature or defining conceptual communities. However, they are not sufficient to resolve the following question: what concept of perception *should* we use? But now, the question is: what is lacking?

I suggest that what we need are *evaluative tools* – a way to evaluate the different uses of a term and the different concepts. Evaluating means using normative criteria, which is why the best approach to resolve a metalinguistic negotiation is a normative approach. The advantage of this kind of approach is also that evaluation is not restricted to actual uses of terms and concepts, but can also be applied to possible uses. It extends the analysis further, taking into account the possibility that the best concepts have not yet been used.

Evaluating concepts in order to find the best one presupposes that there are better concepts than others according to a norm. The normative approach is therefore constituted by two steps: (1) finding the norm of evaluation and (2) evaluating concepts with respect to that norm to find the best one(s). By differentiating these two steps, an objection to the normative approach arises: if the norm of evaluation is truth – namely, there is a true meaning of perception, a true concept of perception – there is no need for a normative approach anymore. What is the best concept to use? The only true concept. I still need to dismiss this solution to fully make room for my normative approach.

Contrary to the descriptive approach, in which finding the concept of perception is a matter of description of uses, finding the *true* concept is a matter of discovery. The presupposition here is that even if there is a multitude of uses, there is one single true concept of perception that we may not have discovered yet.

The discovery approach In this approach, there is one true conception of perception, and one true meaning of perception. This is because the word and concept PERCEPTION directly refers to its referent 'perception' in the world. This is by discovering what perception is that we will discover the concept of perception and the true meaning of the term. The meaning of perception has something to do with how the natural world is, not what people – speakers and hearers - mean by it (like in the descriptive strategy).

Contrary to the descriptive and normative strategies, the particularity of this approach is that the meaning of the term 'perception', and the concept, is *only* fixed by the world. This thesis is a strong form of *semantic externalism*. By explaining how to defend this thesis, I will explain why I discard this approach in the case of 'perception'. Roughly, this thesis requires to presuppose that perception is a natural kind as a starting point, before knowing its meaning. However, I suggested in the previous chapter that we should avoid this assumption.

Strong semantic externalism can be contrasted with theses such as "the meaning of 'perception' is fixed – at least partially – by its use in a specific community", or "the meaning of 'perception' is fixed by the description the speaker has in mind when the term is uttered". Semantic externalism should not be conflated with externalism about mental content. The former states that some words and concepts have a meaning determined by the world, when the latter affirms that some mental states like beliefs have content at least partially determined by external factors. For example, an externalist with regard to mental content (e.g. Burge, 1986) would say that the belief of a researcher whose content can be uttered by 'perception is a mental ability' is not only determined by the internal beliefs of this researcher. Instead, this belief can have a different content (and therefore, different accuracy conditions) depending on what 'perception' means in her scientific community. Hence, the content of beliefs are partially determined by factors external to the subject having the belief. This kind of externalism is not a problem here. On the contrary, strong semantic externalism say that no matter how 'perception' is used, the meaning of the word and concept is entirely determined by the natural world, which is unique, so that there is only one true meaning of perception.

Semantic externalism rests on the *semantic theory of rigid designation* and the *causal-historical view* of Kripke (1980) and Putnam (1975b) applied to natural kind terms. The argument is the following:

- P1. 'Perception' is a natural kind term.
- P2. Natural kind terms have a meaning only fixed by the world.
- C. The meaning of 'perception' is only fixed by the world.

I first explain the second premise by giving you some tools to understand the semantic theory of rigid designation. Premise 2 can be divided as such: natural kind terms are *rigid designators* and the meaning of rigid designators are fixed by worldly facts. In order to explain the theory of rigid designators, I need to introduce very quickly the notion of *possible worlds* (see a good introduction in Lycan (2008, p.46)). Our actual world is constituted by a number of entities and facts, that respectively exist and are true. For example, I exist in our actual world. You – the reader – also exists. The fact that you are actually reading this paragraph is true. However, things might have been otherwise. I may not have written this chapter, you may not be reading this right now. The fact that this chapter is not written is false in our actual world, but we can imagine an alternative world in which it is a true fact. This alternative world is a possible world. Possible worlds are therefore worlds that differ from our actual world in terms of true facts or existing entities, but are nonetheless possible – as they might have been true if something else would have happened. The upshot is that sentences have different truth-values from world to world. For example, "the writer of this paragraph is a woman" is true in our actual world, but in a possible world in which I copied this paragraph from a book written by my father,

this same sentence is false. Likewise, some singular terms designate different referents in different worlds. In our actual world, "the original writer of this paragraph" designates Géraldine Carranante, while in the possible world in which I copied my father's book, "the original writer of this paragraph" designates Pascal Carranante. In the world in which nobody ever writes this paragraph, "the original writer of this paragraph" designates no one at all.

Now, Kripke defines a rigid designator as a term that keeps the same referent across worlds. A rigid designator denotes the same thing in every possible worlds. Kripke (1980) claims that proper names, such as Géraldine Carranante, Kripke, and Pascal Carranante, are rigid designators. It means that whatever the world in which these names are uttered, they always refer to the same individuals, namely to me, to the philosopher Kripke and to Pascal Carranante. In a possible world, Pascal Carranante may have been different, may have done different things, may not have been my father, the name "Pascal Carranante" nonetheless refers to him, and so, independently of its properties that could have been different.

Kripke (1980) and Putnam (1975b) claim that natural kind terms are also semantically rigid. Natural kind terms refer to natural kinds, which means that natural kind terms – according to Kripke and Putnam – refer to the same natural kinds in every possible worlds (modulo their existence). This is the case because natural kinds are characterized in an essentialist way, as I did in chapter 2. I remind you that a natural kind is a non-arbitrary grouping of things that resemble each other in a fundamental and mind-independent way, such that some true facts about one member of this group are also non-accidentally true for the others (see chapter 2). This is explained by the existence of a common essential property – or 'essence', that makes the natural kind what it is. For example, what makes water water is arguably its chemical composition H_2O . This is its essential property. This essential property is the reason why naturalkind terms are rigid designators: they refer in every possible worlds to the things that have this essential property. Members of natural kinds have also superficial properties, that can vary across worlds. For example, in our actual world, water is odorless, but as it is a superficial property, it is possible that in another possible world, water has an odor. By contrast, there is no possible world in which water is not H_2O .

Superficial properties are often part of a descriptive stereotype, shared by a community that has not yet discovered the true metaphysical nature of a natural kind. This fact explains why different communities may have different meanings for the same natural kind, before finding its essential properties.

Now, why rigid designators do have a meaning only fixed by the world? This part of the second premise relies on the causal-historical view also developed by Kripke for proper names, and extended to natural kind terms.

The causal-historical view of meaning answers the following question: in virtue of what a proper name (or a natural kind term) refers to *this* or *that* referent? Kripke's answer (in the case of proper names) is that there is a causal-historical chain of utterances of a proper name that ultimately traces back to the naming of a newborn. Pascal Carranante refers to him in virtue of a chain of utterances of "Pascal Carranante" that begins the day his Mom and Dad named him this way. In the other direction, the story goes: one day a

child was named "Pascal Carranante", then people who were there this day began to use this name to refer to this child, then people who heard from them "Pascal Carranante" as the name of this child used it this way, etc. until I utter it now to refer to him. So my utterance of "Pascal Carranante" succeed in referring to him because it is connected historically and causally to his naming day. What is important here is that the success of reference relies here (only)² on a worldly fact, namely the chain of utterance, and not on any mental states, intentions of the speaker or other internal facts.

In the case of natural kind terms, there is no naming day, but a fact about how Nature is organized. We could imagine a hypothetical constitution day (or period) - to help the analogy - during which perceptual states and processes have been made what they are, namely a natural kind of things. When people use the ordinary language term 'perception' to talk about a paradigmatic exemplar of perception, they actually refer to whatever natural kind this paradigmatic exemplar belongs to, and so, no matter whether they know the true extension of the term. The extension of a natural kind term is not fixed by the psychological state of the speaker, but by the organisation of the world into natural kinds. Finding the extension of a natural kind term is therefore a matter of discovery, and consists in finding the essential property of this natural kind.

Putnam's general methodology is to imagine counterfactual situations in which an item satisfies a stereotype of a natural kind (superficial properties) but do not have the essential properties of the natural kind. This is his famous Twin Earth example. The case is the following: imagine in our world a planet identical to Earth (with a twin you, performing exactly the same behavior as you are). The only exception is that the substance that twin you calls "water", even if it looks exactly like earth water, is not constituted by H₂O, but by XYZ. Putnam argues that twin water is actually not really water, so that the meaning of "water" is actually not determined by the psychological states of the speaker (because twin you and you have the same psychological states when you utter "water"). On the contrary, meaning is determined by the external environment.

We could imagine an application of Putnam's twin earth to perception. The upshot is that it does not really matter what the psychological states of speakers are when they utter "perception". What matters is the extension of the natural kind term "perception", which is determined by the natural organisation of things.

Putnam's theory has been criticized, mostly on the basis that its case relies too heavily on the intuition that twin water is not really water. Dupré (1981) suggests that the existence of opposite intuitions would lead to a dead-end:

If Putnam says "XYZ is not water", and my intuition is that it would be (another kind of) water, how is such a dispute to be settled? Who knows what we ought to say in such fantastic situation? (Dupré, 1981, p.71)

To conclude, the second premise is problematic, as it is not clear that Kripke's theory of rigid designation applies to natural kind terms, as Putnam suggests. But more importantly, I reject this approach because I want to reject the first premise, at least as a

²Kripke adds that "when a name is 'passed from link to link', the receiver of the name must...intend when he learns it to use it with the name reference as the man from whom he heard it" (Kripke, 1980, p.96). This is required to avoid cases of "naming after", in which someone decides to call her cat "Elizabeth II" (see Lycan (2008, p.55) for more details). But these cases can be bypassed by arguing that they break the original causal-historical chain and begins a new chain the day the cat is named.

premise. This whole approach relies on the fact that perception is a natural kind and that 'perception' is a natural kind term and a natural kind concept. The discovery approach actually requires to follow the natural-kind strategy to resolve the metalinguistic negotiation about PERCEPTION. But, I try to solve the metalinguistic negotiation *because* the natural-kind strategy is a deadlock.

This is mainly why I reject the idea that we should rely on scientific discoveries to find the *true* meaning of perception. However, this detour through Kripke and Putnam's theories show the extent to which object-level oriented research and representation-level oriented one are entangled when it comes to scientific terms and concepts.

4.4 Conclusion

In this chapter, I claimed that a part of the disagreements produced by the boundary question are representation-level oriented. There is a metalinguistic negotiation about the concept PERCEPTION, and this metalinguistic negotiation is entangled with other object-level oriented debates about the nature and characterization of perception. I suggest to target this metalinguistic negotiation in order to clarify the whole debate about the boundaries of perception.

By discarding the descriptive and discovery approaches, I claimed that the best way to tackle the metalinguistic negotiation about PERCEPTION is a normative strategy, in which the concepts of perception used in the literature will be evaluated and potentially modified. This kind of strategy is called *conceptual engineering* and aims at finding the best concept of perception.

Another distinction between conceptual engineering, the descriptive approach and the discovery ones is their notion of concepts. For performing conceptual engineering, I will recruit a specific notion of (scientific) concept, in which a concept is a tool in research and scientific theories. As a tool, a concept can potentially be modified to best fit its purposes. This notion of concept is the main subject of the next chapter.

On the contrary, the descriptive approach considers concepts as mental representations and suggests to study them as such. It is a fact that people use different concepts of perception. For the descriptive approach, there is nothing more to say than establishing the list of existing concepts. Finally, for the discovery approach, words and concepts (at least proper names and natural kind terms) are transparent reference to the world. In a sense, pursuing the discovery approach means rejecting the existence of representation-level oriented problems. Every problem is indeed an object-level oriented problem, as the level of representation is a perfect transcription of the world.

Part II

Conceptual Engineering

Chapter 5

The concepts of perception

In the first part of the dissertation, I showed that the debate about the boundary question can be partly read as a metalinguistic negotiation about the concept PERCEPTION. In this second part, I build a method to evaluate and compare several versions of the same concept, and apply it to the concept PERCEPTION. This is a work of *conceptual engineering*.

This chapter is dedicated to narrow the range of possible concepts that can be legitimately compared. I argue first that PERCEPTION should be understood as a scientific concept for my purpose, as opposed to the ordinary concept, used in everyday discussion. Then, I disentangle the terms that should be understood as different versions of the same concept from different concepts tout court, thanks to Brigandt's framework (Brigandt, 2010, 2012). I argue that there are at least four concepts of perception in the literature, and that each of them includes several versions.

5.1 A *scientific* concept

5.1.1 Concept-as-theoretical-tool

The term 'concept' has been used in a lot of different ways in philosophy and in cognitive science. In philosophy of mind, and following Fodor's conception of concepts, concepts are mental representations that constitute – or at least allow us to entertain – propositional attitudes like thoughts. We are able to think about carrots, because we possess the concept CARROT and we use it to form the thoughts "I love carrots" or "Carrots are mostly orange". For cognitive scientists, in a very rough way, concepts are "the bodies of knowledge that are used by default in the cognitive processes underlying our higher cognitive competences" (McCaffrey and Machery, 2012). It means that concepts are whatever allows us to recognize things, to sort things into types, to make inference about things, etc. However, in what follows, I am not interested in concepts as mental representations. I consider concepts as theoretical tools.

Concepts as theoretical tools are constituted by a body of knowledge, an ensemble of properties *used in theories*. I am therefore not interested in the format that these theoretical tools may have in the head of the scientist or the philosopher (whether they are prototypes, exemplars or theories, see Machery, 2005). A concept as theoretical tool can even not have any corresponding mental representation, or if it has, may be quite different than the one mentally available. It may contain more information, or even less, than concepts

as mental representation. Everything that is associated to the concept within a theory is potentially a part of the concept.

For example, the meaning of the concept of ELECTRON as a theoretical tool in theories of physics is constituted by the up-to-date knowledge about it. However, it is possible that nobody in the world possess a mental representation of this conception of electron, either because it is too complex, or because mental representations do not have a format allowing the inclusion of all aspects of a concept as theoretical tool.

It may be objected to my work that the kind of object I am targeting in my investigation is not really a kind of concept. However, this kind of objects are also what natural language calls concept. Furthermore, it shares several properties with concepts as mental representations:

1. It is public. Publicity is the property for a concept to be sharable. A concept-as-theoretical-tool is sharable as it is supposed to be part of a theory that can be discussed, modified, objected by different members of a scientific community.
2. It is productive. The productivity of concepts is the fact that a concept can generate an infinite number of propositions. It is also the case of a concept-as-theoretical-tool, as they are meant to build scientific hypotheses, to participate in scientific induction, etc.
3. It can be characterized with the classical duo of intension and extension.

In this dissertation, I focus on concepts of perception as theoretical tools, used in scientific and philosophical theories. Whether this theoretical tool is used as such in the mind is not in the scope of this work. Yet, I guess that these theoretical tools are internalized in a relatively close way, but this hypothesis remains to be proven, as well as the processes of internalization of complex scientific concepts. Anyway, this work on concepts is neutral on the mental representations used in cognitive processes. This is why some classical debates on concepts are not relevant here – for example on the nature of concepts in general.¹

5.1.2 Perception: A *scientific* concept

When I say that PERCEPTION is a *scientific* concept, I have in mind concepts like ELECTRON, NATURAL SELECTION, MAMMAL, etc. that are used in the *scientific* understanding of Nature. These are concepts-as-theoretical-tools used in scientific theories.

A scientific concept is thus a concept-as-theoretical-tool used in proper scientific theories. In order to give a full-fledged characterization of it, I should enter the debate about what is good science vs non-science vs pseudoscience, which is far beyond the scope of this work. When it comes to the concept PERCEPTION, some research communities working on

¹It may be possible to understand concepts-as-mental-representations and concepts-as-theoretical-tool as part of the same continuum. This could be a way to interpret Nersessian when she says: "talking of concepts provides a means of bridging the psychological and the community phenomena. Concepts are basic way through which human represents the world. They categorize experiences and take note of relationships, differences, and interconnections among them. For the individual, concepts figure in a range of cognitive phenomena including memory, inference, problem solving, language comprehension, and belief systems, to name a few. Scientific concepts provide systematic representations through which individuals and communities understand, explain, and make predictions about phenomena." (Nersessian, 2008, p.186)

it are internationally acknowledged as doing good science. I consider that this is enough to validate a scientific status to the concept PERCEPTION.

One may also mention that scientific research often has a classificatory aspect, so that conceptual apparatus developed by scientists aim to categorize their domain of investigation into kinds. The investigation about the boundaries of perception and the distinction between perception and cognition has this classificatory aspect, as one of the goal of this research could be understood as classifying the states and mechanisms of the mind into kinds.

I will work on this dissertation with the scientific concept(s) of perception, which are those used as theoretical building blocks for broader scientific investigations about the mind. These are the concepts that I will evaluate (in chapter 6) in order to find the best one for theorizing about the mind.²

5.1.3 Objection: what about the ordinary concept of perception?

Some people could object that the relevant concept of perception to study is the *ordinary* one (also called folk concept), namely the one used in everyday life and understood by the layman.

I do not argue here that PERCEPTION is *merely* a scientific concept. PERCEPTION is actually also an ordinary concept. What I argue is that the scientific concept of perception is the relevant one for our inquiry here – namely resolving the metalinguistic negotiation among philosophers and scientists.

Two reasons can be called for asking me to work with the ordinary concept of perception. First, there is a risk of conceptual discrepancy between scientific problems and solutions if I do not use the ordinary concept. Secondly, ordinary and scientific discourses are – or at least should be – connected through corresponding conceptual apparatus. I develop these two lines of reasoning separately in what follows.

Conceptual discrepancy between problems and solutions The objection runs as follow:

- (P1.) The concepts used in asking a question should be the same used in the answer, for it to be an answer to that question.
- (P2.) The ordinary concept of perception sets the terms of our question here, namely the boundaries of perception.
- (C.) Therefore, we should use the ordinary concept of perception in order to answer the boundary question.

²Some people may raise the objection that there is a philosophical concept of perception, different from the one used by scientists (e.g. what should be perception for it to justify knowledge?). However, the literature I am working with is a scientifically informed philosophy which uses scientific results as pieces of evidence in philosophical theories. I assume that the philosophers I study here share a *scientific* notion of perception.

The first premise is quite strong. We all have lived these very frustrated discussions in which we ask questions and our conversation partner answers off track because she misunderstood one term of the question. It reminds me of a french nursery rhyme from Bobby Lapointe about a mom of little fishes, and how nice she is with her offsprings. The song finishes with 'And me, I like her with lemon'. I used to sing this song all the time when I was a child, I found it very funny. This song plays on the broad meaning of the verb "to like". If I ask my partner: "Do you like my mom?" And he answers: "Yes, of course, I tasted a finger last time, I will probably add a little more salt and pepper for the next serving though". I would probably laugh and ask the question again: "Honestly, what do you think of her?". The reason is that he did not answer the question I asked, but another question that could have been asked with the same terms.

When it comes to scientific discourse, Strawson warns us against this practice of giving answers about non-scientific questions with scientific concepts:

[T]o offer formal explanations of key terms of scientific theories to one who seeks philosophical illumination of essential concepts of non-scientific discourse, is to do something utterly irrelevant – is a sheer misunderstanding, like offering a textbook on physiology to someone who says (with a sigh) that he wished he understood the workings of the human heart... (Strawson, 1963, p.505)³

This worry is also found in literature about race, gender, and even knowledge. For example, Haslanger writes:

In asking what race is, or what gender is, our initial questions are expressed in everyday vocabularies of race and gender, so how can we meaningfully answer these questions without owing obedience to the everyday concepts? (Haslanger, 2000, P.34)

Ludlow also says concerning knowledge:

[F]irst, and most obviously, any investigation into the nature of knowledge which did not conform to some significant degree with the semantics of the term 'knows' would simply be missing the point. (Ludlow, 2005, p.13)

Even if the principle of the first premise is quite vague (how is individuated a question among the multitude of questions that may be asked with the same terms), I grant it to the objectors and focus instead on the second premise.

I reject the second premise, i.e. I reject that the boundary question is expressed with the ordinary concept of perception. When I talk with my family about my dissertation subject, and the question I am tackling, I have to explain to them some psychological facts about perception that they may not know but are indispensable to understand the question. For example, I have to take some time to tell them that vision in the human brain does not work as a camera, and that non-linear processes of information analysis appear very early in the brain processes after the retinal registration of light. With these explanations, I can show them why the classical model in which perception consists in a passive capture of an image of the world is put under pressure, and that discussing the boundaries of perception is challenging. The fact that I have to take some time to make

³Quotes in this section are mostly coming from Cappellen's book (Cappellen, 2018).

sure everybody understands the question is a good sign that this question is not to be understood in everyday terms.

It is true that PERCEPTION has not been forged within science, contrary to concepts such as NATURAL SELECTION, or PHLOGISTON. But this does not rule out the possibility that a scientific version of the concept exists and is autonomous. The difference between PERCEPTION and PHLOGISTON or NATURAL SELECTION is that the latter are (arguably) not associated with some ordinary concept. The existence of an ordinary concept of perception is no objection against the fact that some problems or questions are restricted to the scientific concept of perception.

It also may be true that the origin of the questions (in a genealogical sense) has to be found in the ordinary concept. But it is absurd to maintain that original questions never change in the course of scientific or philosophical inquiries. It is a part of the investigation to refine questions, to narrow problems. Part of the job of researchers is to work on questions – not only on explanations. Therefore, even if at some ideal point in the past, questions about perception in general have been framed with an ordinary conception of perception, it is not the case anymore – at least not in my work.

Ordinary and scientific conceptual schemes are connected - if not identical

The idea of this second objection is roughly that within a linguistic community, everybody shares the same concepts, even if only a part of the population (say, scientists and philosophers) knows exactly what it means, and how to determine its extension. The other speakers trust the experts of their community for using the right meaning of that term. The consequence is that there is only one and single concept (there is no distinction between an ordinary and a scientific concept), that is understood more or less exhaustively by different subsets of speakers of a linguistic community. This idea is best illustrated by Putnam's notion of division of linguistic labor. Putnam sets the division of linguistic labor hypothesis as the following:

[A linguistic community] possesses at least some terms whose associated "criteria" are known only to a subset of the speakers who acquire the terms, and whose use by the other speakers depends upon a structured cooperation between them and the speakers in the relevant subset (Putnam, 1975a, p.146)

Putnam takes the example of the concept ALUMINIUM. Most of us know some things about aluminium, such that it is a metal, it is a common material used for pans, and it may be dangerous for health. However, most of us (including Putnam and me) are unable to distinguish a pan made in aluminium from a pan made in molybdenum. Does it mean that we don't have the same concept as the expert in metal who are able to do so? Putnam maintains that it is not the case, we all share the same concept, in virtue of being part of the same linguistic community. What happens is that we rely on experts for knowing the criteria of individuation of some concepts such as aluminium and molybdenum. However, it does not prevent us from using it in circumstances in which it is not important to know how to recognize instances of aluminium. An intuition seems to confirm Putnam's hypothesis: if we gets aluminium and molybdenum mixed up when it comes to identify the material of a pan, we will gladly accept that we were wrong. This would not be the case if the concept we were using was different from the chemical concept of aluminium (like an ordinary concept of aluminium whose criteria are superficial properties such as

color, aspect, etc.).⁴

I suggest two tests to know whether a term is subject to Putnam's division of linguistic labor:

1. In some occasions, a non-expert of the term can ask an expert some help to identify instances of it.
2. A non-expert can be wrong in uttering a sentence with the concept, if the concept is not properly used according to the criteria of experts.

Applied to a concept like GOLD, these two tests work well. First, when we want to know whether a jewelry is made with gold, we go to see a jeweler who is able to identify gold. Secondly, if I affirm that my bracelet is only made with gold but I learn later that it is actually only gold-plated, I would accept that I was wrong. On the contrary, not all concepts meet these tests. According to Putnam, a concept like CHAIR is not subject to the division of linguistic labor. It seems indeed weird to go see an "expert in chairs" in order to identify some chairs. Similarly, if I say that something is a chair and that one tells me that it is not, I would probably resist the comment by saying something like: "it is totally legitimate to call this thing a chair".

Now, the objection against the idea that perception can be studied as a scientific concept is the following:

1. When the division of linguistic labor hypothesis applies, the ordinary and scientific versions of the same concept have the same meaning.
2. The division of linguistic labor hypothesis applies on the concept of perception.
3. Therefore, the ordinary and scientific concepts of perception have the same meaning.

I argue here that the second premise is incorrect, because the division of linguistic labor hypothesis does not clearly apply to the concept of perception. I think that there is an ordinary concept of perception that resists the hypothesis, the same way the concept CHAIR does. The difficulty lies in the fact that there is one concept of perception, namely the scientific one, subject to the division of linguistic labor. This is the source of confusion.

Most of the time, nobody goes to ask an expert in perception whether their mental states are really perceptual or not. Secondly, most of the time, when I utter perceptual terms such as 'to see', it would be very weird if someone tells me I'm wrong in using the term. There are exceptions though, in medical conditions for example in which it is not clear whether some states are perceptual or hallucinatory. In this latter situation, it is the scientific concept of perception that is used.

Let's take Strawson's example with the concept HEART I quoted previously. It is indeed very unlikely that my baker, who has been left by her boyfriend very recently, goes to see

⁴Putnam's hypothesis is based on its externalist account of meaning, in which the meaning of a term or concept is not "in the head of the speaker", but is determined at the level of a sociolinguistic community. As he says: "Whenever a term is subject to the division of linguistic labor, the "average" speaker who acquires it does not acquire anything that fixes its extension. In particular, his individual psychological state certainly does not fix its extension; it is only the sociolinguistic state of the collective linguistic body to which the speaker belongs that fixes the extension." (Putnam, 1975a, p.146)

her doctor to know whether or not her heart is broken. On the other hand, if she goes, complaining about a pain in the chest, she could ask "Doctor, do you think it may be the heart?". If the doctor answers: "yes, you know, sometimes human relations are painful", she would probably not be satisfied. But if he answers: "no, it is probably not the heart because your pain is on your right side and the heart lies more on the left side of your body", she would be more satisfied.

I think HEART and PERCEPTION are similar notions. They both have an ordinary meaning (or ordinary meanings) and a scientific one. The important point here is that the scientific meaning is subject to the division of scientific labor, such that non-expert people may use the scientific concept of heart or perception without knowing much about it. This is the case of my baker when she goes to see her doctor for a pain in the chest, or if someone goes to check his auditory system to know whether she has hallucinations.

The consequence is that there are still two different concepts of perception – one ordinary and one scientific. What I accept from Putnam's hypothesis is that the scientific concept of perception may be used in some degraded states by non-scientific members of the same linguistic community. Could it be source of confusion? I don't think so, the context will most of the time provides clues about which concept is used (e.g. a medical visit).

A last point is important to note: in Putnam's division of linguistic labor, the concept is best understood by experts, and not by laymen. Therefore, in order to know what a concept means, you should go and see the experts. Therefore, if I want to evaluate the best versions of the concept PERCEPTION, I have to work on the expert versions of the concept, not on the ordinary use of the scientific concept of perception.

To conclude, once it is acknowledged that there are both an ordinary and a scientific concepts of perception, and that the division of linguistic labor applies on the latter, Putnam's work supports the idea that I have to focus on PERCEPTION as understood by experts, such as philosophers of mind and psychologists.

5.1.4 Some consequences

Performing conceptual evaluation or conceptual engineering on PERCEPTION as a scientific concept is performing what Cappelen calls a *local* conceptual engineering (vs a *global* one). It means that I target a "specific term for a particular purpose in a particular context" (Cappelen, 2018, p.36). This is different from aiming to change the meaning of a term for the whole language. By restricting my work to the scientific concept(s) of perception, I leave aside the ordinary one.

This restriction is adequate for the metalinguistic problem we want to tackle here, namely a metalinguistic negotiation *among experts*. However, it may give rise to some worries.

First, I am not taking into account some links between scientific and non-scientific world, such as popularization of scientific discoveries.

A second worry is that concepts are not fully independent in terms of meaning. In a strong version, the meaning of a concept is dependent on the meaning of all the other existing concepts, in a kind of quinean holism. In a weaker version, the meaning of a concept is dependent on the meaning of some concepts in its "conceptual neighborhood", so that in

order to fully evaluate a concept, one should evaluate all the neighbor concepts. This should be defended, but I accept that this may be a limit of my work.

5.2 Individuation of concepts

In the previous chapter, I argued that researchers working on perception use different concepts of perception, so that they should enter into a metalinguistic negotiation. Here, I aim to organize the field of perception studies in order to capture the different concepts used by them. This organization applies to the actual concepts of perception, but could also be used to classify new concepts.

The challenge of this organization is to differentiate different versions of the same concept from different concepts tout court. It is not because the same term is used that the concept underneath is necessarily the same.

This distinction is crucial for the next chapter (chapter 6), in which I elaborate a model for comparing different versions of the same concept. I need to be sure that I am comparing what is comparable, namely versions of the same concept. It is possible to ask which conception of APPLE is better between APPLE₁ (e.g. green and round fruit) and APPLE₂ (e.g. the fruit of apple tree), but it is absurd to try to evaluate which concept is better between APPLE and ORANGE.

I use Brigandt's methodological framework to study scientific concepts (Brigandt, 2010, 2012). The particularity of this framework is the introduction of a new constituent of conceptual meaning beside classical intension and extension: *epistemic goals*. This new element enables me to distinguish families of concepts, so that the concepts PERCEPTION encountered in the literature can be organized into four main concepts. These four main concepts are themselves subject to some conceptual variation.

In this section, I begin by presenting Brigandt's framework. I then explain how this framework help us distinguish different versions of a concept from different concepts tout court. In the next section, I will apply the model to the literature about PERCEPTION.

5.2.1 Brigandt's framework

In Brigandt's framework, a scientific concept is characterized by three components: (1) the concept's reference, (2) the concept's inferential role and (3) the epistemic goal pursued by the concept's use.

The reference of the concept is the ensemble of objects, substance or events in reality a concept refers to. I consider this component as synonymous to the classical notion of extension, namely all the objects in reality that fall under the concept. For example, the reference or extension of DOG is the ensemble of dogs.

The inferential role of a concept is similar to the notion of intension. It includes its characterization but also the inferences in which the concept is used in science, in a way that it constitutes its meaning and determine its reference.

Reference and inferential role are classical components ascribed to concepts in order to characterize them. The originality of Brigandt's framework is the introduction of the third component, namely *the epistemic goal of a concept*.

A concept's epistemic goal is the aim scientists pursue when they use this concept the way they use it. Brigandt presents the example of the concept *HOMOLOGY* in evolutionary biology, used by biologists to establish phylogenetic trees. The epistemic goal of the concept *HOMOLOGY* (or at least one of its epistemic goals) is the classification of species. Another example is the concept of *NATURAL SELECTION*, used to account for evolutionary adaptation. The epistemic goal of this concept is therefore something like: "explaining evolutionary adaptation".

It is important to understand the difference between the inferential role and the epistemic goal of a concept, as both may be understood as the "function" of the concept. The inferences and explanations constitutive of an inferential role are the ensemble of actual beliefs and knowledge about the concept, while "the concept's epistemic goal is the kinds of inferences and explanations that the concept is *intended* to support" (Brigandt, 2010, p.24, my emphasis). Whether the concepts *HOMOLOGY* or *NATURAL SELECTION* actually fulfill their epistemic goal is not the question here. Ultimately, the inferential role of a concept should answer the questions raised by its epistemic goal, but this would be the case in mature scientific investigations.

The epistemic goals of concepts can be seen as mere genealogical traits, as scientific concepts are forged in problem-solving processes (Nersessian, 2008). In this perspective, the epistemic goal of a concept would explain its formation but would not be informative of its meaning. However, I am not talking here about an alleged *initial* epistemic goal, but the current epistemic goal of a concept, the one currently pursued by researchers. Following Brigandt, I consider that the meaning of scientific concept is reflected in its scientific use (Brigandt, 2010), that is why the epistemic goal of a concept is an aspect of its meaning.

A concept may have several epistemic goals, as it may be used in different scientific researches. Furthermore, the epistemic goal(s) of a concept has to be understood at the level of a scientific community. I do not deny that some researcher may have idiosyncratic ways of using concepts, and especially have very unusual research goals. But they are not our target here, and their use do not contribute to the meaning of the concept (as a scientific concept).

Another point to note here is that Brigandt's framework is methodological, and not ontological. These three components – reference, inferential role and epistemic goals – are aspects of a concept, used in order to study them and their meaning. Like me, Brigandt does not draw here a picture about the nature of concepts and their ontological constituents⁵.

The particularity of Brigandt's framework, i.e. the introduction of the epistemic goal(s) of a concept, allows me to determine whether different versions of a concepts are actually versions of the *same* concept. This is what I called the challenge of sameness of topic.

⁵"My tenet that a concept consists of three components (reference, inferential role, epistemic goal) is not so much to be understood as a metaphysical doctrine about what a concept is, rather it is a methodological guideline about how actual scientific concepts are to be studied." (Brigandt, 2012, p.5)

5.2.2 The challenge of sameness of topic

The metalinguistic negotiation I identified in philosophy of mind shows that philosophers (and psychologists) have several concepts of perception, but how can I be sure that they are all about the same topic⁶? These philosophers may have different things in mind when they use the term 'perception'. If I want to evaluate concepts of perception, I have to be sure that they are all versions of the *same* concept, compared to different concepts of something close-but-different. This is the challenge of sameness of topic.

This problem of sameness of topic is a classic one in philosophy of science, mostly tackled in the larger debate of theory change. This challenge has been brought to light mainly by Kuhn (1962) and Feyerabend (1975) as the incommensurability challenge. Kuhn holds the view that some theoretical changes are 'revolutionary' when they reassess theoretical assumptions of a field (what Kuhn calls a *paradigm*). A scientific paradigm is a particular scientific world-view, which influences observations, explanations and theories about a field, so that two paradigms have no points of potential comparison. They do not share the same conceptual language. This is why paradigms are said to be incommensurable. As a consequence, concepts – and more generally theoretical tools – of a paradigm are incomparable with concepts of another paradigm, even if they share the same term. They are simply not about the same thing, not about the same world.

Kuhn says:

Since new paradigms are born from old ones, they ordinarily incorporate much of the vocabulary and apparatus, both conceptual and manipulative, that the traditional paradigm had previously employed. But they seldom employ these borrowed elements in quite the traditional way. Within the new paradigm, old terms, concepts, and experiments fall into new relationships one with the other. The inevitable result is what we must call, though the term is not quite right, a misunderstanding between the two competing schools. (Kuhn, 1962, vol II, n2, p.149)

For example, the concepts of space before and after Einstein's general theory of relativity are not comparable according to Kuhn. Before Einstein, space is "necessarily flat, homogeneous, isotropic, and unaffected by the presence of matter." (ibid.) Without this characterization of space, Newtonian physics could have not be developed. The Einsteinian scientific revolution replaces this old concept with a new way of conceptualizing space (along with new concepts of time, matter and force). The Newtonian and Einsteinian concepts of space may share the same term, but they are not comparable, as they stand in completely different scientific world-views.

Coming back to our problem of sameness of topic, it can be said that the Newtonian and Einsteinian concepts of space are actually not about the same thing. They are not two different versions of the same concept SPACE, but two different concepts tout court.

In the case of perception, I acknowledge that (most of) the concepts of perception in contemporary philosophy of mind are actually concepts of PERCEPTION, broadly speaking. They (arguably) belong to the same paradigm, and their authors share probably a close world-view about the mind. Therefore, they do not fall under the suspicion of incommensurability. However, I think that there is a way in which some researchers on perception

⁶I use the term topic to refer vaguely to a tackled subject. This vagueness is on purpose as it is precisely the challenge of sameness of topic to define what a topic is to individuate concepts.

are not using the same concept, while others are using the same but in different versions. The upshot is that some researchers are not talking past each other, while others do.

Conceptual variation across the same research community is not rare, so that the concept of perception is not an exception. Brigandt points the same pattern of semantic variation of the concept GENE in contemporary molecular genetics, in which different characterizations lead to different intension and extension of the concept. To cite an example of variation among many, what is called a gene could be a DNA element in one theory, but can be the product of a DNA element in another. The problem arises because there is a many-many relation between a DNA elements and products, so that the classification of genes in the two theories are not isomorphic (Brigandt, 2012).

I use Brigandt's epistemic goal to make this distinction between two versions of the same concept versus two different concepts of close-but-different topics. I use epistemic goals as the criterion for individuating concepts – so that theoretical terms that have the same epistemic goals but differ in terms of reference and/or inferential roles, are different versions of the same concept. In other words, I equates the notion of 'topic' with the problems tackled by a specific epistemic goal.

Using epistemic goals is important for me because most of the concepts of perception used by researchers do not share the same reference or the same inferential role. If I had to be limited to these two aspects of meaning to organize concepts, I would have ended up with all instances of concepts as tackling different topics. The upshot is that I would not have been able to perform any evaluation. Conceptual engineering would have been impossible.

My solution also allows me to explain why communication between researchers is not broken across variation. They are understanding each other because they are trying to solve the same problems, with broadly the same assumptions. Their sharing of the same epistemic goals explains also why they tackle the same debates, the same questions. It also explains why some researchers are really not understanding each other, when they are actually not aiming the same epistemic goal.

Here is a summary of how I individuate concepts and versions of concepts:

- Two different concepts are concepts that have different reference, inferential role and epistemic goals.
- Two theoretical tools that share the same epistemic goals are different versions of the same concept.

My solution is close to Cappelen's use of *coarseness* of topic to secure sameness of topic. He argues that a topic is more coarse-grained than extensions and intensions, so that "sameness of topic doesn't track sameness of extension" (Cappelen, 2018, p.109). He uses literature about samesaying, arguing that in a lot of cases (gradable adjective, context-sensitive expression, etc.), we grant that two people are saying the same thing, even if what they are saying are both sensitive to change of context, and uttered in different contexts. In a sense, my solution is a specific case of Cappelen's one as mine is only applicable to scientific concepts. The meaning of a scientific concept, in Brigandt's framework, is more complex than its intension and extension, as it must include epistemic goals pursued by the concept's use. As I chose to identify sameness of topic (or subject of study) with the family of concepts that share the same epistemic goals, my topics are coarser than

intensions and extensions. The advantage of my solution is that I have a clear criterion to disentangle different versions of the same concept, while coarseness of topic is too vague. The disadvantage of my view is that it can only be applied to scientific concepts that can be studied with Brigandt's framework. This drawback is acceptable for me in this work, as I consider the concept of perception as a scientific concept in this context of inquiry. On the contrary, Cappelen's solution is meant to be applied to all sorts of concepts and topics.

More traditionally, sameness of topic across conceptual variation is secured with sameness of reference (Devitt, 1979, for a review). In this perspective, two theoretical tools are two versions of the same concept when they refer to the same things in the world, even if they are used differently in different theories. This solution, however, is not applicable in the case of perception, as I showed indeed that the different concepts of perception used in the metalinguistic negotiation have different intensions (or inferential roles) that pick out different extensions (or references). Therefore, the way to characterize perception actually changes the reference of the concept.

Richard (2019) holds that the intension and extension of a concept are essential to it, so that changing one or the other necessarily leads to a switch of concept. My answer is that my characterization of concept is not ontological, but methodological. In my framework, the three components of concepts (reference, inferential role and epistemic goals) are not essential parts of the concept. They are ways of studying it. In this perspective, Richard's extension and intension are not the same thing as what I call reference and inferential role.

Another worry, raised by Haslanger, is that by changing the extension and intension of a concept, we "are in danger of providing answers to questions that weren't being asked" (Haslanger, 2000, p.34). As an example, if we are changing the intension and extension of the concept WOMAN, and ask the question "why women are paid less than men?", we may have the illusion to have answered the question with the new concept, when the question was actually asked with the ordinary concept of woman (example found in Cappelen, 2018, p. 101). As a consequence, we just do not have answered the same question. I suggest that by focusing on epistemic goals as criterion of sameness of topic, I secure a continuity of inquiry across different versions of the concept. When concepts are essential to answer some questions, so that change of extension endangers their identities, these questions are part (or direct consequence) of the epistemic goals of these concepts.

A last comment is that epistemic goal may change while reference remains the same (when a concept is used in a new investigation for example). I do not think this case arises for perception, but it may be the case in the future. The consequence is that the sameness of epistemic goal is sufficient but not necessary for sameness of topic. Sameness of topic may be said to be preserved if one of the three components of Brigandt's framework is unchanged. Using sameness of epistemic goal to secure sameness of topic is therefore a particular case of this more general rule. This particular case is enough for my purpose though, as I restrain my work to versions of the concept PERCEPTION that share the same epistemic goals.

5.3 The multitude of scientific concepts of perception

5.3.1 The epistemic goals of PERCEPTION

I identified (at least) four different epistemic goals for the scientific concept of perception. Before presenting them, I want to specify that my investigation is not exhaustive, as my research is limited to analytic philosophy of mind and psychology. The four following epistemic goals are the ones I encountered most often in my readings and discussion.

The four epistemic goals are the following:

1. The phenomenal goal: describing the phenomenal character of perceptual experience
2. The knowledge goal: explaining belief/knowledge acquisition
3. The computational goal: solving the under-determination problem, namely explaining how the perceptual system performs the determination task from proximal stimuli to the representation of the world.
4. The biological goal: explaining how an organism survives and adapts to its environment

The first two epistemic goals (phenomenal and knowledge) are mostly coming from philosophy, while the two others (computational and biological) are mostly coming from cognitive psychology. I begin with short descriptions of the four epistemic goals, and then I add some remarks about their overlapping.

The phenomenal goal The phenomenal goal of the concept PERCEPTION is to describe and explain the phenomenal character of perceptual experiences, or in other words, the "conscious aspects" of experiences (Fish, 2010, p.17). In Nagel's famous terms (Nagel, 1974), there is "something it is like" for an individual to have a specific perceptual experience. There is "something it is like for me" to see a red light, and "something it is like for me" to hear a piano. Researchers that are trying to capture this specificity of perceptual consciousness follow what I call the phenomenal goal in studying perception.

This epistemic goal is more generally nested in studies about consciousness, in which researchers focus on experiences (i.e. conscious experiences) in general. It is actually a whole program of description, research of conditions of possibility, function, laws, etc. This broader epistemic goal is well expressed – in its descriptive side – by Nagel:

At present we are completely unequipped to think about the subjective character of experience without relying on the imagination-without taking up the point of view of the experiential subject. This should be regarded as a challenge to form new concepts and devise a new method - an objective phenomenology not dependent on empathy or the imagination. Though presumably it would not capture everything, its goal would be to describe, at least in part, the subjective character of experiences in a form comprehensible to beings incapable of having those experiences (Nagel, 1974, p.449).

With these methods and concepts, Nagel hopes that we might "develop concepts that could be used to explain to a person blind from birth what it was like to see" (ibid.). In this general research program of consciousness studies, perception is thought as a particular way of being conscious. As Moore says: "the fact that I am seeing something now is

obviously related to the fact that I am conscious now in a peculiar manner” (Moore, 1925, pp.46-47). Moore is understanding ”perceiving” as a determinate state of the determinable being conscious”. It is no surprise that philosophical theories within this epistemic goal focus on the items a perceiver is aware of when she is perceiving, as well as their intrinsic and relational properties. This is the case of sense-data theorists (a perceiver is aware of sense-data and their properties) and naive realists (a perceiver is aware of external objects and properties).⁷

In the phenomenal goal, what is typically called unconscious perception does not fall within the extension of the concept PERCEPTION, and unconscious phenomena (blindsight, suppression effects, etc.) are indeed missing in most, if not all, debates. This lack has been given as a reason to give up relationalist theories such as naïve realism (Block, 2010). But Phillips (2018) suggests that this reason is flawed:

There is something puzzling about this complaint. For, insofar as unconscious perception is incoherent on a relationalist approach, only a theorist who had already rejected that approach to perception could possibly interpret a given empirical case in such terms. But then it is unclear whether the issue is, after all, empirical or, instead, a disagreement concerning which conception of perception one ought to adopt in the first place. This issue concerns not just relationalist views, but any experientialist view, that is any view on which perception is treated as constitutively connected with conscious experience (Phillips, 2018, pp. 4-5).

In other words, and following Phillips suggestion, it is not that theories such as sense-data or naïve realists are ill-equipped to give an account of unconscious perception, it is merely outside of their research domain. For the followers of the phenomenal goal, perception is studied as conscious experience. The fact that they do not have the theoretical resources to accommodate unconscious perception does not endanger their theories, as they never had the ambition to study something else than perception-as-conscious-experiences.

A similar story goes in a recent article of Kriegel (2019). He is looking for an introspectable difference between perceptual experiences and thoughts. Therefore, the difference has to be conscious – not in virtue of an ontological necessity, but because it is the domain of Kriegel’s research in this article. He says:

The discussion to follow ignores entirely the phenomena of unconscious perception and thought. I will not consider much phenomena as blindsight or even subliminal perception, not tacit or dispositional belief. The reason for this is that I do not expect there to be any first-personally manifest difference between such mental states (Kriegel, 2019, p.7)

In this article, Kriegel also focuses on perception in the realm of consciousness. Should we say that his research is inadequate because it is silent on unconscious perception?

Some may object that the phenomenal goal cannot be the epistemic goal of a *scientific* concept of perception. But this difficulty may be explained by the fact that phenomenal character of experiences is necessarily linked to the individuals who experience them. On the contrary, classical scientific methods try to cut loose from the subjectivity of a specific point of view. Nagel reminds us this point:

⁷See Fish, 2010 for a review of literature on sense-data theories and naïve realism.

If the subjective character of experience is fully comprehensible only from one point of view, then any shift to greater objectivity – that is, less attachment to a specific view point – does not take us nearer to the real nature of the phenomenon: it takes us farther away from it. (Nagel, 1974, p.445)

Today, scientific research on consciousness is thriving, proving that scientists found scientific ways to study it. Furthermore, my way of understanding scientific concept does not rule out in principle any method of investigation. A scientific concept is required to be used in theories within a research community. Refusing that the phenomenal concept of perception is scientific is refusing that the whole community of researchers working on consciousness is not really doing legitimate scientific research.

There is variability among concepts of this goal. One important variability relates to the accuracy of experience. For disjunctivists (including naive realists), PERCEPTION only includes veridical perception while sense-data theorists include also illusions and hallucinations – even if both theories acknowledge that phenomenal character of veridical and non-veridical experiences is indistinguishable by introspection. I do not have the ambition to deeply analyze the consequences of this difference. However, I suspect that this variability exists (and is maintained) because veridicality of perceptual experience – in the sense that it leads to the acquisition of veridical beliefs – is not an important property of experience when it comes to the phenomenal goal (even for naive realist – as the rejection of illusory and hallucinatory experiences is a consequence of their theory, not a requirement). On the contrary – and this is my transition to the next section – the acquisition of (veridical) belief is the main point of the epistemological goal of perception.

the knowledge goal The knowledge goal of perception is the use of the concept of perception within the general goal of explaining human knowledge. Perception is indeed thought as the primary source of knowledge. The knowledge goal of the concept perception is twofold. First, it focuses on the role of perception in belief acquisition, and secondly on the explanation of how perception may lead to knowledge (understood as justified true belief). The two horns are related but independent in principle. Indeed, knowledge may arguably be acquired without belief (e.g. Russell's Knowledge by acquaintance). I nonetheless put them together under the notion of knowledge goal as the goal of explaining belief acquisition from perception is often tackled as a foundational step for further investigations about knowledge (in foundationalist theories, but not only). Put it another way, the knowledge goal is constituted by all the questions and problems arising from the idea that perceptual experience may be considered as a "tribunal" for beliefs - in Quine's terms (Quine, 1951). Problems about knowledge settle questions such as how experience can "return a particular kind of verdict", namely a good one, towards beliefs (McDowell, 1996, p. xiii). Problems about belief acquisitions are broader as they question "how experience can return any verdict on our thinking at all" (ibid.)

McDowell's position about the content of perception – namely the conceptuality of perceptual content – is mainly justified by the idea that perceptual content must stand in justificatory relations with beliefs. McDowell's goal is thus to explain the contribution of perception in justifying belief. The knowledge goal is also followed by Dretske as he writes:

My description concerns itself with those features of perception which have served, and continue to serve, as a touchstone for epistemological controversy.

In a word, the topic is what we see and how we see in so far as the answers to these questions help us to get clear about what we know and how we know it (Dretske, 1969, pp.1-2).

Similarly, Fodor affirms that "the point of perception is the fixation of belief" (Fodor, 1983). I suspect that Fodor studied perception in the research perspective of its contribution to his "Language of Thoughts".

Brewer is another follower of the knowledge goal. Here are the first introductory words of one of his books:

This book is about the role of conscious perceptual experiences in the acquisition of empirical knowledge. I take it for granted that a person's beliefs about the way things are in the world around him causally depend to a large extent upon the course of his perceptual experiences. My central concern here is with the epistemological dimension of this relation, its provision of a peculiarly fundamental source of knowledge: perceptual knowledge. How is such knowledge, that particular mind-independent things are objectively thus and so, even possible? What is the nature of the conscious experiences upon which it is based? How should we conceive of the epistemic contribution which such experiences make to it? These are the questions which drive my enquiry. (Brewer, 2002, p. xiii)

From the general form of the knowledge goal is derived a range of epistemological questions in which the concept of perception is recruited. For example, Peacocke builds a theory of perception in order to solve questions about concept acquisition and possession (Peacocke, 1992). One may also name the famous debate in philosophy of science on whether scientific observation is theory-neutral. On this question, one may quote Dretske again: "One of my chief interests is the relation of observation to scientific practice" (Dretske, 1969, p.3).

The alleged epistemological role of perception is the main resistance towards the cognitive penetrability of perception, as there is a sense in which if perception is penetrable, then it is not possible to consider it as a robust foundation of knowledge anymore.

When it comes to the consequences of this epistemic goal on the concept PERCEPTION, researchers following the knowledge goal must accept that the extension of PERCEPTION includes illusion and possibly hallucination, as the problem of perception as source of knowledge is partly constituted by the apparently contradictory facts that (i) perception is the main source of knowledge and (ii) perception may be illusory.

There is also some variability in the intensions and extensions associated with PERCEPTION within the knowledge goal. This variability is mainly due to differences in theories of perception considering what in perception explains its relation to belief and knowledge. For example, in a dretsian theory of perception in which the non-conceptuality and analog format of perceptual content explains why perception is source of both knowledge and illusion, all states with conceptualized content are consequently post-perceptual. However, in a conceptualist picture of perception like McDowell's one, in which the justificatory role of perception is based on its conceptuality, all perceptual states have conceptual content. In a sense, if some states prove to have non-conceptual content, McDowell would probably reject them as genuinely perceptual.

the computational goal The computational goal consists in finding what operations are performed within the mind for it to represent the world through sensory organs. This epistemic goal is mostly tackled by psychologists and neuroscientists (vs philosophers) as the processes studied are said to be unconscious and require sophisticated methods and tools to be characterized. The computational goal won his popularity thanks to the rapid rise of vision science in this last century.

This epistemic goal is best understood as the aim of resolving what is called the under-determination problem. The under-determination problem in perception (or inverse problem) stands for the discrepancy that exists between the proximal stimulus (the retinal image for vision) and the representation of the distal stimulus in the mind, such that it is impossible to determine with certainty the latter from the former only. In other words, the information registered by the senses (the proximal stimulus) under-determines its source (the distal stimulus), so that in order to form a representation of the world, it is necessary to transform it in some ways.

Phillips describes the under-determination problem as follow:

For any given proximal stimulus, there is a multitude of environmental scenes that could have caused it. For example, a convex object illuminated from above, and a concave object illuminated from below, can produce the same pattern of light on the retina. (Phillips, 2019, p.7)

Here is another example:

[T]he spectrum of light wavelengths reflected from an object's surface into the observer's eye is a product of two unknown spectra: the surface's color spectrum and the spectrum of the light illuminating the scene. Solving the problem of "color constancy" - inferring the object's color given only the light reflected from it, under any conditions of illumination - is akin to solving the equation $y = a \times b$ for a given y , without knowing b . No deductive or certain inference is possible. (Griffiths et al., 2008)

Perception of depth, volume and distance also requires an extraction and transformation of information from the retinal image.

Contrary to the phenomenal and knowledge goals, the computational goal may seem to be an *internal* problem to perception (when phenomenal and knowledge goals target respectively broader questions about consciousness and epistemology). It is not the case, though. The computational goal aims at explaining the computations performed by mental systems, in which the notion of perception is recruited to play a certain role. Within this epistemic goal, perception refers to the states and processes that produce a representation of the world through our senses. The mistake to think that the computational goal is different from the others comes from the fact that the under-determination problem is often conceived as a problem faced *by the visual system*. It produces the idea that the resolution of the under-determination problem is actually the *function* of perception, understood as its inferential role. But the under-determination problem is faced by the researcher (not the visual system !). This is an epistemic goal. *The visual system is not resolving a problem, it performs a task* - which is a task of determination. The computational goal of perception is to resolve the under-determination problem, which consists in explaining the determination task performed by perceptual systems.

That is what Maar does in his theory of vision. He describes processes building successive sketches from retinal information (raw information) to fully understandable image. He famously argues that vision consists in building what he calls a 2 ½ sketch, in which depth is represented, but not volumetric shapes and distance. These latter features are a matter of interpretation and are post-perceptual (Maar, 1982). He therefore recruits the concept of perception (more specifically vision) for explaining the broader problem of how retinal information is transformed into objects and properties of the world.

Historically, the computational goal has been at the origin of a profusion of psychological theories of vision during the twentieth century. To only name a few :

- the Cybernetic theory of McCulloch and Pitts (1948) tried to explain – among others – ”how the experience of form is preserved despite differences of position of projection of objects on the retina” (Bartley, 1969, p.466)
- Brunswick’s probabilistic-functional theory (1947) describes the reconstruction of the external object by perceptual processes as a probabilistic approximation of it.
- Gestalt theories express ”laws of Gestalten” that are supposed to describe the rules underlying percept formation.

Some of the ”laws of Gestalten” are taken by Maar (1982) to build his theory of vision. For example, the construction of the full primal sketch makes use of combination principles such as local proximity and similarity (spatially close and similar elements are combined). More recently, and in a philosophical tradition, Pylyshyn (1999), Burge (2010) and Raftopoulos (2009) have, among many others, tried to characterize visual processes in their ”operational constraints”.

The computations leading to the formation of the percepts of objects are determined by the input to the visual system (that is, the optical array), by the physiological mechanisms involved in vision, and by the computations they allow and certain principles that restrict and guide the computation. These principles are constraints that the system must satisfy in processing the input. These constraints are needed because perception is under-determined by any particular retinal image ; the same retinal image could lead to distinct perceptions. (Raftopoulos, 2009, pp.103-104)

Within the computational goal, perception is understood as the set of processes that perform the determination task from proximal stimulus to the representation of the world. In the specialized literature of the perception/cognition divide, Phillips’ definition of perception is a good example of the computational goal:

A process is perceptual just in case it has the function of producing representations of environment entities by being causally controlled by those proximal stimuli that these entities produce. (Phillips, 2019, p.7)

Caveat: the computational goal may seem to be quite close to the knowledge goal, as they both aim to explain how perception represents the world in some ways. However, the knowledge goal aims to explain how perception represents the world in a way that is assessable through accuracy conditions, while the computational goal aims to resolve

an equation. Given the available information (proximal stimulus), what are the operational steps performed in the mind to reach a state of representation of the world. In other words, the question is: how does perception come to say something about the world?

Conceptual variability within this epistemic goal comes mostly from the processes postulated by different theories in characterizing the determination task of perceptual systems.

One type of processes commonly considered as contributing to the determination task is constancy capacities. They are said to disentangle variations in the proximal stimulation caused by properties of a distal object from variations due to other features of the environment. For example, color constancy capacity allows us to see a red chair as uniformly red even if illumination actually causes a lot of variation in the information detected by the retina. There is a huge literature on constancy capacities, especially in vision. Constancy effects have been shown for color, size, shape, distance, or motion.

Burge (2010) is inclined to think that constancy capacities fully explain the determination task of perception ("objectification" in his own term, 2010, p. 339). However, other capacities may be thought to participate in the determination task. I think of amodal and modal completion or phonemic restoration (what Burge calls the "extraction of form", (2010, p. 418)). One may also think of multi-sensory integration like multi-sensory binding, in which different modality specific proximal stimulus perceived as being caused by the same source (Treisman, 1998; Revonsuo, 1999; O'Callaghan, 2019). In multisensory integration, information from a modality-specific proximal stimulation (for example audition) can be used to disambiguate information from another modality-specific proximal stimulation (for example vision).

Finding which processes are used in the determination task is ultimately an empirical matter. It is no surprise that followers of this epistemic goal favour empirical methods. In this line, Burge says:

One cannot distinguish cognition from perception in any warranted way from the armchair. The processes for forming attributives on the basis of perception are too fast, inaccessible to consciousness, and complex to allow phenomenological or other armchair methods to distinguish perception from cognition. Only sophisticated use of experimental evidence bears on these issues in a way that goes beyond uninformed playing (Burge, 2014a, p.583)

However, variation in the concept does not come only from operational solutions in the problem resolution. It also comes from disagreements about the nature of these processes (whether they have to be understood at the level of the individual or at the level of a computational system) or from disagreements about which sensory states can actually represent the world (smell is famously rejected by Burge as non-perceptual, 2010).

For example, Burge rejects what he calls sensation as part of perception, on the basis that it lacks the "objectivity" of perception (understood as the ability to represent the world). On the contrary, all the processes from the retinal image are probably genuinely considered as perceptual by Raftopoulos, for which burgian sensations would probably be an early stage of perceptual processing.

The biological goal The biological goal of perception is its contribution to survival and adaptation of an organism to its surrounds. I called it biological goal because this line of research often focus on an alleged biological function of perception, and its relation to the whole body of an organism.

Here is an example in the literature:

[T]he biological function of perception is to give the organism information about the current state of its own body and its physical environment, information that will assist the organism in the conduct of life. (Armstrong, 2002, p.209)

The biological goal understands perception as a provider of *useful* information about the world, mostly in order to guide actions of individuals. On the contrary, the knowledge goal focuses on its possibility to provide *accurate* information, while the computational goal is attached to explain how perceptual information is about the world simpliciter.

One may say that only accurate information is useful for an organism, so that the knowledge and biological goals eventually merge. But other sources of usefulness for perceptual information may be considered such as the speed of processing. This is one of Fodor's consideration in defending modularity of perceptual processes: "Roughly, we want the perceptual identification of panthers to be very fast and to err, if at all, only on the side of false positives." (Fodor, 1983, p.70). We may add "in order to run as soon as possible and survive".

Akins (1996) offers several examples of perceptual systems that do not aim at producing accurate descriptions of the world, but instead at warning the organism about any relevant or dangerous event happening to it. This is the case of thermoreception - the perceptual mechanism underlying our perception of external temperature. Human thermoreception does not work like a thermometer, i.e. producing a signal correlated with external temperature that preserve the structure of properties (this temperature is greater than that one). On the contrary, our perception of temperature is quite complex (composed of four kinds of receptors, unevenly distributed over the body surface) and varies depending on the area of the body, as well as temperature change at the surface of the skin. The upshot is that the representation of temperature through our thermoreceptors cannot be said to be *accurate*. It is however quite useful and corresponds to human's needs for survival.

The system as a whole constitutes one solution to man's various thermal needs - that he be warned when thermal damage is occurring or before it is likely to occur, when temperature changes are likely to have specific consequences, and so on. (Akins, 1996, p.350)

The biological goal is a very broad investigation, and I do not have time here to explore all of its derived forms⁸. What is interesting is that the biological goal may lead to concept of perception quite different from the other goals. For example, Forgas and Melamed write:

We have decided to place the process of perception within the context of man's general need to adapt to his environment so that he can cope effectively with the demand of life (Forgas and Melamed, 1976, p.1)

They are clearly following the biological goal. They continue with:

For our purposes perception will be defined as the process of information extraction (ibid.). [W]e [...] conceive of perception as the superset, with learning, memory and thinking as subsets subsumed under the perceptual process. (Forgus and Melamed, 1976, p.1-3)

We see here how the biological goal may completely inverse the relation between perception and thoughts compared to the knowledge goal. Within the knowledge goal, perception is understood as a required step that lead to thoughts, and especially true beliefs. On the contrary, the biological goal may lead (as in Forgus and Melamed's theory) to an understanding of the ability to think as a subpart of a larger perceptual ability understood as an ability of extracting information.

5.3.2 Independence and synergies

Epistemic goals are not always explicit in the work of a researcher, so that it is not always clear which one she is following. Furthermore, as it might have strike you during the reading of this section, most of the goals effectively pursued by researchers are often a combination of several epistemic goals. For example, proponents of architectural strategies (see chapter 2) can be said to follow both the computational and knowledge goals. The modularity of perception is indeed a theoretical step towards the resolution of the under-determination problem, but also a way to secure a foundation for knowledge, free from belief influences.

In this last section, I explain why the four epistemic goals I presented before are *in principle* independent. It means that the can be followed without any mention of the others. This is true that, *in practice*, they are often pursued jointly. I will show that this joint research may be fruitful in some conditions, thanks to the existence of synergies between the epistemic goals. I will however warn against a widespread practice of considering the epistemic goals of the concept as the roles of perception in general.

Independence By following more than one epistemic goal, one makes the theoretical assumption that the same mechanism/phenomenon plays a role in several scientific investigations (e.g. the same mechanism/phenomenon is supposed to both explain the phenomenal characteristics of a perceptual experience and resolve the under-determination problem). This is not something that should be obvious. When researchers are asking questions about consciousness, knowledge or representations, they use PERCEPTION as a building block in their theory, but nothing secure the fact that they are referring to the same thing. They are actually shaping the concept PERCEPTION to their own theoretical needs. As researchers are not asking the same questions, and do not have the same theoretical needs, they actually create different concepts of perception that are not necessarily referring to the same mechanisms and/or phenomena. I suspect that this idea is usually overpowered by the assumption that perception is a natural kind, which would secure the

⁸I wonder whether Gibson's ecological approach and enactive theories of perception (Noë) have to be understood within the biological goal. It is true that they both give a fundamental role to bodily movement, action and environment in their theory of perception, but I think their research goal was primarily to explain how perception comes to represent worldly properties that are not present in the retinal image. In other words, I think that the primary epistemic goal of ecological and enactive approaches of perception is the computational goal. Gibsonian answer is that all the information is included in the optical array, when the dynamic of movement is taken into account. A typical enactive answer would be that part of perceptual content is built through past and present actions of the organism with the perceived environment.

fact that concepts of perception used in answering different questions are actually referring to the same entity or class of entities in the world. However, as I shown in the second chapter, the idea that perception is a natural kind is not well-established, and it is more cautious to refuse it.

Once perception is not considered as a natural kind anymore, I am free to accept that the mechanism recruited to answer an epistemic goal is different from the mechanism that answers another one. What explains the phenomenal character of a perceptual experience may be different from what resolves the under-determination problem, and most importantly, this is absolutely not a problem.

This independence of epistemic goals of perception may shed a new light on some historical disagreements and objections. For example, sense-data theorists have been said to throw a "veil" on reality – in Bennett's famous expression. The objection relies on the idea that sense-data theories do not explain at all how perception may lead to knowledge about the world. Sense-data are indeed compatible with a skepticism a la Berkeley towards the external world. The objection is something like: (1) Perception is (one of) the foundation of knowledge, (2) Sense-data theories do not have the theoretical resource to explain this fact, (3) Therefore, sense-data theories must be false. With the idea that sense-data theories have been developed within the phenomenal goal of perception, we now see why sense-data theories may not have the resources to explain the knowledge goal of perception. Sense-data theories are not fit for the epistemological agenda of perception. But there is a difference between "saying nothing about x", and "saying something false about x". As Dancy says:

[Sense-data theory] does not have the consequence that external objects are unobservable; it purports simply to tell us something about what it is to observe them. (Dancy, 1985, p.65)

On the contrary, belief-acquisition theorists developed theories of perception within the knowledge goal of perception, and it happens that they do not have the resources to meet the phenomenal goal. As Fish says:

I haven't yet said anything about what the phenomenal character or presentational character of an experience is according to the belief acquisition theory. This is, I confess, because I'm not sure that it applies. The belief acquisition theorists identifies visual experiences with episodes of acquiring beliefs, and I suspect that this theorist would deny that the episode of acquiring beliefs has a phenomenal character. (Fish, 2010, p.56)

Of course, some philosophers tried to pursue both the phenomenal and knowledge goals at the same time. One can cite the sensory core theory, the percept theory, Pitcher's theory of perception, etc. (for an overview, see Fish, 2010)

According to my view, it is quite unfair to dismiss either sense-data theories or belief-acquisition theories on the basis that they do not have the theoretical resources to meet an epistemic goal they never tried to follow. They may well be both quite powerful in their own domain.

There is one important consequence of the independence of the four epistemic goals: the four corresponding concepts are actually not comparable – one cannot say that a phenomenal concept of perception is better or worse than a knowledge concept of perception. However, two concepts elaborated within the same epistemic goal (what I call two versions of the same concept), are evaluable and comparable. This consequence will be crucial for the next chapter, in which I develop a method to evaluate different versions of a concept.

Synergies Epistemic goals of perception are independent in the sense that they can be followed independently. They are however related, so that it is possible to follow several of them at the same time. But the theoretical transfer between epistemic goals must be done with some care.

There is indeed a very strong resistance to the idea that the epistemic goals are really not related such that there is no possible communication between them. First of all, even if I leave open that perception is a natural kind, I do not want to deny that researchers working on perception within different epistemic goals are nonetheless talking about very similar things, corresponding to a capacity of the mind vaguely defined as the ability to bring information about the external world inside us.

Furthermore, I think that there are real synergies between epistemic goals. They are coming from the fact that all theories within different epistemic goals must be consistent with the physical implementation of mental capacities. Therefore, discoveries within an epistemic goal become constraints within another one. The idea is that all theories should be eventually compatible, at some level of understanding. This is why one may find some synergies across all epistemic goals.

What is the difference between accepting synergies and accepting dependence between epistemic goals? The point is that there is no direct theoretical transfer between epistemic goals: a conceptual adaptation is required, as conceptual schemes are different across epistemic goals.

The notion of an indirect theoretical transfer seems unproblematic when a new conceptual scheme is explicitly offered, with different terms. For example, Shea (2014) uses new conceptual tools (top-down and bottom-up effects) in order to tackle some of the questions typically raised in philosophy of perception. This new conceptual scheme does not map properly the typical distinctions between perception and cognition (and that's actually the point of it), but has nonetheless something to say (and something new) about these (old) questions. In Shea's words:

The distinction between top-down and bottom-up effects can be put to philosophical work. It allows us to ask, for any psychological mechanism, how much its operation is driven by current input and to what extent it is influenced by prior information. Answering that question will form an important part of spelling out how sensory information is processed. It also provides a basis for asking epistemological questions about the relations between sensory systems and the world, relations amongst sensory systems, and their relations to other psychological capacities. These issues also arise in the philosophical literature on the cognitive penetrability of perception by cognition. The top-down vs. bottom-up distinction allows us to pose some of those questions more generally,

not just about paradigmatically perceptual states. So the distinction would allow some of the concerns in the cognitive penetrability debate to be preserved, even if a distinctive category of the perceptual turned out to be theoretically unsustainable. (Shea, 2014, p.2)

A direct theoretical transfer consists in taking the discoveries from one epistemic goal and putting them directly at work within another epistemic goal. It assumes that the concepts used by all theoretical parts are the same, and that the domain of investigation is identical. In an indirect theoretical transfer, one has to be careful about discrepancies of conceptual scheme and domain of application. I found this special care in Raftopoulos' work, when he draws some consequences of his view (within the computational goal) on epistemological questions (knowledge goal). He points that algorithmic work on how vision produces outputs does not directly apply to belief justification. He says:

In perception, the epistemic support or justification for perceptual beliefs accords with externalism and not internalism. A viewer does not have to state the reasons on account of which they are justified in holding a perceptual belief. Instead, one's visual experience as if X is before them is a reason for believing that there is an X before them only because in one's world such a visual experience is reliably related to an X being before them. (Raftopoulos, 2015, p.168)

Computational discoveries on perception are useful for the epistemologist, as they may explain why a perceptual experience is reliably related to the existence of some object in the world. However, the overlapping of the two theories is not total: the theory that meet the computational goal does not directly explain epistemological problems of perception. Some theoretical work is necessary for putting at work the computational discoveries.

Here are some other example of synergies between epistemic goal. This list is not exhaustive, and only cross two goals when it is also possible to imagine crossing 3 or 4 goals.

Phenomenal and Knowledge: How things look may give some ideas on the form of justification people are given to their beliefs. Imagine a dialog between two friends: "I believe that Steve is in Paris - Why ? - Because I saw him there this morning - Are you sure ? He told me he was in Tokyo. Can you describe me what you saw ?"). However, it is not clear that:

1. the process of belief acquisition is directly explained by a kind of phenomenal endorsement (I believe how things look),
2. what explains how things look also explains how people justify their beliefs, and
3. the best way to describe the phenomenology of vision (or another sense) is what will help us understanding how people comes to believe and justify their beliefs based on the phenomenal character of their perceptual experiences.

Phenomenal and Computational: Phenomenology of perception may inform the computational theories on some of the outputs (at least the conscious ones) of perception. Phenomenal work on perception can provide thus some of the explananda to the computational community (but not all). The other way around, computational work may constrain and guide some phenomenal work. For example, the phenomenology of color perception can based some of its thesis and guides its research /experimental design on

knowledge of how colors are actually coded in the retina and the brain.

Knowledge and Biological: As I previously pointed, the accuracy of content in belief (studied by the epistemologists) is important for a biological point of view, as it is better to perceive accurately the world in order to survive. However, other features of perception, such as high speed, may be crucial to survival and may reduce in some sense the need for accuracy (see Fodor's panther example).

These examples are very common ways to relate different epistemic goals together. But I must emphasize that these synergies are not direct and require a change of perspective. In contrast, let's consider an intuitive and largely spread alternative understanding of the four epistemic goals of perception. This alternative understanding is well illustrated in the introduction of Schellenberg's book. To better understand the mistake, I need to make you read a large part of it for you to follow the narrative:

Perception is our key to the world. It plays at least three different roles in our lives. It justifies beliefs and provides us with knowledge of our environment. It brings about conscious mental states. It converts informational input, such as light and sound waves, into representations of invariant features in our environment. Corresponding to these three roles, there are at least three fundamental questions that have motivated the study of perception: Epistemology question: How does perception justify beliefs and yield knowledge of our environment? Mind question: How does perception bring about conscious mental states? Information question: How does a perceptual system accomplish the feat of converting varying informational input into mental representations of invariant features in our environment?

[...] Theories motivated by addressing the mind and information questions have been developed largely independently of concerns about how perception furnishes knowledge of our environment and how it justifies our beliefs. Similarly, theories motivated by addressing the epistemology question have been developed largely independently of concerns about how perception brings about conscious mental states. To be sure, most accounts of perceptual justification rely heavily on the idea that perception justifies beliefs in virtue of its phenomenal character. However, such accounts typically take it as given that perception provides evidence and immediately proceed to addressing the question of what the relationship is between such evidence and relevant beliefs. This split between philosophy of mind and cognitive science on the one side and epistemology on the other has hindered our understanding of perception. Questions in philosophy of mind are intimately connected with questions in epistemology in particular with regard to perception: the role of perception in yielding conscious mental states is not independent of its role in justifying our beliefs and yielding knowledge. If this is right, then perceptual experience should be studied in an integrated manner. (Schellenberg, 2018, p.17)

Schellenberg's "three fundamental questions that have motivated the study of perception" (Epistemology, Mind and Information question) are very similar to three of my epistemic goals (knowledge, phenomenal and computational). But she says that these three questions are corresponding to three roles perception is playing in our lives, which motivate –

according to her – an integrated study of perception.

In a more schematic picture, Schellenberg's reasoning is the following:

1. Resolving a question about perception (Epistemology, Mind, Information, etc.) is explaining one of the roles of perception.
2. A good theory of perception is a theory that explains all these roles (with the same theoretical apparatus).
3. Therefore, in order to develop a good theory of perception, one must study the questions about perception altogether.

According to my view, there is a deep misunderstanding of the nature of questions about perception in this picture. The mistake is to think that questions about perception (what correspond more or less to my epistemic goals) reveal "roles of perception in our lives" (in Schellenberg's vocabulary). These roles are therefore considered as explanandum, and questions widespread in the literature are symptoms, signs of these roles. I suggest that the picture should be inverted. Questions are fundamental – not questions about perception specifically, but bigger questions about consciousness, knowledge and representation of the world. Within these questions, the concept of perception is recruited to play a theoretical role. These theoretical roles are the roles corresponding to the questions Schellenberg cites. However, these roles are not the roles of perception strictly speaking, these are roles of the concept of perception within a theory of knowledge, of consciousness, etc.

The difference of pictures may be reformulated this way: In Schellenberg's picture, there is one thing – perception (probably a natural kind), that possesses some roles in our lives in virtue of its nature. Therefore, the best way to understand perception is to give an integrated theory that explain all of these roles. If a theory leaves out one of these roles, it means that it does not fully capture the nature of perception.

In my view, researchers engage in different questioning and postulate the existence of a mechanism or perceptual phenomena that participates in answering questions about consciousness, knowledge, or computational processes of the mind.

What I suggest is that Schellenberg confuses the epistemic goals and the inferential roles of perception. Brigandt warns against this confusion. He points the fact that both "inferential role and epistemic goal are aspects of a concept's use", however, he says:

[U]se has usually been identified with how a term is used (inferential role), though what a term is used for (epistemic goal) is likewise to be taken into account. Most importantly, labels such as 'concept use', 'function of a concept' and 'conceptual role' could be seen as ambiguously referring to both inferential role and epistemic goal, even though the two must be clearly distinguished. (Brigandt, 2012, p.5)

The confusion is therefore quite easy. To the question "what is perception?", one may answer without much objections: "what justify belief, yield knowledge of our environment, bring about conscious mental state and convert informational input into mental representation of our environment, altogether". But one is actually pointing out to what perception *is used for*.

It is possible to pursue several epistemic goals at the same time but either it requires an additional theoretical assumption (that perception is a natural kind), or it should be done carefully as the concepts do not necessarily refer to the same class of entities. Theoreticians have to build bridges across conceptual apparatus.

5.4 Conclusion

In this fifth chapter, I have offered a structured description of the concepts of perception in cognitive science. The concepts can be organised in (at least) four independent epistemic goals and display variations inside each epistemic goals. Following Brigandt's framework, I consider them as four different concepts of perception with several versions for each of them.

In terms of comparability, only versions of the same concept can be genuinely compared. They are indeed competing for the best way to *conceptually grasp their object* (while the other concepts are merely referring to another object).

The second reason why it is more interesting to compare two versions of the same concept is that the notion of epistemic goal can be used as a powerful and fundamental ground for building an evaluative apparatus. Within the same epistemic goal, a concept is indeed better than another if its extension and intension allow to better meet (or at least follow) the epistemic goal. For example, within the knowledge goal of perception, concepts that will participate better in understanding mechanisms of belief and knowledge acquisition will be preferred to concepts that are less efficient in this task. Put it another way, an epistemic goal provides what Brigandt's call *epistemic standards*, in the light of which a concept can be evaluated – or two concepts can be compared.

[T]he epistemic goal pursued by a concept's use provides the required standard: a change in the concepts inferential role (definition) is rationally warranted if the new inferential role meets the concept's epistemic goal to a higher degree than the previous inferential role. (Brigandt, 2012, pp.15-16)

Brigandt uses the notion of epistemic standard in order to explain how scientific conceptual change can be rational. In the next chapter, I borrow the notion to explain how we may rationally prefer a version of a concept over another.

Chapter 6

Evaluation, Pluralism, Fragmentation

The goal of this chapter is to resolve the metalinguistic negotiation about PERCEPTION. Put it another way, it aims at determining which concept should be used in scientific and philosophical research, among the variety of different ones identified in the literature.

It is not possible to evaluate which concept is best between the phenomenal, knowledge, computational and biological concepts of perception, because they are not about the same topic. It is possible that some of them eventually refer to the same mental mechanisms, but this is something that need to be defended with empirical and theoretical evidence. To emphasize the difference between these concepts, I call them p-PERCEPTION, k-PERCEPTION, c-PERCEPTION and b-PERCEPTION, respectively for phenomenal, knowledge, computational and biological (I add x-PERCEPTION to talk about them in an indefinite way).

Each x-PERCEPTION nests several versions. These versions share the same epistemic goal, and that is why they can be evaluated with respect to it: the best concept is the one that meet best their shared epistemic goal. This is the evaluative method I elaborate in this chapter. This method is a technical tool of conceptual engineering that provides a way to arbitrate between versions of the same concept. For the sake of clarity, I will first present the theoretical model for the evaluation. It is mostly based on Egré and O'Madagain's work on concept utility (Egré and O'Madagain, 2019). Then, I will apply it to the concept k-PERCEPTION¹. I let the evaluation of other x-PERCEPTIONS for a future investigation. The application of the theoretical model requires a lot of decision on the dimensions of evaluation, that are specific to each evaluated family of concepts. That is why the application on k-PERCEPTION is absolutely not applicable as it is to other x-PERCEPTIONS.

This evaluation gives us which version is best between the different versions of k-PERCEPTION. I will however nuance the result and show that other criteria could enter the evaluation and potentially change the verdict. Furthermore, the evaluation is relative to the chosen dimensions of evaluation, as well as the evaluated items, so that a change in dimensions or new empirical discoveries can change the evaluation.

After performing the evaluation, we have an answer to the metalinguistic negotiation: it is possible to provide which concept is best to use. However, the answer depends on the

¹I chose the knowledge concept of perception (k-PERCEPTION) as a personal preference. I suspect that my adhesion to this specific concept is influenced by my interest for its epistemic goal.

epistemic goal of the researcher. We are left with a *pluralism* of concepts, that are all legitimately the best concepts of perception. I argue that we should accept this form of conceptual pluralism, as a decent form that avoid the usual pitfalls that one may reproach it.

Conceptual pluralism is actually frequent in science, and I will show that the pluralism of PERCEPTION is the result of a process of conceptual fragmentation that regularly happens to scientific concepts. Conceptual fragmentation is a scientific phenomenon in which a concept thought as having a single meaning, turns out to be used with several different meanings for different purposes. The upshot is that the conceptual pluralism of PERCEPTION is not to be feared as a bad consequence of my conceptual engineering, but instead is the sign that scientific research about perception is currently progressing.

This chapter is organized in two main parts: (1) Evaluation, and (2) Pluralism and Fragmentation. I begin with the evaluation of the versions of k-PERCEPTION (section 6.1). In this section, I present the theoretical model, the application to k-PERCEPTION and the limits of it. In the second section, I explain why we should accept conceptual pluralism for PERCEPTION as a consequence of my conceptual engineering (section 6.2). I will show that this pluralism is harmless, useful and even expected as a part of a classical process of conceptual fragmentation happening to scientific concepts in general.

6.1 Evaluation

Evaluating versions of a concept means that some of them are better than others. However, intuitively, a concept is a mere representative tool and cannot be better or worse than another one. Here, what I call a good (or virtuous) concept is a concept that promotes good thinking and behavior. By contrast, a bad (or defective) concept participates in bad thinking and behavior. For example, I consider that, all things being equal, a concept of WOMAN that promotes social justice is more virtuous than another one that justifies social inequalities.

This way of evaluating and comparing concepts is based on how concepts are used, so that I do not evaluate *concepts per se*, but their use. However, in the case of scientific concepts and according to Brigandt's framework, I considered that their use - at least their epistemic use - has to be understood as an aspect of their meaning. Therefore, evaluating how a scientific concept is epistemically used provides a judgment on the concept (and not *merely* on its use).

The epistemic virtue of scientific concepts are their good contribution to science, in the sense that better scientific concepts are those that better contribute to improve scientific knowledge. There are several ways to contribute positively to science. Some concepts may improve knowledge in specific areas of science, some may produce more coherence between different scientific areas, some may help designing good empirical designs, etc.

For the evaluation here, I focus on the contribution of concepts to science in terms of their contribution to the formation of reliable expectations (Goodman, 1955) in the pursue of their epistemic goal. I will mention the other virtues at the end of the section. This way to capture the epistemic value of a concept is evoked in the literature on perception. Phillips

writes:

The most important requirement is that [the] members [of the class of perceptual states] must share certain distinctive properties the appeal to which gives us scientific explanations and inductive inferences that we wouldn't otherwise have at our disposal. (Phillips, 2019, p.2)

Phillips uses the notion of natural kind to secure that his class of perceptual states is epistemically useful: by sharing distinctive properties, perceptual states are contributing positively to science in providing explanations and inductive inferences. When the assumption of naturalkindness is discarded, what remains is the aim to provide a scientific class of perceptual objects - in other words, a concept of PERCEPTION - that generate good and useful inferences and explanations. This goal can be captured, secured and evaluated with the notion of *concept utility* (Egré and O'Madagain, 2019), without any mention of natural kinds.

6.1.1 Concept Utility

I suggest to study the contribution of a concept to scientific explanation and inductive inference with the model of concept utility developed by Egré and O'Madagain (2019).

According to them, a concept utility is a combination of two competing aspects of a concept: its *inclusiveness* and its *homogeneity*. The inclusiveness of a concept corresponds to the range of its extension. The inclusiveness of a concept increases with the number of items included in its extension. For example, the concept of animal is more inclusive than the concept of dog, which is more inclusive than the concept of labrador. The second aspect - homogeneity - is a measure of the similarity between the items falling under its extension. In this perspective, the concept of labrador is more homogeneous than the concept of dog, because all the labradors share more common properties than all the dogs.

Inclusiveness and homogeneity of a concept are competing against each other as the more inclusive a concept, the less homogeneous it tends to be. Concept utility is the product of its inclusiveness and its homogeneity, so that a concept is better than another when the combination between its inclusiveness and homogeneity is better.

But why do these two aspects of concepts say something about their usefulness in scientifically categorizing a domain of inquiry? Egré and O'Madagain derived their theory from epistemic utility theories, so that the "value of a concept [is understood] in terms of the value of the beliefs we are inclined to form using that scheme" (Egré and O'Madagain, 2019, p.3).

According to epistemic utility theories, the value of a belief (or what makes a belief useful) is constituted at least partly by its plausibility and its informativeness (Huber, 2008). The plausibility of a belief is how likely it is to be true, while its informativeness concerns how much information it gives about the world.

According to Egré and O'Madagain, homogeneity and inclusiveness of concepts correlate respectively with plausibility and informativeness of beliefs they enable to form. The idea is that the more homogeneous a concept, the more plausible the generalization made with it. For example, if I have a very affectionate labrador and want to generalize this

property, the proposition "labradors are very affectionate" is more plausible than "dogs are very affectionate". On the other hand, saying that "labradors are very affectionate" is less informative than "dogs are very affectionate", in the sense that in the latter I say something that apply to more situations or things in the world.

Theoretical model

Egré and O'Madagain's model is made to evaluate competing conceptual schemes of a specific domain (i.e. an ensemble of objects or phenomena). A conceptual scheme is a set of categories in which items of a domain are classified.

Let's take a toy example (similar to Egré and O'Madagain's one). For a domain constituted of 3 items (o_1 , o_2 and o_3), classified into 2 categories (C1 and C2), the model allows us to compare which classification is best between:

- $P1 = \{C1 = \{o_1, o_2\} ; C2 = \{o_3\}\}$ (read as the conceptual scheme in which a concept C1 applied to the items o_1 and o_2 and a another concept C2 applied to o_3)
- $P2 = \{C1 = \{o_1\} ; C2 = \{o_2, o_3\}\}$
- $P3 = \{C1 = \{o_2\} ; C2 = \{o_1, o_3\}\}$ ²

The conceptual schemes are evaluated with respect to several *features* (let's call them F_x) that are relevant in categorizing the items of the domain. One may build a table of description of the domain such as:

	F_1	F_2	F_3
o_1	1	1	1
o_2	0	1	1
o_3	1	0	0

Figure 6.1: table of description (toy example)

(from Egré and O'Madagain, 2019, p.10 – the value "1" means that the item possesses the feature)

From this table, one may calculate the concept utility of each conceptual scheme P1, P2 and P3. I explain the calculation for P1.

I first measure the inclusiveness of C1 and C2. A good measure of inclusiveness is the proportion of items of the domain included into a concept:

- $Incl(C1) = 2/3$, as C1 includes 2 items among 3
- $Incl(C2) = 1/3$, as C2 includes the remaining item o_3

Now, I measure the homogeneity of C1 and C2. The homogeneity is calculated with respect to each feature. It is measured by the proportion of items falling under the extension

²I rule out the conceptual schemes with a single category $C = \{i_1, i_2, i_3\}$ and with 3 singleton-categories, as I will not use them in my application.

of the concept having or lacking the feature (the highest proportion). The idea is that having a feature in common is as informative as lacking a feature in common for measuring the similarity between the members of a category. The result is rescaled by multiplying by 2 and subtracting 1, so that the measure is comprised between 0 and 1, with 0 meaning no similarity at all.³

- $Hom(C1, F1) = 0$, as o_1 has F1 and o_2 lacks it. The proportion is $1/2$, rescaled to 0 (we want to say that there is no similarity between o_1 and o_2 with respect to F1)
- $Hom(C1, F2) = 1$, as o_1 and o_2 both share the feature F2
- $Hom(C1, F3) = 1$

I average the score to find a value of homogeneity for C1: $Hom(C1) = (0 + 1 + 1)/3 = 2/3$

For C2, all measures of homogeneity is 1, as there is only one item in C2, so that $Hom(C2) = 1$

Now, the measures of inclusiveness and homogeneity are combined to get a measure of concept utility for each concept:

- $U(C1) = Inc(C1) \times Hom(C1) = 2/3 \times 2/3 = 4/9$
- $U(C2) = 1 \times 1/3 = 1/3$

The average of $U(C1)$ and $U(C2)$ gives a score for P1: $U(P1) = 7/18$.

By calculating the same way for P2 and P3, I get: $U(P2) = 1/6$ and $U(P3) = 5/18$

Comparing these score gives us a comparison of the three conceptual schemes in terms of concept utility.

The best conceptual scheme in this toy example is P1, and the worse is P2, as $U(P1) > U(P3) > U(P2)$.

6.1.2 Concept utility of versions of k-PERCEPTION

I apply here the model of concept utility to three different versions of k-PERCEPTION (the concept of perception built in the knowledge epistemic goal of PERCEPTION). Before performing the evaluative calculation, I need to build the evaluative matrix, with the list of items and features. I begin by explaining how I determine the list of items, then I justify my choice of features. In a third step, I present the four versions of k-PERCEPTION I chose to evaluate, and finally I perform the evaluation.

Determination of items

A lot of things have to remain fixed while proceeding to a conceptual evaluation. As I already mentioned, all object-level disagreements are banned for the moment. Furthermore, all the conceptual schemes have to be evaluated with the same rules. That is why, both the items and the features that are deployed in the model should remain the same across

³see explanation in Egré and O'Madagain, 2019, note p.11

all evaluated schemes.

The fact that the items should remain fixed seems unproblematic. They are the set of item that constitute the domain of inquiry. The evaluation is about different conceptual schemes of the same domain. Therefore, the set of items to classify should remain the same.

However, determining the domain for my purpose here and its items is challenging. Egré and O'Madagain's model can directly be used for sets of concrete objects. Their own application focus on classifications of celestial bodies in the solar system. In their application, the domain is the ensemble of celestial bodies of the solar system. All the objects can be listed individually and can directly fill in the evaluation table. The case of mental objects – states or processes - is different. I cannot individually list all instances of mental states. I have to reduce the list of items to make the evaluation tractable. I reduce the list in two ways: (1) by narrowing the domain of conceptualization and (2) by using "kinds of mental objects" as items in my evaluation.

In my evaluation of conceptual schemes, the domain should be the ensemble of *all mental objects*, namely the mind⁴. PERCEPTION would indeed be one of the concepts of a broader conceptual scheme gathering all kinds of mental objects. I chose to narrow the domain in focusing on the alleged perceptual states/processes and some mental objects in their *conceptual neighborhood*. The notion of conceptual neighborhood is quite vague and I suspect that most of mental objects can be said to be in the neighborhood of perception with an acceptable justification (e.g. emotion, imagery).

However, I chose to focus on the distinctions between perception, sensation and belief, partly because they constitutes the relevant borders for the knowledge goal of perception⁵, and also because these boundaries received more attention in the literature (which provides me different concepts to feed the evaluation). A more complex evaluation could be done by extending the domain to other legitimate mental objects in the neighborhood of perception.

I used the concept SENSATION here to express sensory states or mechanisms that are not considered to be perceptual, but anterior to perception. My evaluation will be performed on conceptual schemes constituted of three concepts: SENSATION (or pre-perceptual objects), PERCEPTION and BELIEFS (or post-perceptual objects). Framing perception with two other neighbor concepts allows me to highlight *boundaries* of perception.

Ideally, the model should be adaptable to all the sensory modalities but here it is exclusively concerned by vision.

The second way to reduce the list of items is to already group individual mental states or processes into kinds and use these kinds as items in my evaluation. This move is con-

⁴This could be a wonderful, but gigantic research program, in which results from all the disciplines in cognitive science would be taken into account to conceptually map the mind. This program would answer a question like: given our current state of knowledge about mental mechanisms, which are the best ways to classify them. This research would help organize research, but could also be generative in identifying "conceptual holes". Imagine this map like the periodic table of mental mechanisms.

⁵By including beliefs in the conceptual scheme, I also restrict the evaluated concepts to anthropocentric versions of perception - or at least to creatures that possess beliefs. I accept this limit as I am comparing versions of k-PERCEPTION, which is a concept of perception used in explaining knowledge. I suppose that working on knowledge is already anthropocentric in a sense.

troversial as the kinds I use are themselves conceptual categories that could be evaluated through a similar conceptual evaluation. Furthermore, I have to be careful in grouping the states for not *counting them twice*, as they have to reflect 'flatly' the conceptual domain. Consequently, the kinds of mental objects that can be used as items in my evaluation should (1) be quite consensual, (2) follow a principle of grouping so that two individual objects cannot be put in two distinct groups.

Finding a principle of grouping here is not easy as it has to not pertain to the evaluation I am trying to perform. It means that the principle must be different from the features used in the evaluation table. First I need to be clear on what my individual objects are. As my focus is on the knowledge goal of perception, I am interested in whether and how information can be conveyed, transformed, preserved from retinal receptors to mechanisms of belief formation. Therefore, an individual object would be for me a piece of information. It may be easier to think of it as mental states, but only as informational states (and not as contentful states, in the philosophical sense of mental states with content)⁶. In order to group these informational states and to avoid grouping them twice, I follow a schematic informational route from the proximal stimulus to belief formation processes. I group together informational states produced in a specific stage along this informational route. I assume that psychological mechanisms transform information, so that after each 'manipulation', the information has changed in some ways (content, format, relation between its parts). These alleged differences will be the targeted features in the evaluation. My items are:

1. Retinal information
2. 2 ½ Sketch (Early vision)
3. Perception of volumetric shape and distance (3D Model)
4. Perceptual Object Representation
5. Scene perception
6. Event Perception
7. Perception of basic-level categories
8. Perception of expert categories
9. Perceptual demonstrative thoughts
10. Paradigmatic Belief

This informational route does not have to be thought in a series. I just assume that there are some mechanisms in the neighborhood of perception that transform information into one of these types of informational states. This is really quite schematic, and assume quite a lot. I am more than open to discuss this determination of items, but let's say that it is sufficient to perform a relevant evaluation of conceptual schemes. Remember that all conceptual schemes will be evaluate with the same list of items.

Here are some additional comments on the items:

⁶Having a content is a feature that could be used in the evaluation, therefore I do not use it in determining the items.

- The $2^{1/2}$ D sketch is defended by Maar (1982) as being the output of early vision. It includes depth, but not volumetric shapes and distances. The latter are included in the 3D model, that is partly constituted by representations of unseen surfaces. According to Maar, this 3D model is formed following a matching process between the incoming perceptual signal and a catalogue of 3D models.
- Scene perception is the perception of the spatial structure of the real world in terms of position and functional grouping of objects (Kaiser et al., 2019). Scene perception should be distinguished from low-level group perception, in which group of objects are perceived together as a whole following some Gestalt rules of similarity (like a flock of birds flying together in the same direction or the speckles of the speckled hen). This low-level group perception is included in early-vision.
- Event perception corresponds to the perception of structured segment of time (with a beginning and an end). It includes preception of simultaneity, succession, causality. In a broader perspective, I also include the perception of types of event (e.g. containment, occlusion events, see Strickland and Scholl, 2015) and perception of action and animacy (Proklova et al., 2016; Scholl and Gao, 2013).
- Basic-level categories (Rosch, 2002) are the quickest retrieved category of a perceived object, and the more efficient information to detect an object. "Dog" is more basic than "animal" (supra-level) or "labrador" (infra-level).
- Expert categories are high-level properties that do not belong to the group of basic-level categories. They are roughly categories recognized only by experts of a domain. For example, "pine tree", "the schematic representation of benzene", or "gala apple". They are assumed to be more difficult to retrieve than "basic-level categories" and require a long learning process.⁷
- Perceptual demonstrative thoughts are demonstrative thoughts that are perceptually grounded, i.e. a part of their content is taken from perception (e.g. thinking "That house is big" when looking at a big house).
- Paradigmatic belief is a belief about something that is not in the immediate environment of the subject. Something like "Paris is the capital of France".

Of course, this list of items could be attacked in many fronts, especially in not being *exhaustive* of the domain of interest. They may be some other informational states that are difficult to categorize in one of these ten groups (evaluative perception, social perception, etc.). However, I think that this categorization of the domain is less controversial than classical distinctions between beliefs and perceptual states. These categories here are less inclusive than PERCEPTION and BELIEF (in Egré and O'Madagain's vocabulary), so that the members of each group are supposedly quite similar, and share some structural features. Furthermore, by using this list of items, I can build some of the controversial conceptual schemes in literature that we want to evaluate.

Some may be worried that I do not include in my items "illusions" and "hallucinations", as they are debated perceptual states. I did not include them as they do not follow my

⁷Both basic-level categories and expert categories may include natural kind properties, and artificial kind properties. They may also both include evaluative properties such as "looking disgusting", as long as evaluative properties can behave like basic-level categories in their psychophysical profile (e.g. quick and easy retrieval).

principle of grouping individual mental objects. In other words, they are not types of informational states that are formed at a specific stage of my schematic informational route. But, don't worry. First, they have been counted. They are distributed across all the other groups – at least those whose members can be illusory or hallucinatory. Secondly, these notions will nonetheless play a role in the choice of the relevant features to perform the evaluation.

Choice of relevant features

Unfortunately, Egré and O'Madagain do not explain how to select the features (see Egré and O'Madagain, 2019 note n°7, p.9). Yet this is a crucial issue for the evaluation insofar as the outcome can completely change depending on the features that are selected. Once again, I rely on the epistemic goal of the partitions I evaluate, namely explaining knowledge and belief formation. This epistemic goal gives indeed an epistemic standard to rationally compare two concepts or models. As the epistemic goal of a concept sets research questions for which the concept is used, two versions of a same concept can be compared in their way they fulfill their role in answering these research questions. My choice of features is done in this perspective: within the knowledge goal of PERCEPTION, evaluative features are *properties of informational states that are expected to play a role in epistemological questions*. The features must reflect important notions recruited in epistemology, especially in grounding beliefs in perception. The question is: what are the features of informational states that are relevant for a theory of belief acquisition/knowledge through perception? In this way, the the scores of evaluated partitions will correspond to partition utility in answering epistemological questions, as highest scores are reached by partitions that allow the best inferences (in terms of informativeness and plausibility) when it comes to epistemological questions.

Let's go very shortly in the epistemological field. Perception is related to belief and knowledge in three different ways (Audi, 2011):

1. Causal relation: beliefs are causally grounded in perception. It means that perceptual states or mechanisms cause beliefs. Questions on the causal ground are questions about beliefs acquisition through perception. A simple view on this question is that acquisition of beliefs is done by mere endorsement of perceptual content. But this simple view is debated based on the alleged difference in kind of content and format between perceptual states and beliefs. Therefore kinds of content (non-conceptual vs conceptual) and kinds of format (iconic vs discursive) are relevant features for questions on the causal relation between perception and beliefs.
2. Justificatory relation: perceptual states give reasons to believe things about the world. Two questions arise about justification. First, what kind of states can count as a reason? Questions about format and content are still present here. Secondly, why do we think that perceptual states give such reasons? Different theories will of course give different answers to this question, but one may identify a fundamental idea in the fact that perception connects us with the world in some ways. Perception gives us access to objects and properties of the world, and does it in a correlational way, such that usually, when I see an elephant in front of me, there is an elephant in front of me. Two features reflect these ideas: perceptual states present in some sense objects and properties of the world, and in normal cases, perception arise from sensory inputs caused by worldly objects and properties.

3. Epistemic relation: perception also lead in some sense to knowledge (understood as justified true belief). Putting aside the question of justification (tackled by the justificatory relation), the epistemic ground is more specifically about the involvement of perception in the truth (or falsehood) of beliefs. The question arises because of the existence of sensory illusions or hallucinations. Sometimes, it seems as if there is an elephant in front of me, when actually there is something else (illusion) or nothing at all (hallucination). Illusions arise in cases of misidentification or bad structuration of information into worldly objects and properties. It is therefore possible each time a worldly object or property is presented through perception. Hallucinations arise in cases in which there is identification (but not structuration of sensory input, as there is no sensory input).

Given this quick overview of epistemological questions linked to perception, I identified five relevant features of evaluation for partition utility:

1. Conceptuality
2. Iconicity
3. Ability to represent the environment
4. Stimulus-driven
5. Subject to hallucination

Conceptuality grasps the alleged difference in content between perception and beliefs that raises doubts on the soundness of explaining belief acquisition by mere endorsement of perceptual state content. *Conceptual content* may be understood differently across theories. Here, I equate conceptual content with a content that requires the possession of a concept. It therefore says nothing on the structure or format of the informational state (like being constituted by some mental representations called concepts). I favored this understanding because the format (iconic vs discursive) is another feature of the evaluation, so that the two features about content and format are not redundant.

In terms of features ascription, I consider that the items from perception of basic-level categories to paradigmatic beliefs are conceptual in this sense (items n°7 to 10). For example, it seems obvious that in order to categorize an object as a dog in your visual field, you must possess to some extent the concept DOG. On the contrary, it is not too controversial to assert that people are able to perceive animacy, some actions, causality, objecthood and some features of natural scene, without knowing the concepts of ANIMACY, CAUSALITY, OBJECTHOOD, etc.

Iconicity is understood here the same way I present it in the second chapter (section 2.3.2). In particular, information is delivered in iconic format when there is no canonical decomposition of it. I therefore consider that retinal information, 2 ½ Sketch and the 3D model provide iconic information (items n°1 to 3), while all the other items are non-iconic. Like conceptual content, iconicity of information is recruited in questions about belief acquisition and justification. How may we acquire discursive beliefs or justify discursive beliefs through our senses if the information is obtained in the first place in an iconic format?

The third feature, **the ability to represent the environment**, distinguishes information that *is about* the *distal* stimulus of perception, versus the *proximal* stimulus of perception. Paradigmatic beliefs are about the world, however retinal information is the direct consequence of the proximal stimulus (the light captured by the retinal receptors) and cannot directly be understood as representing the environment. Therefore, mere retinal information can hardly justify beliefs or lead to knowledge. The states that represent the environment are also those that display constancies (Burge, 2010), as constancies are the mechanism par excellence that allow transformation of proximal information to representational one (see chapter 4 on the computational goal, p.89). I consider that informational states from the 3D model to paradigmatic beliefs are able to represent the environment (items n°3 to 10).

Stimulus-driven informational states are those that are caused by a *distal* stimulus (via a proximal one). I count veridical perception and illusions as stimulus-driven states. However, I exclude hallucinations. This feature is here to say something about the origin of information (the world vs the perceiver). This feature also capture the connection of perception with the world, which in turn may be used for theorizing about justification or knowledge. All informational states from retinal information until perceptual demonstrative thoughts (all informational states except paradigmatic beliefs) are stimulus-driven. It means that they arise in presence of a corresponding sensory input (proximal *and* distal).

The feature of being **subject to hallucination** applies to informational states that can exist without proximal stimulus (in hallucinatory cases). This feature tracks the involvement of perceptual states in truth of true beliefs. Contrary to the previous feature, this one reflect the breakable connection of perception with the world. This feature is ascribed to items that can lead to false beliefs about the world. I consider that states subject to hallucination are also potentially subject to illusion. Items from object perceptual representations to paradigmatic beliefs are subject to hallucinations.

As a consequence, I get the following chart:

	F_1	F_2	F_3	F_4	F_5
Objects	Conceptual	Iconic	Representational	Stimulus-driven	Hallucination
1. Retinal information	0	1	0	1	0
2. 2 ½ Sketch	0	1	0	1	0
3. 3D Model	0	1	1	1	0
4. Perceptual Object Rep.	0	0	1	1	1
5. Scene perception	0	0	1	1	1
6. Event perception	0	0	1	1	1
7. Basic-level cat.	1	0	1	1	1
8. Expert categories	1	0	1	1	1
9. Dem. thoughts	1	0	1	1	1
10. Beliefs	1	0	1	0	1

Figure 6.2: table of description k-PERCEPTION

One may object that I forgot a lot of relevant features such as voluntariness, conscious-

ness, subject to adaptation, etc. I considered that they were not (or less) relevant to the knowledge goal of PERCEPTION than the five features I chose for the evaluation.

Versions of k-PERCEPTION

The conceptual schemes I evaluate reflect different borders of perception that have been defended in the literature within the knowledge goal. There are two alleged boundaries of perception: the *lower* border and the *upper* border. The former is the border between perception and something that is-not-yet-perception, that I subsume under the concept SENSATION. The upper border is the famous but controversial border between perception and beliefs. These two borders are reflected by conceptual schemes constituted by three categories: pre-perceptual items (sensation), perception and post-perceptual items (belief-like states). For short, I call these concepts respectively SENSATION, PERCEPTION, and COGNITION.

I evaluate four conceptual schemes (or partitions of the items), constituted by the concepts SENSATION, PERCEPTION, and COGNITION. I present them one by one and provide a summary at the end for clarity.

- **P1** (Limited perception): PERCEPTION applies to informational states that have the ability to represent the environment (vs SENSATION which contains states constituted by proximal information and its transformation). This lower border of perception is advocated by Burge (2010), for whom perception is constitutively a representational competence, for it to be capable of veridicality (or failure of veridicality) (Burge, 2010, p.379). Concerning the upper border, perceptual states may have conceptual content, in the sense that their output is basic-level categories (the quickest categories retrieved from a perceptual stimulus). The idea that perceptual states includes basic-level categories is defended among others by Fodor (1983) and Mandelbaum (2017), for which it is the output of the perceptual module. Therefore, PERCEPTION applies to informational states from 3D Model to basic-level categories (items n°3 to 7), COGNITION applies to informational states including expert categories, demonstrative thoughts and paradigmatic beliefs (items n°8 to 9) and SENSATION is constituted by retinal information and early-vision (items n°1 to 2).
- **P2** (Very early perception): This partition considers that perception is constituted by the sensory information before any individuation of objects. It reflects positions of authors rejecting 'perceptual object representations' as genuinely perceptual (Carey, 2009; Piaget, 1954; Spelke, 1988) and for which perception presents "a continuous layout of surfaces in a state of continuous change" (Spelke, 1988, p.229). Perception is said to be restricted to iconic states. In this partition, SENSATION includes only retinal information (item n°1), PERCEPTION contains 2 ½ sketch and 3D model (items n°2 and 3), and COGNITION begins with perceptual object representations (items n°4 to 10). In this partition, COGNITION includes *core cognition* (Spelke, 2000; Carey, 2009; Jenkin, 2020), which is famously known to blur the classical border between perception and cognition.
- **P3** (Early perception): This partition is similar to the previous one for the lower border (they have the same concept SENSATION). The difference lies in its upper border, that includes in perception *perceptual object representations* and *event perception*. This partition corresponds to the view of some modularists (Quilty-Dunn, 2019;

Scholl and Gao, 2013), which excludes cognitive influence of perception, but argues for a perceptual capacity able to make sense of the sensory information (especially in parsing objects, and identifying events and natural scenes - abilities performed by core cognition in P2). Therefore, PERCEPTION includes items from 2 ½ sketch to event perception (items n°2 to 6).

- **P4 (Enlarged perception):** This partition follows the lead of proponents of high-level perception such as Siegel (2006), for which high-level properties can be part of the content of perception (and are not restricted to basic-level categories). For this partition, PERCEPTION includes the representation of basic-level categories *and* expert categories, as well as the representation of low-level categories in the 2 ½ Sketch. (items n°2 to 8). SENSATION is also composed of retinal information, and COGNITION applies to demonstrative thoughts and paradigmatic beliefs (items n°9 and 10)

Here is a summary chart of the four partitions:

P1 (limited perception)

SENSATION	i_1, i_2	retinal information, 2 ½ sketch
PERCEPTION	i_3, \dots, i_7	3D model, POR, scene perception, event perception, basic-level categories
COGNITION	i_8, i_9, i_{10}	expert categories, demonstrative thoughts, paradigmatic beliefs

P2 (very early perception)

SENSATION	i_1	retinal information
PERCEPTION	i_2, i_3	2 ½ sketch, 3D model
COGNITION	i_4, \dots, i_{10}	POR, scene perception, event perception, basic-level categories, expert categories, demonstrative thoughts, paradigmatic beliefs

P3 (early perception)

SENSATION	i_1	retinal information
PERCEPTION	i_2, \dots, i_6	2 ½ sketch, 3D model, POR, scene perception, event perception
COGNITION	i_7, \dots, i_{10}	basic-level categories, expert categories, demonstrative thoughts, paradigmatic beliefs

P4 (enlarged perception)

SENSATION	i_1	retinal information
PERCEPTION	i_2, \dots, i_8	2 ½ sketch, 3D model, POR, scene perception, event perception, basic-level categories, expert categories
COGNITION	i_9, i_{10}	demonstrative thoughts, paradigmatic beliefs

Figure 6.3: Versions of partitions of the domain around k-PERCEPTION

Model Application, results and interpretation

After calculation (see Appendices), the comparison between partitions utility gives the following ranking: P4 (Enlarged perception) < P3 (Early perception) < P2 (Very early perception) < P1 (Limited perception), with respective final round scores of 0.23, 0.25,

0.27 and 0.28.

These final scores should be taken with care, as they do not represent any measurement in concrete units. Some elements are useful to understand the results:

1. The theoretical maximum of partition utility for a 3-categories partition is $1/3$. This theoretical maximum is not reachable with the matrix here. However, this score of $1/3$ can be considered as the highest score in terms of optimality for a 3-categories partition.
2. The theoretical minimum of a partition utility is 0. It corresponds to partitions in which every category has a homogeneity of 0, which is possible only when the members of each category have no resemblance with respect to each feature of the evaluation (half of the members display the feature, and half of them does not). This theoretical minimum is very unlikely, and in our application seems even impossible.
3. In order to give a better idea of the range of possible applied results, I evaluated a random partition with the following categories: for i_x the item $n^\circ x$, $C1 = \{i_3, i_5, i_6, i_8\}$, $C2 = \{i_2, i_7\}$ and $C3 = \{i_1, i_4, i_9, i_{10}\}$. The final round score of this partition is 0.16.

These scores should not be over-interpreted, and I suggest to focus on the ranking, which gives us a tool to compare the optimality of partitions, given all the decisions taken in the evaluative apparatus. In the limits of this evaluation, the first partition (limited perception) is more optimal in terms of partition utility than all the other ones. Limited perception includes in PERCEPTION: 3D model, POR, scene and event perception as well as perception of basic-level categories. The retinal information and the 2 ½ sketch comprise SENSATION, and COGNITION begins with the representation of expert categories. This partition of the domain 'sensation-perception-belief' is more likely to produce useful theoretical beliefs within the knowledge goal of PERCEPTION.

What is interesting is that the concept of PERCEPTION from the first partition, including informational states from the 3D model to the basic-level categorisation, is the best concept in terms of concept utility, without having any essential feature (none of the features is specific to items included in PERCEPTION). In particular, the items do not share a common kind of content or format, and are not equal with respect to hallucinations. However, these features nonetheless play a role in grounding the category PERCEPTION. *Most* of the items including in PERCEPTION share nonetheless common features, so that grouping them together allows an epistemologist to use this category with success. In other words, this concept of PERCEPTION will lead to build better theories of knowledge (with a better balance between informativeness and plausibility) without having a distinctive feature shared by all and only members of the category.

I am very aware of the high number of limits that such evaluative apparatus encounters. I review some of them in the next section and answer to some objections. The main goal of this conceptual evaluation is to give rational grounds to prefer a conceptual scheme over another one in absence of obviousness of superiority. It is a way to accept a scheme among several acceptable schemes, without yielding to wild pluralism or complete relativism in which all conceptual schemes are equal and reflect a deep theory-ladenness of all theories.

In this enterprise, I tried to show that, even if determining the extension and intension of PERCEPTION is impossible following the naturalkind strategy (see chapter 2 and 3),

there is nonetheless rational grounds, linked to real features of minds, for shaping the concept in one way instead of another, so that PERCEPTION still refer to a group of states/processes/mechanisms that share common properties (even if these properties are not strictly speaking essential to the kind). However, the possession of properties is not the only judge in building a useful conceptual scheme. This evaluation is partially based on rational choices in (1) the relevant features and (2) the domain of conceptualization.

6.1.3 Comments and limits

First, I remind you that my evaluation leaves aside all object-level debates. Beliefs about the world are supposed to be hold fixed. In practice, the filling of the evaluative matrix (the possession of a feature by an item) is controversial for some cases. This is problematic because it has influence on partition utility of evaluated conceptual schemes. The first limit is therefore that an evaluation takes place given specific data about the world. However, this limit is not a mere disadvantage, as it says that a conceptual scheme is not disconnected to empirical data. If new scientific discoveries modify the domain of conceptualization, either by adding an item, adding a relevant feature, or changing the matrix of feature ascription, the ranking of old conceptual schemes can be deeply changed, so that it may be better to abandon a conceptual scheme in the light of scientific discoveries.

The first limit is general to Egré and O'Madagain's model of conceptual evaluation. In what follows, I elaborate limits specific to the application of the model to the conceptualization of the 'sensation-perception-belief' domain. I already mentioned some of them, but I remind them here in a more systematic way:

On the items First, the number of items to be classified influences the computation of partition utility (as it modifies the score of inclusiveness of a concept). When the domain of conceptualization is constituted by individual concrete objects (like celestial bodies in the solar system), the number of items is not subject to debate within the framework of the evaluation (it may be debated how many celestial bodies are in the solar system, but this is an object-level debate). However, in our application here, the items are already kinds of mental objects, so that their number is not a reflection of any real feature of the domain of conceptualization. It may be objected for example that the item " $2^{1/2}$ sketch", understood as the ensemble of informational states produced by the transformations of retinal images into $2^{1/2}$ sketches, is a broader kind (in terms of extension) than the others items (it includes several transformations). As the number of items classified influences the final score of utility, the whole evaluation could be impacted by a simple division of " $2^{1/2}$ sketch" into smaller kinds such as color perception, shape perception, luminance perception, motion perception, edge perception and orientation perception. I do not see another answer to this objection than computing a new evaluation with this new evaluative matrix (note that these all 6 items are similar in terms of feature possession in the matrix, which is why I grouped them together in the first place).

After calculation (see Appendices), the scores are: P1 (0.30), P2 (0.27), P3 (0.24) and P4 (0.22), so that the ranking is not changed. What is interesting is that by dividing " $2^{1/2}$ sketche" in smaller items, the difference between the best model and the worst is even bigger.

Furthermore, as individual states are already grouped into kinds, the list of items is not theory-neutral (or at least is more theory-laden than a list of celestial bodies). Contrary to ways of counting celestial bodies, my principle of grouping is not highly consensual. It could be interesting to find another principle of grouping items, and perform a new evaluation with this different list of items. I leave this task to someone else.

On equi-optimality How should we interpret equi-optimality of two conceptual schemes in this evaluation? Two conceptual schemes are equi-optimal when they have the same final score of partition utility within the same evaluative apparatus. The four schemes previously evaluated are not equi-optimal, so that a ranking is possible (even if some differences are quite small). However, it is possible to find two equi-optimal conceptual schemes. For example, the following partition is equi-optimal with P1: for i_x the item $n^\circ x$, $C1 = \{i_1, i_2, i_3\}$, $C2 = \{i_4, i_5, i_6, i_7, i_8\}$ and $C3 = \{i_9, i_{10}\}$. This conceptual scheme could be interpreted as a partition of the domain with the concepts EARLY-VISION, LATE-VISION and COGNITION, close to what could be found in Raftopoulos (2014, 2015). I did not use this partition in my evaluation in the first place as this is a conceptualization found within the computational goal of PERCEPTION (and less on the knowledge goal - but this case of equi-optimality invites pursuers of the knowledge goal to contemplate this partition).

The possibility of equi-optimality means that such evaluation will not always produce a clear ranking of different conceptualizations. Is this a problem? I think not. First, it is important to recall that this conceptual evaluation is always a comparison, so that the scores have no real 'absolute' meaning. Secondly, this comparison is dependent on the choices and data given through the evaluative matrix. Therefore, if an equi-optimality is not welcome in the results, it is always possible to modify (in a meaningful way) the list of features, in order to decide between two equi-optimal schemes.

Furthermore, I think equi-optimality gives a genuine and valuable information about conceptual schemes. It says that within a same epistemic goal, two different sets of concepts may be equally useful in terms of concept utility. It leads the way to a strong defense of conceptual pluralism, without yielding to a conceptual relativism. If an evaluation highlights two equi-optimally good conceptual schemes, this is a good reason to accept them both. However, this evaluation should give bad scores to meaningless partitions, so that it supports the idea that not all partitions are equal. I will say more on conceptual pluralism later in this chapter.

On the temptation of building a maximally homogeneous conceptual scheme

Given the evaluative matrix presented in this application, one may be tempted to build a maximally homogeneous conceptual scheme with the minimum number of categories. Here, it would be the following 5-categories partition: $C1 = \{i_1, i_2\}$, $C2 = \{i_3\}$, $C3 = \{i_4, i_5, i_6\}$, $C4 = \{i_7, i_8, i_9\}$ and $C5 = \{i_{10}\}$. This partition is appealing as one can build an essential characterization of each category (at least by conjunction of the features). However, there is a price to pay. The categories are less inclusive, so that the overall partition utility is quite low ($1/5$). This partition and the random 3-categories partition have the same partition utility. In a sense, when it comes to the participation of a maximally homogeneous conceptual scheme into theories, the gain in plausibility of inferences is counter-balanced by a loss in terms of informativeness.

Furthermore, and this is a more general comment, it is not recommended to forge new conceptual schemes based on an evaluative apparatus like this one. A conceptual scheme should be forged in the light of a theoretical background. The partitions I evaluated are built this way for some theoretical reasons, not because they fit well in a conceptual evaluation. I think it is reasonable to use this kind of conceptual evaluation to give some directions of thought, but this would be flawed to give it the power to forge concepts. It would result in artificial concepts, that would have a very high score of optimality, while they actually do not participate in any theory.

In other words, this kind of evaluation is not meant to grasp the richness of a conceptual scheme, but to highlight their respective theoretical power.

Let's draw a parallel with the Academic Ranking of World Universities (the famous "Shanghai Ranking") in order to illustrate the flawed move of building a conceptual scheme based on the evaluation. The Shanghai ranking of Universities is roughly based on an agregative calculation of five criteria: the number of international Prizes in alumni, in academic staff, the number of papers published in Nature and Science, the number of highly cited researchers identified by Thomson Scientific and the total number of papers published by the academic staff of the institution). The question whether these criteria are relevant for evaluating Universities is highly debated. However, I think that it is consensual that one cannot build a university by only having these criteria as objectives. As Billaut et al. ironically illustrate:

Suppose that you manage a university and that you want to increase your position in the ranking. This is simple enough. There are vast areas in your university that do not contribute to your position in the ranking. We can think here of Law, Humanities and most Social Sciences. Drop all these fields. You will surely save much money. Use this money to buy up research groups that will contribute to your position in the ranking. Several indices provided by Thomson Scientific are quite useful for this purpose: after all, the list of the potential next five Nobel prizes in Medicine is not that long. And, anyway, if the group is not awarded the prize, it will publish much in journals that count in the ranking and its members are quite likely to be listed among the highly cited researchers in the field. This tends to promote a view of Science that much resembles professional sports in which a few wealthy teams compete worldwide to attract the best players. We are not fully convinced that this is the best way to increase human knowledge, to say the least. (Billaut et al., 2010, p.18)

The take home message is that an evaluation is meant to evaluate, not to produce essential grounds or objectives.

Other virtues Epistemic utility is one virtues of scientific concepts, but it is not the only one. It is therefore possible to build other evaluative apparatus that take into account other virtues. I list here several other considerations that may lead to prefer a conceptual scheme over another. Once again, this list is not exhaustive. My main goal here is to lessen the protuberance of epistemic utility in the global evaluation of a concept. Epistemic utility is indeed a crucial point for a scientific concept, but not the only one.

- **Determinacy and vagueness:** The version of a concept may be less determinate than another in the sense that its extension may be vague. At first glance, a determinate concept should be preferable to a vague one in a scientific context, all things being equal. However, I suspect that the determinacy of a concept should be seen in the light of the maturity of a scientific field. In a 'young' scientific context, a vague concept may be source of scientific creativity, while it could be source of confusion in a more mature science.
- **Empirical identification:** It may be more or less easy to identify empirically members of the extension of a concept. For example, identifying a bird by external properties such as "having a beak, wings and a tail" is easier than "having a bird DNA". The facility of empirical identification may be seen as depending on (1) the complexity of empirical apparatus needed for identification, and (2) the reliability of such apparatus. The facility of empirical identification has to be clearly separated from epistemic utility: it is not because the intension of a concept makes its member easy to identify that this concept is epistemically useful. However, all things being equal, a concept allowing an easier empirical identification seems preferable.
- **Relation with broader scientific domains:** I performed the evaluation of conceptual schemes on a mental domain constituted by sensory, perceptual and belief-like states. However, in the idea that cognitive science is looking for a unified theory of the mind, the different conceptual schemes could also be evaluated with respect to broader mental domains. A conceptual scheme that would fit within a broader classificatory structure (e.g. a conceptual scheme which follows a generative principle of classification for the whole mind) would therefore be preferred to a local categorization.

These other scientific virtues are legitimate in entering into a broader evaluation of conceptual schemes, so that the evaluation I performed is far from being exhaustive.

I have to warn however against an *intuitive but deceptive* virtue, namely the *similarity with the ordinary concept*, which is sometimes used to justify a scientific concept. I'll end up with commenting the over-estimated virtue of *matching with empirical data*.

Similarity with the ordinary concept I already justified why my investigation should focus on the scientific concept of perception, and not the ordinary one (see chapter 5, section 5.1.3). However, some authors seem to maintain that it is a virtue for a scientific concept to resemble the ordinary one. There are three reasons why the similarity of a scientific concept with the ordinary one may seem virtuous in science. First, it is a way to secure a topic. Second, it is the starting point of scientific investigation and finally, it grounds some important intuitions. In what follows, I argue that none of the three reasons are convincing, so that the similarity with the ordinary concept is a deceptive virtue.

Phillips is one of the author using the similarity as a virtue. He writes:

One key desideratum is to avoid a perception/cognition distinction that is purely stipulative. Take the explanatorily fruitful distinction between bottom-up and topdown processes. Construing this as the perception/cognition distinction would do too much violence to the folk conception. There is certainly nothing in the ordinary understanding of perception that rules out top-down influences of the intraperceptual variety. (Phillips, 2019, p.3)

Phillips rules out the bottom-up / top-down distinction as a good ground for the scientific concept of perception based on the fact that it is too far from the ordinary concept of perception. I think Phillips is a proponent of the naturalkindness of perception, associated with the idea that both ordinary and scientific concepts of perception track this natural kind. He indeed continues in his article with:

Beyond this, views on just how far a folk notion can diverge from the natural kind in question, while still referring to that kind, are highly contentious. (ibid.)

In this second quote, Phillips highlights that one of the reasons why scientific and ordinary concepts should resemble each other is a worry about sameness of topic.

In his article, Phillips offers a new way to draw the distinction between perception and cognition. In a sense, he suggests a new intension of the concept PERCEPTION in the 'perception-belief' domain. For him, being similar to the ordinary concept is a virtue as it anchors his concept as a concept *of perception*. On the other hand, and for this reason, a concept whose intension would be forged on the bottom-up/top-down distinction is not a version of PERCEPTION, it is just another concept (even if it could be, as he acknowledged, a useful concept). Being close to the ordinary concept is therefore a way to secure the topic of the concept, and this is why it should be a scientific virtues.

I argue that an alleged similarity with the ordinary concept is a bad way to secure the topic of a concept in presence of a better alternative.

First, it is not clear in what consists the similarity and how to measure it. Is it an overlap of extension? An identity of paradigmatic cases? It is easy to find two concepts with an overlap of extension or sharing paradigmatic cases that are not 'similar' in the right sense (e.g. BELIEF and MENTAL STATE). Secondly, securing the topic is made through an alleged sameness of reference. Yet, this sameness of reference holds only thanks to the assumption that both the scientific and ordinary concept refer to a natural kind. However, I want to avoid this assumption in my work. To sum up, a similarity with the ordinary concept secures the topic of a concept in a vague way and rests on a controversial assumption when it comes to PERCEPTION.

The good news is that I already developed a way to secure sameness of topic, without the assumption of naturalkindness, nor the use of the ordinary concept (chapter 4). Two concepts tackle the same topic if they share the same epistemic goal (or the same extension/intension). Phillips, or other researchers who would like to suggest a new version of the concept PERCEPTION (with a new extension/intension), should indeed shows that its suggested concept meets or has the ambition to meet the epistemic goal(s) of PERCEPTION. This would be better than arguing for a vague similarity with the ordinary concept of perception.

The second reason to advocate for the similarity is that the ordinary concept has a role to play in the scientific inquiry as a starting point. In this view, science is considered as a refinement of everyday thinking, as a process of improvement of conceptual schemes and explanations. The similarity of a scientific concept with its ordinary counterpart would therefore be the mark of *being on the good track of science*.

The scientific concept of perception would be an improvement of the ordinary clumsy concept that has been shaped through time, science and empirical data in an incremental process. Here, the similarity is not a sign of sameness of topic, but the sign of a good scientific practice. It could be therefore considered as a scientific virtue. I argue that it is not.

There are two ways of understanding 'being a starting point'. On the one hand, the ordinary concept may be considered as a starting point in an individual conceptual development. In the individual development of a researcher, the ordinary concept of perception may play the role of an anchor to build more complex scientific concepts. However, in this case, this starting point is not the sign of a good scientific practice, but maybe a good learning/teaching practice. It has nothing to do with the development of science. The second understanding, on the contrary, would be to take the ordinary concept as the starting point of the whole scientific inquiry about the topic, at the level of past and present scientific communities. But I think it is false to think that our current ordinary concept of perception is genetically the starting point of our current scientific concepts of perception. Both ordinary and scientific conceptual schemes have evolved through time, sometimes separately, sometimes connected. If there is some old concept at the origin of our scientific inquiry about perception today, I suspect it is long lost in history.

There are some good scientific practices in building new conceptual schemes, such as taking into account recent empirical data, or being aware of ontological assumptions behind conceptual schemes. However, similarity with the ordinary concept does not reflect any good scientific practice. Paying a tribute to a putative antique ordinary concept that gave birth to the research field is nothing but mythological rite.

Finally, the last reason to value similarity pertains to the requirement to ground some intuitions about PERCEPTION. Intuitions are arguably generated by an ordinary conception of the mind. They are used in grounding some important argument about the nature of perception. Let's take for example arguments in favor of non-conceptual content of perception based on perception in children and non-human animals. The argument is roughly the following: there is something in common between our perceptual states and perceptual states of some non-human animal. Yet, some of these non-human animals do not possess concepts. Therefore, there is something about perceptual states that is not conceptual. The first premise of the argument is largely based on an intuition. Intuitively, babies or dogs are seeing the world more or less like us (human adults). How explaining that this intuition is strong enough to ground some philosophical arguments if the ordinary concept (from which it comes from) has nothing to do with the scientific investigation?

One way to answer is to simply reject the value of these intuitions. Consequently, this argument in favor of the non-conceptual content of perception would be unsound because it rests on an intuition. I am quite sympathetic with this answer.

However, Intuitions seem to have a creativity power, useful in science. It is difficult to accept that this power is mere serendipity. Valuable intuitions may be explained by similarities between scientific and ordinary *questioning about behavior* in general (and not a similarity of concepts). Even if the ordinary concept of perception does not have epistemic goal as such (as an epistemic goal is dependent on a scientific community), it is used in everyday inferences in order to explain everyday behavioral facts. The behavior of a neighbor, or a cat, is explained by a combination of perceptual states, beliefs, decisions

and actions: "He saw me but he didn't say hello, yet he must have recognized me. The only explanation is that he avoids me, for sure!". The everyday need of explaining the behavior of others is source of questions: "Is it possible that he merely did not see me, even if he looked in my direction for a few seconds?" The question is legitimate, and can feed a scientific questioning (and the epistemic goal of the scientific concept). Therefore, intuitions coming from an ordinary conception of the mind may be useful in science as it guides scientific inquiry and epistemic goals. But this usefulness is not grounded in any similarity in extension or intension between the ordinary and scientific concept.

I conclude that the similarity between the scientific and the ordinary concepts PERCEPTION is not a scientific virtue, because the three justifications one can suggest to defend it are all flawed.

Still one may claim that this similarity is an advantage for conceptual transmission, in cases of teaching or popularization (which is a non-scientific virtue). It seems always easier to teach a new concept if it is anchored in another – even mistaken – concept. For example, to teach the concept WOLF to a child without pointing at it, one may use the concept DOG and add that a wolf looks like a wild dog, living in the forest. One may construct the similarity as a degree of usefulness in explanation. A concept A is similar to the concept B if explaining B based on A is easier than explaining B as a brand new conceptual tool.

This advantage applies in non-scientific contexts though (such as conceptual learning and popularization of scientific theories). It should also be counter-balanced by a risk of confusion. If I teach the concept WOLF this way to my child, nobody would be very surprised if next time I go to the forest with him and meet a jogger with her dog, my son cries out: "look Mummy, a wolf!".

A good match with empirical data I would like to end this section by a short comment on a scientific virtue that is often over-estimated: a good scientific concept should be a concept determined by empirical data, such that 'matching' empirical data is a scientific virtue.

What does it mean to *match empirical data*? It means that concepts should carve a domain of conceptualization by following the empirical properties of the items. There is a sense in which this thesis is unproblematic. It is indeed one of the principle on which is based the notion of concept utility developed previously (homogeneity of concepts captures the resemblance between items with respect to some features). However, this virtue is dangerous when it is considered as the only virtue of scientific conceptual schemes. It is dangerous because it can never be sufficient in building a conceptual scheme. Properties must be, in one way or another, pointed out as relevant for conceptualization, and mere empirical data cannot give such weighting.

For example, in a workshop (Paris, 2019, April 1st), Quilty-Dunn suggested that the items in the extension of perception can be tracked with retinotopic adaptation. But how ad hoc is it to favour this property as the mark of perception? What I mean is that it is not retinotopic adaptation alone that can track perception, but a whole theoretical construction that ground the idea that retinotopic adaptation is a good sign of what is supposed

to be perception. Empirical data is blind without the eye of the scientist.

Of course, all properties are not objectively equal, only some of them are adequate to build scientific conceptual scheme: this is Sider's view:

We should use those concepts that match up with the objective joints in reality. In rough terms, these will be concepts that pick out those properties that David Lewis describes as natural ones. Such natural properties (as opposed to jerrymandered quark-or-elephant) are ones that mark out objective similarities between things, and that have real explanatory import (Sider, 2011, p.830)

What Sider says is that Nature already chose those properties on which we should build classifications and conceptual schemes. I agree with the idea that not all properties are equal, but I reject that the best conceptual scheme is the one that is based on natural properties, in an *exhaustive* way. The problem is that there are too many natural properties in Sider and Lewis sense. Building useful conceptual schemes requires not to take into account jerrymandered properties, but also to take into account a *limited* number of natural properties, in order to make some understandable explanations emerge.

This thesis deserves a better defense than a few words, but this is not the place to do it. I will emphasize these ideas in the next chapter, in which I analyse Sider's notion of joint-in-nature. For now, I wanted to highlight that it is not sufficient to rely on an alleged good matching with empirical data to justify a choice of conceptual scheme.

6.2 Pluralism and Fragmentation

After performing the evaluation, we can answer the metalinguistic negotiation. The concept of perception of the winner partition (limited perception) is the best concept to investigate the knowledge epistemic goal of perception. The answer is however partial and depends on the epistemic goal of researchers. The upshot is that there is a 'best concept of perception' for each x -PERCEPTION, leading to a pluralist view in which several concepts are all legitimate. This is a form of *conceptual pluralism*. Furthermore, I also showed that two concepts can end up being equi-optimal within the same evaluative apparatus. In this case, the conceptual engineering strategy faces another source of conceptual pluralism.

By accepting that the metalinguistic negotiation can help answering the boundary question, it also means that there could be several different boundaries of perception. The latter claim is a form of pluralism with respect to perceptual boundaries (by contrast with conceptual pluralism). Recently, pluralism about the perceptual border has been suggested by some philosophers (Beck, 2018; Phillips, 2019). Beck writes:

Given the wide array of debates that presuppose a perception–cognition boundary, it would be rather surprising if there were only one legitimate way to draw the boundary (Beck, 2018, p.3)

Here, I defend conceptual pluralism for PERCEPTION⁸. I postpone remarks about pluralism about the boundary after the third part of the dissertation, in which I explore

metaphysical grounds of the perceptual boundaries.

Conceptual pluralism is a recent position for PERCEPTION, but has been defended for other psychological concepts (e.g. CONCEPTS, EMOTION, MEMORY), and is pervasive in philosophy of biology (e.g. SPECIES, GENES)⁹. I will get resources from these other debates to present criticism that such positions already received, and show that the conceptual pluralism produced by my conceptual engineering is rigorous and does not fall into the presented pitfalls.

I even go further and show that conceptual pluralism of PERCEPTION is actually a part of a classical process happening to scientific concepts - conceptual fragmentation (Taylor and Vickers, 2017) – in which it is expected that a concept with a single meaning is fragmented into several sub-concepts used legitimately for different scientific purposes. I will show that recent literature about the boundaries of perception nests all the signs (not only pluralism) that PERCEPTION is currently subject to fragmentation.

In this section, I first present with more details the conceptual pluralism I produce with my evaluation, especially the two sources of plurality (*x*-PERCEPTIONS and equi-optimality). Then, I defend this form of conceptual pluralism against some usual complains made against pluralism in general. Finally, I will present conceptual fragmentation and show that PERCEPTION is currently under a process of fragmentation.

6.2.1 Pluralism and its two sources

I define conceptual pluralism as the thesis according to which there are several *equally legitimate* concepts of perception, *not hierarchically nested* and which demarcate *extensionally different* categories of states. The three conditions (equal legitimacy, absence of hierarchy and different extension) distinguishes conceptual pluralism with more common forms of conceptual plurality.

A first example of conceptual plurality can be found in Block (ms). He argues that there are several criteria to distinguish perceptual states from other states, such as adaptation and constancies. However, they all converge (they classify the states the same way). One could say that Block's view offers different concepts of perception, with different intension, but with the same extension (i.e. defining extensionally equivalent classes of states). I do not consider this as a genuine form of conceptual pluralism. It is important to specify that the concepts in a pluralist position should demarcate *extensionally different* classes of states.

Another example is found in Phillips (2019). He explicitly defends a pluralist view for perception but argues that a general criterion heads all its notions. He writes:

Each legitimate way of marking a border between perception and cognition invokes a notion I call 'stimulus-control'. Thus, rather than being a grab bag of unrelated kinds, the various categories of the perceptual are unified into a superordinate natural kind (Phillips, 2019, p.2)

⁸I actually suspect that Beck's pluralism is best understood as a form of conceptual pluralism as he suggests that the reason for accepting pluralism is the variety of uses of the perception-cognition boundary in scientific debates. This could be interpreted as a mention of the different epistemic goals of PERCEPTION.

⁹See Machery (2009) for CONCEPTS, Griffiths (1997) for EMOTION, Ereshefsky (1992), Kitcher (2007) and Dupré (1981) for SPECIES. See Taylor and Vickers (2017) for a review of conceptual pluralism in science.

Phillips' view is not a genuine pluralism. He argues instead that there are interesting sub-categories in the class of perceptual states, but they are all subsumed under a general concept of perception. In conceptual pluralism, the different concepts should not be subsumed under each other, but be *alternatives* to each other. They should not be hierarchically nested.

Finally, it is important to remind that all the concepts have to be equally legitimate, or a pluralist position could be compared to a list of possible concepts. It will then fall under some critique such as Hull's one:

If pluralism is anything more than the reminder that other serious options exist or may crop up in the future, I do not see what it is (Hull, 1987, p.179)

In this chapter and the previous one, I listed a lot of different concepts (and versions of concepts) for PERCEPTION. I mentioned in the text when pluralism arises, but for the sake of clarity, I summarize here the two sources of pluralism, i.e. the two reasons why I end up with equally legitimate, not hierarchically nested and extensionally different concepts of perception. The two sources are (1) the individuation of concepts individuated with epistemic goals, and (2) equi-optimality in conceptual evaluation.

I presented in the previous chapter four concepts of perception that cannot be compared because they are not use for the same purpose: p-PERCEPTION, k-PERCEPTION, c-PERCEPTION and b-PERCEPTION. They are concepts of perception forged in pursuing different scientific inquiries, respectively the description of conscious experience, the explanation of knowledge, the decoding of the mental computational operations and the explanation of fitness of organisms. These four concepts constitute a first layer of conceptual pluralism. They are indeed all equally legitimate - none of these four inquiries are better, or more important, than another. None of these concepts subsume the others. Finally, I defended that they have different extensions (chapter 5, section 5.3.1). For example, unconscious perceptual mechanisms are not part of the extension of p-PERCEPTION but are contained in k-PERCEPTION.

The second layer of conceptual pluralism appears within these four concepts of perception. I showed that they all have internal variations, leading to several versions of the same concept. Not all the versions of a concept should be accepted as legitimate, but it is possible that several versions are equally good in terms of scientific virtues. The evaluation I performed at the beginning of this chapter was supposed to give us the best version of each concept. This is what happened in the specific application I offered on k-PERCEPTION. However, I mentioned that equi-optimality is possible. I showed indeed that another partition is equi-optimal with limited perception but is not a partition elaborated within the knowledge epistemic goal of perception (section 6.1.3). The upshot of the possibility of equi-optimality in the evaluation is that it is possible that two versions of a concept get the same concept utility (Egré and O'Madagain, 2019). It means that, given the same evaluative apparatus and for different reasons, they produce inference and explanation with the same epistemic utility (with an equivalent ratio between informativeness and plausibility). It is possible that other scientific virtues arbitrate between equi-optimal concepts, but the possibility remains that two versions of a concept are equally legitimate, given all the scientific virtues on which they can be compared. Furthermore, versions of a concept are real alternatives (and not hierarchically nested), as they follow the same epistemic goal (in a sense, they are interchangeable where x -PERCEPTIONS are not). Finally,

they do not have the same extensions (that is why they are different *versions* of a concept).

At the end of the conceptual engineering on PERCEPTION, I end up with a two-layer conceptual pluralism composed of at least four legitimate concepts of perception with different epistemic goals, that can respectively accept several legitimate versions.¹⁰ This may be seen as a bad consequence for the boundary question, as this plurality of extensions does not help in providing a unique answer to feed the debates that require clear perceptual boundaries.

It would not be the first time that pluralism is under a cloud of suspicion. In other areas of science and philosophy, pluralist views have received a lot of criticism (e.g. pluralist views on SPECIES in biology). In the next subsection, I defend the conceptual pluralism of PERCEPTION against some usual complaints made to pluralist views in general.

6.2.2 A decent pluralism

Objections have been developed to block pluralist stances, both in the philosophy of biology and in other philosophical domains (against value pluralism or pluralist theories of truth). I present here some common complaints (in a non-exhaustive way) addressed to pluralism and argue that my conceptual pluralism of PERCEPTION adequately answers the worries. Therefore, it is a decent pluralism, so to speak.

The objections are:

1. Arbitrariness and the risk of over-generation
2. The reduction of explanatory power
3. Inconsistency

These objections (or worries) for pluralism are expressed by proponents of monism (i.e. people arguing for the existence of a single concept). Another powerful complaint exists against pluralism – the risk of confusion between the concepts – but this is a criticism in favor of the replacement of the concepts by others (eliminativism). I tackle this objection in chapter 8.

Arbitrariness and the risk of over-generation The first worry addressed to pluralist views is their alleged arbitrariness. It is not enough to grant legitimacy to several concepts, this legitimacy requires a solid ground. Pluralism is often considered too liberal, in accepting conceptions as legitimate, when they are mere possibilities. Without a strong justification for the legitimacy of its concepts, pluralism seems arbitrary, and exposes the view to over-generation. A decent pluralism should give means to discriminate between legitimate and illegitimate concepts.

Bradley (2009) expresses this worry against value pluralism¹¹:

¹⁰This conceptual pluralism is open to new concepts and new versions, either because I forgot some epistemic goals, or because future research will create some new goals and new conceptual variation.

¹¹This objection also arises against pluralist views of SPECIES in philology of biology (Ghiselin, 1987, p.136).

Pluralism seems objectionably arbitrary. Whatever the composition of the list, we can always ask: why should these things be on the list [of legitimate conceptions of good]? What do they have in common? What is the rational principle that yields the result that these things, and no others, are the things that are good? (Bradley, 2009, p.16)

This alleged arbitrariness has a dogmatic flavor, in contradiction with philosophical thinking. Bradley adds further:

We wanted enlightenment, but we are provided instead with a list, and are told not to look any deeper. This is not theorizing, but a refusal to theorize. (Bradley, 2009, p.17)

The conceptual pluralism of PERCEPTION I defend in my work is however not arbitrary. In each of the two layers of pluralism, there is a clear principle grounding the legitimacy of a concept. In the first layer (*x*-PERCEPTIONS), a legitimate concept of PERCEPTION is a scientific concept, used by scientific communities in the pursue of a specific epistemic goal. There are therefore several ways for a concept to be illegitimate: (i) not being *scientific* (in the sense that it does not have any epistemic goal) (ii) being *pseudo-scientific* (used in pseudo-scientific investigations), and (iii) not being used in an epistemic goal requiring the concept in an explanatory manner. The risk of over-generation is also limited in this first layer, as all concepts sharing the same epistemic goal are gathered together. The legitimacy of the concepts in this first layer is in a sense transferred to their epistemic goals: the legitimate concepts are those having a legitimate epistemic goal. Epistemic goals of PERCEPTION are legitimate when they are pursued by scientific communities using PERCEPTION in an explanatory manner.

The ground of legitimacy is different in the second layer of pluralism (versions of *x*-PERCEPTION). Two concepts are equally legitimate when they are equi-optimal in terms of scientific virtues in pursuing their shared epistemic goal. Contrary to Bradley's worries that the list of legitimate concepts is given without any resources to compare them, the legitimacy of concepts in the second layer is the *product of a comparison*, in which two or more concepts have been evaluated as the best (for different reasons). Considering the evaluation of the versions of *k*-PERCEPTION I did, equi-optimality is the exception, not the rule. For two concepts to be equally legitimate because of equi-optimality, it means that they really are scientifically virtuous compared to the other alternatives.

One may object that arbitrariness remains in the building of the evaluative apparatus, in the choice of items and features. But, even if these choices require a lot of theoretical decisions (that can be objected), these decisions are all justified. The items are chosen given the domain of conceptualization, and the features are chosen for their relevance in the pursue of the epistemic goal of the evaluated concepts.

Consequently, none of the layers of my pluralism are arbitrary. There are good reasons for a concept to be accepted as a legitimate of PERCEPTION.

The reduction of explanatory power Pluralism has been said to reduce the explanatory power of scientific theories because it prevents generalization. If there are several concepts of perception, extensionally different, laws about perceptual states will be relative to a concept, what will reduce the scope of the law. Hull says:

The [...] mistake is to overlook regularities that actually exist, to opt for multiplicity for the sake of multiplicity. The safest strategy in intellectual pursuits is to point out the endless differences that exist in nature without venturing any generalizations about them [...] To the extent that science involves generalization and categorization, the infinite multiplicity is not open to scientists (Hull, 1987, p.168)¹²

It is true that in my view discoveries within an epistemic goal cannot be directly generalized to theories about other *x*-PERCEPTIONS, but it is for the best. First, being able to generalize is great, but only when one is likely to perform good generalizations (generalizing for mere generalizing does not add any explanatory power). Furthermore, I defended in the previous chapter that there are synergies across epistemic goals, and therefore bridges between concepts. Recognizing that concepts are different in my conceptual pluralism highlights that the theoretical transfer across epistemic goals has to be done carefully, in taking into account the difference of perspectives. I do not see that as a reduction of explanatory power, but as an increase of rigor.

Inconsistency Some philosophers of theories of truth argue that pluralist positions are inconsistent (Lynch, 2005; Davidson, 1984). Lynch says:

[Pluralism about truth] is not clearly a pluralist view of truth at all. After all, the idea that there is more than one way for propositions to be true just implies that there is something these ways have in common that makes them all worth calling 'ways of being true'. In denying this, [pluralism about truth] arguably undermines its own credentials as being both pluralist and a theory of truth (Lynch, 2005, p.42)

When applied to the conceptual pluralism of PERCEPTION, the worry consists in saying that it cannot be both *pluralist* and *about perception*. The idea is that all the concepts of perception should share a commonality in virtue of which they are all concepts *of perception*. However, if they can all be subsumed under a common characterization, the position is not really pluralist anymore (it would be a view like Phillips (2019), in which his several concepts of perception are hierarchically nested into a sub-concept). The upshot is that a pluralist view about perception is inconsistent.

This objection is actually a worry about sameness of topic (again). Being pluralist would prevent securing that all the concepts are about the same topic. I have a double answer concerning my conceptual pluralism. In its first layer, I recognize that the four *x*-PERCEPTIONS are not about the same topic. But this is not a problem, because it is part of what I claim. I argue that there are actually (at least) four different topics, that are called 'perception' in research. They are all concepts of perception - broadly construed - because it is how they are thought. But it is actually a mistake.

One could say that my first layer of pluralism is therefore not pluralist, as the concepts are not about the same thing. But there are enough people arguing that these concepts should be merged, rejected or subsumed under another one (Phillips, 2019; Schellenberg,

¹²This quote highlights a hint of disdain for pluralism as an *easy solution*. A decent pluralism is the fruit of a careful analysis. It seems to me easier to declare (without much justification) that a concept is not yet discovered, primitive or vague.

2018), that I think that my view can be said to be pluralist compared to the widespread monist perspective.

My second layer of pluralism (equi-optimality between versions of the same concept) secures sameness of topic through the sharing of the same epistemic goal. It is a way to resolve the inconsistency: the concepts are extensionally different (this is a pluralism), but they share the same epistemic goal (they are about the same topic).

A decent pluralism Based on the answers to the previous objections, I claim that my conceptual pluralism about PERCEPTION is a decent one, as it is non-arbitrary, consistent and useful for scientific research. A last objection will be raised in the last chapter, namely that the multiplicity of perceptual concepts actually increases the risk of confusion in scientific debates and research. The solution offered by enemies of pluralism is to eliminate PERCEPTION from science and replace it by other concepts (in our case, it would be *x*-PERCEPTIONS, or other concepts like top-down and bottom-up processes.)

More than decent, pluralism about PERCEPTION is actually to be *expected*, because it is part of a common scientific process, i.e. conceptual fragmentation. By accepting conceptual pluralism as a normal step in scientific research, I block the misconception about pluralism: the idea that conceptual pluralism is an "exotic doctrine" (Davidson, 1984, p.5). Conceptual pluralism of PERCEPTION is actually a result of the conceptual fragmentation of PERCEPTION, a process that happened to a lot of other scientific concepts before PERCEPTION.

6.2.3 Fragmented PERCEPTION

Conceptual fragmentation is a phenomenon in which a scientific term, thought as having a single meaning, appears to be scientifically used with several distinct meanings. I use here the notion developed by Taylor and Vickers (2017). In their own words:

We will use the term 'conceptual fragmentation' to refer to any case where:
 (i) a certain term, originally widely assumed to enjoy a single meaning, has been found to have multiple distinct meanings no one of which is privileged, and (ii) different definitions are adopted for different theoretical uses. (Taylor and Vickers, 2017, p.20)

Conceptual fragmentation is different from conceptual pluralism. The former is the process of multiplication of definitions that happened to be adopted depending on scientific contexts, while conceptual pluralism is the statement that we should accept these definitions as being legitimate definitions of the concept. I already advocated conceptual pluralism for PERCEPTION. What I want to show here is that the multiplication of perceptual concepts is an expected result in scientific research, so that pluralism is not an exotic doctrine (in Davidson's words).

Pluralism does not follow directly from fragmentation, as conceptual fragmentation can also lead to eliminativism (because a concept is fragmented, we should abandon it). In the last chapter, I argue that we should resist eliminativism for PERCEPTION.

Conceptual fragmentation usually occurs because a term is ambiguous between different approaches or objects of reference. It may also arise when a kind of entity is identified as being constituted by different subkinds. This is a common step in scientific investigation: a group of phenomena or entities, superficially homogeneous, is discovered to be heterogeneous and actually consists in several more specific classes of homogeneous phenomena.

Here is a historical example that I like for its poetry: in its *Novum Organum* (1610), Sir Francis Bacon offers a new method to find the causes of a phenomenon X. He suggests to list all the positive instances of X and to discover the commonalities between them. He applies this method to what he calls "the form of heat". Here is his list of hot things:

The rays of the sun, fiery meteors, burning thunderbolts, eruptions of flame from the cavities of mountains, all bodies rubbed violently, piles of damp hay, quicklime sprinkled with water, horse-dung, the internal portions of animals, strong vinegar which when placed on the skin produces sensation of burning, keen and intense cold that produces a sensation of burning. (Adams and Aizawa, 2010, p.58)

According to Bacon, the common factor between these cases is the high degree of molecular vibration. However, science showed that this list is actually composed of three distinct phenomena with no common underlying mechanisms: (1) heat produced by friction, (2) heat coming from biological decomposition and (3) chemical exothermic reactions. The first phenomenon is today studied by physics, the second by biology and the last one by chemistry.

Taylor and Vickers (2017) gives numerous examples of recent conceptual fragmentation in several scientific domains, and they suggest a three-step causal explanation for the rise of conceptual fragmentation¹³ in all domains of (philosophy of) science. Here is a reconstruction of their proposal for a concept X:

1. First step : There is a debate about X focused on the question "what is X ?". Answers to this question are judged by their extensional adequacy, i.e. whether they "include any cases which certainly are X, and not include any cases which certainly are not X" (Taylor and Vickers, 2017, p.22). This question structured the debate in a succession of definitions-counterexamples-new definitions. This ends up by multiplying the definitions of X, none of which are perfect.
2. Second step : several of the definitions suggested in the first step are actually considered as useful, important, valuable in certain contexts.
3. Third step : Participants in the debate stop believing in the classical theory of concepts, i.e. the idea that a concept is best captured by a definition (following the criterion of extensional adequacy)¹⁴

In the heart of this process lies the death of the classical theory of concept (i.e. the content of our concepts have a definitional structure). It is allegedly because the classical theory is false that the first step is a dead end : there is no perfect account of "what is X ?" in a definitional form (it is always possible to find a counterexample). Fodor writes:

¹³There proposal constitutes a sufficient cause for conceptual fragmentation, but is not necessary.

¹⁴This does not include stipulative definition, or pragmatic definition.

[T]hese days almost nobody thinks that concepts are definitions. . . There are practically no defensible examples of definitions; for all the examples we've got, practically all words (/concepts) are undefinable. » (Fodor, 1998, p.44-45)

The second step is the emergence of new criteria to judge accounts of X, different from extensional adequacy. A concept may be valuable because it is theoretically useful (even if it does not meet the extensional adequacy criterion), or may refer to a subkind of phenomena that are theoretically important in their own way. For the second step to be complete, the community should end finding that more than one concepts are valuable. Finally, for the third step, it is not sufficient that the classical theory of concept is false, it is also necessary that some of the debaters come to believe it.

The result is that the field is now furnished with several useful definitions X and debaters do not feel the urge to pursue the quest of finding the only good definition of X anymore. This is conceptual fragmentation.

This story applies roughly to the history of concepts such as SPECIES, DISEASE or SCIENTIFIC METHOD (Taylor and Vickers, 2017). Considering SPECIES, the concept is used in different ways, depending on whether it refers to ecospecies, biospecies or phylopecies, i.e. lineage individuated by respectively ecological, interbreeding and phylogenetic approaches of species taxonomies (Dupré, 1981; Ereshefsky, 1992; Grant, 1981). These different available definitions became used for *different theoretical investigations*. For example, the notion of BIOSPECIES, i.e. group of interbreeding organism producing fertile offspring, structures and explains the stability and evolution of many organisms. However, some species of organisms do not reproduce sexually, which makes the concept of biospecies ineffective in classifying and explaining evolution of these groups of asexual organisms. Other concepts, such as PHYLOSPECIES or ECOSPECIES are therefore theoretically useful in this specific case.

Ereshefsky defends the idea that each approach of species has its own value, and advocates a form of pluralism:

A taxonomy of monophyletic taxa provides a framework for examining genealogy. A taxonomy of interbreeding units offers a framework for examining the effect of sex on evolution. And a taxonomy of ecological units provides a structure for observing the effect of environmental selection forces. A systematic study that considers just one of these taxonomies provides an overly coarse-grained picture of evolution. (Ereshefsky, 1992, p.678)

Each of the steps suggested by Taylor and Vickers can be tracked by specific kinds of literature. The first step is constituted by works trying to define a concept X, with arguments (or objections) based on the criterion of extensional adequacy. One of the symptoms of this step is the existence of a metalinguistic negotiation among the debater. People may agree on the kinds of phenomena in need of classification, but disagree on how to classify them.

The second step is characterized by a literature in which stipulative or pragmatic definitions are used (of the form, "In this context, it is useful to use this characterization of x" or "For this work, I use X in this (stipulative) definition").

Finally, the last step will see the rise of debates about pluralism and/or eliminativism, which are the two possible roads when one finally reach the conclusion that it is impossible to forge a single perfect definition of X. Note that the steps are not necessarily nicely defined in time. I suspect that they are usually mixed, but the second step and third step are expected to appear at least slightly after the first one in the literature, because they both need the generation of a multitude of definitions produced during the first step.

Each of the three steps of conceptual fragmentation have been illustrated in recent philosophy of perception. It makes me conclude that perception is currently subject to a process of conceptual fragmentation.

Step 1: There is a debate about "what is perception?", in which proponents of a position are using the extensional adequacy criterion as an argument in favor or against a specific account of perception. These arguments are the one used in the metalinguistic negotiation I identified across the positions about the hallmark of perception. I remind you that a metalinguistic negotiation is a representation-level oriented issue, in which what is at stake is how we should use correctly the concept of perception. This metalinguistic negotiation is characterized by a specific objection that I called "the disagreement on extension". According to it, a specific conception of perception is wrong because it does not generate an adequate extension for PERCEPTION. Quilty-Dunn's argument for defending the perceptual nature of object perception is a nice recent example: (i) "Without object representations, vision would be as William James imagined it to be for infants, a 'bloom-ing, buzzing confusion'" (Quilty-Dunn, 2019, p. 22), (ii) yet vision is not this kind of thing, (iii) therefore object perception should be considered as genuinely perceptual.

Step 2: The second step is reached when necessary and sufficient conditions for defining perception are not considered as compulsory anymore for saying something about perceptual phenomena. Consequently, several definitions are considered useful, even if they are not complete or perfect. The acknowledgment that sufficient and necessary conditions are not required is illustrated by the recent effort of Block to define perception with the notion of "joint-in-nature" (as opposed to the notion of natural kind). Hence, Block says that "perception is constitutively iconic, non-conceptual and non-propositional, but these properties are not sufficient for perception. For example, hallucination has all these properties but is not perception." (Block, ms).

Researchers also clarify their terminology at the beginning of their articles, recognizing that there are several uses of PERCEPTION in the literature. Here are two examples:

I have spoken of perception and observation, or vision. These terms are not employed consistently in the literature. Sometimes 'perception' purports to signify our phenomenological experience, and thus includes the recognition and identification of objects and events. Since I do not use the terms the same way - I adopt Dretske's (1985) and Shrager's (1990) usage - I will introduce some terminology to explicate my usage of the terms. [...] I call these processes that transform sensation to a representation that can be processed by cognition *perception*. Perception includes both low-level and intermediate-level vision and is bottom-up. (Raftopoulos, 2001, p. S188).

My usage of the terms 'cognition' and 'perception' is consonant with much of the recent literature on perception (Firestone and Scholl, 2016a; Pylyshyn,

1999), but some restrict the term perception to what I am calling low-level perception (Linton, 2017) and others use 'cognition' to encompass mid-level and high-level perception as both perception and cognition (Cavanagh, 2011). (Block, ms, p.5)

Furthermore, there is a kind of "pragmatic turn" in philosophy of perception, in which the goal is "to uncover those notions of perception and cognition that earn their explanatory keep in cognitive science" (Phillips, 2019, p. 318). The goal in characterizing perception is not to find the right concept of perception anymore, but instead to find a characterization that "will prove to be of theoretical interest" (Phillips, 2019, p. 318).

Step 3: The last of the three steps is characterized by the rise of pluralism. This is a sign that people are giving up the classical theory of concept, and accept the idea that a concept can be grasped by a multiplicity of contextual characterizations. I said earlier that some pluralist views about perception emerge in philosophy of perception (Beck, 2018; Phillips, 2019). Their view would lead to genuine pluralism only if they accept that different extensions of PERCEPTION can be *legitimate*. They have to hold an anti-definitionist view, i.e. a view that reject the classical theory of concept. By contrast, Phillips (2019) argues that a general criterion heads all of his borders of perception, what unifies his concepts under a single superordinate natural kind. His pluralism is still shaped by the idea that a pluralistic account of a notion should serve an ultimate goal of unification, based on the common property of the subphenomena he describes. This is no surprise, from a historical point of view, if we consider that the third step of conceptual fragmentation is in its infancy for the concept PERCEPTION in contemporary western philosophy of perception. I allow myself a small prediction: pluralistic accounts of PERCEPTION and the perceptual border will multiply in the near future. My dissertation will be one contribution among others. ¹⁵

As I evoked earlier, conceptual fragmentation can lead to different positions about what we should do with the concept PERCEPTION, and how we should use it now. Two powerful roads are open: pluralism and eliminativism. The former states that we should keep the concept PERCEPTION, while acknowledging that this concept can only be grasped to a multiplicity of distinct and legitimate definitions (this is my view in defending conceptual pluralism). On the contrary, the latter argues that conceptual fragmentation is a good reason to get rid of the original concept PERCEPTION and replace it with more specific concepts corresponding to specific definitions or characterizations of it. I tackle eliminativism about PERCEPTION in the last chapter. I wanted to show here that conceptual pluralism about PERCEPTION is understandable as a classical step in science, and should not be demonized as a bad consequence of my conceptual engineering. Conceptual pluralism, when it is rigorous, is actually good news: it means that the scientific field is making progress in discovering new mechanisms, new properties and new explanations. Conceptual pluralism is not the sign of disorganization, but the sign that the field produces new data and theories that call for being organized in broader conceptual schemes. My conceptual pluralism is actually evidence that we are not completely limited by our current conceptual apparatus in waiting for the next kuhnian revolution: we have the possibility to stretch our concepts to accommodate an increasing amount of knowledge.

¹⁵Perceptual pluralism is said to be in the "Zeitgeist". Conceptual fragmentation of PERCEPTION can be seen as a cause of it (in addition with academic influence, and probably a lot of other social criteria).

Part III

Joint-in-Nature

Chapter 7

Three ways to carve nature at its joints

The goal of the third part of my dissertation is to investigate the metaphysical assumptions behind the notion of "boundaries of perception". I showed in the first two parts that considerations about concepts may be the place for genuine disagreement in distinguishing perception from the rest of the mind. In this part, I show that metaphysical considerations also have an influence on the boundary question. In this chapter, I present three different ways to conceive the boundary of perception in a realist framework. In the following chapter, I will show how these different metaphysical conceptions of boundaries produce different understanding of an empirical debate such as the cognitive penetrability of perception.

This chapter will unfold as follows. I begin by some preliminary remarks on ontological boundaries. I spend time on the distinction between conceptual and ontological boundaries. I then sketch the landscape of realist positions and introduce the notion of "joint-in-nature" (section 6.1). In the second part, I use the notion of joint-in-nature as a metaphysical tool to grasp the idea of the boundary between perception and cognition. I draw a minimal characterization of it. (section 6.2) In the third part, I develop three different versions of joints-in-nature that fit this minimal characterization (section 6.3).

The conclusion of this chapter is that ontological boundaries of perception are metaphysically ambiguous between (at least) three different conceptions.

7.1 Ontological boundaries: preliminaries

7.1.1 Conceptual vs Ontological boundaries

What are boundaries? Until now, I have leaned on an intuitive understanding of this term. But let us now explore further the nature of boundaries.

Boundaries are what makes concrete the distinction between two things. Two things are different, distinct, when there is one or several dimensions in which these two things can be distinguished. In these dimensions, the two things can be separated by a boundary (like a metaphorical dividing line). In order to characterize a boundary, one need therefore to specify the dimensions and the location of the distinction. For example, there is a bound-

ary between my garden and my neighbor's garden in the spatial dimension – whether or not there is a fence between them. There is a boundary between my child self and my adult self in time, if one considers that some event (like the age of majority) stands for the beginning of adulthood. Dimensions are not restricted to spatio-temporal dimensions. I have a friend who likes ordering her books in her bookshelves by size and color. In her bookshelves, there is a spatial boundary between her blue and green books. In my bookshelves, books are ordered by a complex algorithm taking into account the place I need on my desk to work, the place I don't have in the shelf and the number of books I have to keep to feel safe (I unfortunately do not have any order). However, I can say that there is a boundary between my blue and green books, but on the color dimension only.

With these general considerations, the *conceptual boundary* of a concept A (with its neighbor concepts - or merely non-A) can be characterized as the gap/distance in the conceptual space (constituted by the relevant properties) between the objects that fall under the extension of A (the members of the category), and the objects that do not fall under its extension (or fall under the extension of its neighbor concepts).¹

We can now turn to the ontological notion of boundary. The *ontological boundary* of A is the mind-independent distinction between A and non-A (or its neighbor entities) in reality (vs conceptual space).

These two notions of conceptual and ontological boundaries are independent in principle, so that the conceptual boundaries of PERCEPTION are *not necessarily* related to the ontological boundaries of perception. Conceptual boundaries are relative to the way we represent the world with concepts. Ontological boundaries are what and where they are in virtue of how Nature is ordered, organized, in a mind-independent way. Of course, we can make conceptual boundaries dependent on the organization of Nature if we make our conceptual representations match the organization of Nature. However, this is an indirect dependence, depending on how successful we are in finding ontological boundaries and reflecting them in our representations.

In the previous part, I focused on the variety of concepts for representing perception and the class of perceptual states. It allowed me to argue that there are several concepts, with different extensions, that may be good concepts of perception. This conceptual pluralism for perception leads to the existence of several legitimate *conceptual boundaries* for the class of perceptual states (e.g. I found that two concepts of perception are equi-optimal with a different extension - "expert categories" is excluded in one of them). However, this conceptual pluralism is compatible with a range of ontological views on the boundaries of perception, such as ontological monism, ontological eliminativism and ontological pluralism.

In ontological monism, perception is ontologically unique, so that there is only one ontological boundary between perception and the rest of cognition². This boundary is mind-independent, grounded in the architecture of reality. This position is nevertheless com-

¹Characterization of the "conceptual gap" depends on how concepts are thought to be structured. If concepts are structured as homogeneous kinds sharing essential properties, the conceptual gap is constituted by the shift between having and not having this essential property. If concepts are thought as built by family resemblance, the gap is constituted by the shift (sharp or vague) between "resembling" and "not resembling". If concepts are structured as the ensemble of cases close to some paradigmatic cases, then the gap is reached when the distance becomes "too far".

patible with conceptual pluralism.

In a first version of compatibility, conceptual pluralism constitute a provisional step towards the ultimate goal of finding the unique concept that matches perception and its ontological boundary. This position is associated with a humble realist scientific attitude: science can find the ontological architecture of the world, but this requires extensive scientific investigation. We should investigate all promising hypotheses, until one will proved to be true. Until this day of final verification, let's consider as legitimate all representations that fit equally well our current knowledge.

A second version of compatibility states that ontological boundaries are inaccessible to human knowledge. As our representations will always fail to match the architecture of reality, the best way to capture ontological boundaries is to multiply approximate conceptual boundaries. This pessimistic attitude towards our ability to know natural organisation of things can be traced back to Locke (1690), for which 'real essences of things' are not accessible to us. We are representing things and organizing them thanks to perceptible ideas that we observe, and not according to their real essences. Locke does not deny the existence of real essence, which could be considered as an acceptance of ontological monism (i.e. things have real and natural boundaries). However, as these are unknowable, the only available solution is an approximation through what we can observe (Rosenberg, 1994, for a similar position in philosophy of biology). Like theories in Popper's metaphor, the conceptual division of reality can be thought as a human instrument - a fishnet - designed to "catch the world". Our instrument can be improved but it will always be imperfect. His metaphor is about theories, but nicely applies to conceptual apparatus:

[W]e may succeed in improving our theories – even as instruments: in making nets which are better and better adapted to catch our fish, the real world. Yet they will never be perfect instruments for this purpose. They are rational nets of our own making, and should not be mistaken for a complete representation of the real world in all its aspects; not even if they appear to yield excellent approximations to reality. (Popper, 1995, p.42)

In this second version of compatibility, ontological monism is associated with conceptual pluralism because of the limitation of our cognitive abilities to represent the world.

Another ontological view on the boundaries of perception is ontological eliminativism. In this view, perception does not exist as such in the world (as an entity or class of entities). Consequently, our conceptual representation of it does not have any referent. There is therefore no boundary to draw between perception and the rest of cognition. However, ontological eliminativism is compatible with conceptual pluralism about perception.

One way to associate these two views is to consider that concepts of perception are useful fictions constructed to organize our knowledge and investigations about the mind. In this kind of instrumentalism, nothing prevents that several concepts are more useful than one. One idea could be that our knowledge about the world is complex enough to be hardly captured by a single ensemble of well-organized concepts. More than one conceptual schemes are therefore necessary to grasp the ensemble of nuances contained in our representation of the world.

In a more cynic perspective, ontological eliminativism is compatible with the idea that the nonexistence of a worldly entity in the vicinity of the supposed referent of our concepts of

²Seeing perception as a natural kind of mental states is only one way to be an ontological monist about perception.

perception dooms us to conceptual pluralism until our scientific theories are strong enough to go without it. Concepts of perception are fictions - useless fictions - that should be dismissed by a better conceptual framework. However, they haunt our conceptual space for now in absence of better alternatives.

Finally, ontological pluralism is compatible with conceptual pluralism. Conceptual pluralism may straightforwardly reflect ontological pluralism (the concepts refer to worldly entities such that boundaries of the entities are structurally analogous to boundaries of the concepts). This is however possible to endorse both ontological and conceptual pluralism with boundaries that do not correspond to each other. All reasons suggested previously are here applicable: (1) ontological boundaries are unknowable, (2) ontological boundaries are still under investigation, (3) conceptual boundaries should be useful, which does not lead necessarily to a perfect correspondence to ontological boundaries or (4) conceptual boundaries are there, different from ontological ones, but are unfortunate (and hard to dismiss) legacy of history and common sense.

I hope that it is now clear that conceptual and ontological boundaries are distinct, and that their correspondence requires additional theoretical commitment. In the previous chapter, I questioned the representations of perception and their diversity, what constitutes a step towards a better understanding of the broader debate about the boundaries of perception. My goal was to shed light on some representation-oriented level sources of disagreement between debaters. However, studying the nature of the conceptual boundaries of perception³ will not say anything about its ontological boundaries. Yet, these ontological boundaries are those whose location and nature (sharp, fuzzy, porous, etc.) are important in the debate. In this part, I question the metaphysical assumption behind these alleged ontological boundaries of perception. In what follows, I question the nature of the ontological boundaries of perception and show that there are several different entities that can stand for the perceptual boundary. In the next chapter, I will show that metaphysical views on boundaries can shape reasoning all the way through the interpretation of empirical debates such as the cognitive penetrability of perception.

7.1.2 Realist framework and the joint-in-nature hypothesis

In this chapter, I question the nature of the alleged ontological boundaries of perception. This part takes place in a realist framework, i.e. I take for granted that perception is an entity, several entities or a class of entities, so that perception has one or several ontological boundaries. From now and until the end of this chapter, I will use 'boundary' simpliciter for 'ontological boundary'.

In this section, I show first that realism leaves room for divergence about the nature of perception and its boundaries. Then, I suggest that all realist positions commit to a minimal version of the joint-in-nature hypothesis, i.e. there is a joint-in-nature between perception and the rest of the mind.

³It is possible that some disagreements about the boundaries of perception in general are mixing issues between conceptual boundaries and ontological boundaries, so that analyzing the nature of conceptual boundaries could clarify them (and extract what comes from representational-oriented level and what comes from metaphysical views). I did not identify such mixing in the literature, but would be very interested to read someone who did.

Realism is often thought as committed to naturalkindness of perception. However, critics of perception as natural kinds (and natural kinds in general) make room for alternative realist views about perception. I have shown in the first part of my dissertation that natural kind is a demanding notion, and that some facts about perception put doubt on its naturalkindness. Should we give up on realism? If perception is not a natural kind, should we concede that there is no boundaries to draw between it and the rest of the mind, or at least that its boundary with cognition is not really a boundary?

There are actually alternative realist views, in which distinctness does not require a sharp boundary like between classical natural kinds (in the essentialist tradition).

A famous non-classical account of kinds is the homeostatic property cluster view (Boyd, 1991). According to Boyd, members of a (natural) kind must not share essential properties. It is sufficient that (1) clusters of co-occurring properties are displayed more by the members of the kind than by non-members, and that (2) there are some mechanisms that ground the clustering of these properties, so that the resemblance between the members of a kind is non-accidental. Mechanisms are thought to be dynamic and context-dependent. Consequently, they do not generate identical instances of a kind depending on the subject and the environment. If perception and cognition are thought as boydian kinds, the boundary between them does not need to be sharp anymore. Two (or more) underlying mechanisms may generate mental states that tend to share two different clusters of properties, but depending on circumstances, mental states do not share the properties of paradigmatic members of their kind, and may superficially appears as intermediary or hybrid.

Dupré's promiscuous realism is another non-classical account of kinds (Dupré, 1981, 1995). In Dupré's view, originally applied to biology, kinds are *mind-dependent* organized classes of *mind-independent* clusters of properties. In other words, the boundaries chosen as separating kinds are natural boundaries between different classes of things, but other boundaries only slightly different could be as legitimate in separating genuine different classes of things. Promiscuous realism hold that sameness (and distinctness) relations are numerous in the natural world, and that none of them are more fundamental or more important in grounding kinds. The kinds identified by science, or even in ordinary life, are chosen according to human's interests and use, but are still organizing things into real and natural sameness relations.

These two non-classical views of kinds are nonetheless realist and accept the existence of one or several boundaries between perception and cognition. They enlarge the notion of boundary between kinds. While the boundary between two natural kinds has to be sharp, these non-classical kinds accept a notion of boundary as transitional space, in which a set of features is gradually transformed into another set.

The architectural strategy (chapter 2) suggests another alternative to the naturalkindness of perception, but is still a realist position about the distinctness of perception. The boundary between perception and cognition has to be found in the architecture of the mind in terms of modularity of processes. Proponents of a strong modularity (the encapsulationists) would draw a sharp boundary between perception and other parts of the mind (with no intermediary cases, or this would ruin the notion of encapsulation). However, modularity has been understood at a functional level, without the high demand of informational encapsulation (Carruthers, 2006). In this understanding of modularity,

the distinction between perception qua module and the rest of the mind opens space for intermediary cases and hybrid states and phenomena.

I presented these three alternative views (Boydian kinds, Dupré's kinds and functional modularity) in order to illustrate the diversity of realist positions. They do not constitute an exhaustive list. These views are all different realist position, but they have all in common that they are realist about the boundaries of perception. My goal is to catch this commonalities, in order to dig up the metaphysical assumptions behind realist positions.

And one way to present this boundary realism is to claim that there is a joint-in-nature between perception and the rest of the mind. This is what I call the "joint-in-nature hypothesis".

The next section focuses on this notion of "joint-in-nature", in order to draw a minimal characterization of it. The ultimate goal is to specify the metaphysical commitments behind realist views about the distinctness of perception. I will show in the third section that there are actually (at least) three different ways to understand "joint-in-nature" and therefore to accommodate the joint-in-nature hypothesis.

7.2 Joint-in-nature

7.2.1 Literature

PHAEDRUS: What is the other principle, Socrates?

SOCRATES: The second principle is that of division into species according to the natural formation, where the joint is, not breaking any part as a bad carver might. (Plato, Phaedrus 165e)

Under a plane-tree, by the banks of the Ilissus, Socrates expresses the two powerful tools of the dialectician to Phaedrus: the generalization, or the art of grasping the unifying idea behind scattered notions; and the division of a whole into species. Plato illustrates this latter principle by the metaphor of the good butcher, carving a piece of meat at the joints, while the bad one breaks parts arbitrarily. Dividing things into species in the right way requires therefore to know where the joints of reality are. This powerful metaphor found his way through the history of ideas until contemporary analytic philosophy. One finds today expressions such as "carving nature at its joints" and "joints-in-nature" in contemporary metaphysics (Campbell et al., 2011; Sider, 2011, 2020), philosophy of Science (Franklin-Hall, 2015; Wimsatt, 2007) and philosophy of mind (Allen, 2017; Block, 2014; Firestone and Scholl, 2016; Mandelbaum, 2017; Phillips, 2019; Siegel and Byrne, 2017). The spirit of the original metaphor is kept, but is today thought as talking about the success of science in discovering and identifying the distinct kinds of things that will provide us inferential and explanatory power in understanding the world. But there are actually two distinct uses that should be clarified: a use relative to natural kinds and another to the fundamental structure of reality.

When "joint-carving" divides natural kinds In contemporary philosophy, the metaphor is often used in debates related to natural kinds. "Carving nature at its joints" means organizing our description of the world in natural kinds, by putting together the things

that nature made with the same mold. This interpretation comes from the aristotelian tradition, in which Plato's theory of kind-membership in terms of participation in a Form is transformed in terms of individuals and essence (individuals are organized in kinds in virtue of shared essential attributes). This broad notion of essential attribute is what is common to all classical views of natural kinds (Slater and Borghini, 2011). The book "Carving nature at its joints" (Campbell et al., 2011) in the MIT Press serie "Topics in Contemporary Philosophy" is a good example of this use. The subtitle makes explicit the topic of the book : "Natural Kinds in Metaphysics and Science".

When joints-in-nature are the fundamental structure of reality In 2011, the same year as the publication of the volume edited by Campbell et al., Sider publishes "Writing the book of the world" (Sider, 2011), in which he uses Plato's metaphor to talk about the "fundamental structure" of the world. He argues that the primitive notion in metaphysics should be fundamentality : talks about genuine features, that really exist and ground necessary resemblances between entities can be reformulated as a single question: what is the structure of the world at the fundamental level?. He uses the term fundamentality more or less interchangeably with "joint-carving" and "part of reality's structure" (Sider, 2011, p.5). Importantly, he explicitly states that his project is different from naturalkindness and essentialism (as well as modal metaphysics, ontology and conceptual analysis): "Metaphysics, at bottom, is about the fundamental structure of reality. Not about what's necessarily true. Not about what properties are essential. Not about conceptual analysis. Not about what there is. Structure." (Sider, 2011, p.1).

These two uses of the metaphor and derived expression (natural kind divider or structural joint) are often mixed when they are used in a non-metaphysical context, and the expressions acquire then a weaker meaning, such as "being mind-independent". This mixed and weaker notion can be found in philosophy of Science, especially in the literature about non-classical views of natural kind (natural kind talks without essentialism). As a good example, Franklin-Hall mixes in one sentence "mind-independence", natural kind talk, and a quote from Sider:

On the realist approach, as traditionally construed, the universe possesses a mind-independent natural-kind structure, such that whether a category or classification "carves at the joints" has nothing to do with its place "in human languages, conceptual schemes, biology or anything like that" (?) (Franklin-Hall, 2015, p.927)

In philosophy of mind, the distinction between perception and cognition is said to *carve nature at its joints*, or even to *carve the mind at its joints*. This notion is best understood here as the idea that the distinctness of perception is a fundamental distinction for the study of mind (i.e. in cognitive science). This use borrows Sider's fundamentality (i.e. built in the structure of the world) but confines it to cognitive science (Sider doubts that psychological concepts can be fundamental in his sense). At the same time, the expressions keep a slight flavor of natural kind talk, but without strong commitments to essentialism.

Ned Block is the central figure of the use of "joints-in-nature" in the debate about the boundaries of perception:

Burge, Fodor and I agree that there is a joint in nature between percepts and

concepts, so there is an important theoretical question of how to characterize it. (Block, 2014, p.2)

His forthcoming book is about characterizing "the border between what is fundamental to perception and what is fundamental to cognition" (Block, ms, p.10). According to him, "a joint is a fundamental and explanatory significant difference between the kinds that are separated by the joint." (Block, ms, p.9), but these kinds are not necessarily natural kinds in the classical view as "such a joint is compatible with causation from one to the other and with the existence of borderline cases" (Block, 2014, p.2). Like Sider, he considers that the notion of "joint-in-nature" is primitive, cannot be properly defined and should rather be grasped through examples, e.g. the distinctions between living and non-living (Block, 2014), between animals and plants (Block, 2015), between liquid and solid, hydrogen and helium, lepton and quark (Block, ms). Finally, one crucial idea behind joints in nature is that they "are discovered and not stipulated" (Block, ms., p.14).

Following Block's use, other philosophers of mind quote the metaphor.

there may be no more foundational distinction in cognitive science than that between seeing and thinking. [...] [T]here is a "joint" between perception and cognition to be "carved" by cognitive science. (Firestone and Scholl, 2016, p.17)

there has been renewed support for the idea that there is a joint in nature separating perception from cognition. (Mandelbaum, 2017, p.3)

Mandelbaum characterizes this distinction as a "break between perception and cognition" (ibid, p.25). Finally, in Phillips (2019), one can find a mention of natural kinds but with a quite liberal meaning. "Searching for a border" between perception and cognition is understood as "carving the mind at its joints", which in turn means finding natural kinds behind the concepts of perception and cognition. However, these natural kinds are the "notions of the perceptual and the cognitive [that] have explanatory significance in cognitive science" (Phillips, 2019, p.2), which is not classical way to characterize natural kinds.

Following Block and the other participants to the debate about perception/cognition divide, I use the term joint-in-nature to characterize the boundary of perception, but I want to make it a useful metaphysical tool, instead of a mere metaphor.

7.2.2 A metaphysical tool

I suggest to use the notion of "joint-in-nature" to express the common underlying commitment among realist views, i.e. perception is *really* distinct from the rest of the mind. It is a tool in the sense that it is a primitive notion (I'll provide a tentative characterization in this section, but no definition), and it allows to grasp the metaphysical issue about perception:

- (i) The distinction between a realist and an anti-realist view is whether the boundary of perception is a joint-in-nature or not.
- (ii) The boundary of perception is a joint in nature (the joint in nature hypothesis) is the minimum requirement for all realist positions : realist positions can therefore be differentiated according to the additional commitments they have.

In my use, saying that the boundary of perception is a joint-in-nature is not (only) an empirical matter, but a commitment to a certain view about the mind, i.e. (i) the mind has a mind-independent structure that can be discovered and (ii) the boundaries of perception are metaphysically related to this structure.

The minimal characterization of a joint-in-nature is the following:

A joint-in-nature is a distinction between things that:

- (C1) captures the organization of reality (the "architecture" or "structure")
- (C2) plays an explanatory role in science
- (C3) may or may not allow interactions
- (C4) may or may not accept a certain fuzziness (a non-sharp boundary with intermediary cases)

(C1) expresses the "fundamentality" a la Sider. A joint-in-nature is mind-independent, has to be discovered (instead of invented). When a distinction is a joint-in-nature, it has nothing to do with human affairs. (C2) expresses the idea that joints-in-nature are what science should target to build good theories and explanations. I do not target here a specific kind of explanation. I have in mind a broad notion of explanation as something that increase the understanding of a domain of phenomena. It may consist in discovering causes, in unifying a domain, or in theorizing general laws or mechanisms.

There is no primacy of (C1) or (C2) here, fundamentality and explanatory power of joints-in-nature are two sides of the same coin. (It explains well because it is fundamental, and it is fundamental because it explains well).

Contrary to (C1) and (C2), (C3) and (C4) are dispensable requirement. They remind that there is no constraint on impermeability or sharpness of the distinction. I prefer keeping them explicit in the characterization to mark the difference between a joint-in-nature and a distinction between classical natural kinds.

This characterization is really close to Block's use of joint-in-nature in the literature of the perception/cognition divide. It captures the minimal requirements for realism about the distinctness of perception, so that the joint-in-nature hypothesis is now more specific. Saying that the boundary of perception is a joint-in-nature (the joint-in-nature hypothesis) means that perception is a fundamental and explanatory part of the mind.

As an illustration, let's go through the realist positions mentioned before and check how they comply with the hypothesis⁴:

Natural kind Realism The distinction between two natural kinds is not only fundamental, but also has to be sharp (with no intermediary or hybrid cases). In other words, natural kind realists state that reality has an architecture, and this architecture is organized in neat and sharp classes of things. The explanatory power of natural kinds is associated with their role in inferences. As essential properties are shared by all the members of the same natural kinds, the members are "interchangeable parts of the cosmic

⁴I focus on the perception-cognition boundary in the examples.

machine” (Quine, 1969, p. 20), which allows generalizations. A distinction between two natural kinds like perception and cognition is therefore a joint-in-nature. The particularity is that the distinction does not allow any intermediary cases (C4).

Boydian Realism A distinction between perception and cognition as two boydian kinds is a distinction between kinds grounded in distinct mental mechanisms. These mechanisms are distinct in virtue of the architecture of reality, shaped by natural selection and individual development in a specific environment. Two boydian kinds are therefore distinct in virtue of the architecture of reality. Even if no property are shared by all members of the kind, they shared a common cause, so that they also allow some generalizations, taking into account contextual considerations. Furthermore, intermediary cases are possible as the context will produce variability among products of the same mechanism. Mechanisms can also interact so that the effect of several mechanisms can be combined in a single product. If perception and cognition are boydian kinds, mental states can be both perceptual and cognitive.

Promiscuous Realism In Dupré’s view, all similarities in the world can be used to build kinds with a specific sameness relation. All these kinds are legitimate parts of the architecture of reality, as soon as the similarity is objective. If perception and cognition are Dupré’s kinds, their boundary is therefore a joint-in-nature (among many), even if its salience in cognitive science is due to a scientific choice and not to a natural fundamentality in comparison to other possible kinds. This position underlines the fact that the joint-in-nature hypothesis is compatible with the possibility that the distinction is both grounded in the architecture of reality and mind-dependent (in the sense that it follows human interests). The explanatory power of Dupré’s kinds has to be found in the association between objective similarities and the interest of the classifier. The similarity chosen to classify a domain of phenomena is the one which provide a useful tool for the classifier: for cooks, classifications highlight gustatory or nutritional properties. For scientists, classifications highlight properties that participate in scientific explanations. As a multitude of kinds in the vicinity of a same class of phenomena are legitimate kinds, some particular phenomena may seem intermediary or hybrid.

Modularism Modularists ground the distinction between perception and cognition into the modular architecture of the mind. For the encapsulationists, the joint-in-nature hypothesis is quite obvious, as the architecture of the mind is part of the architecture of reality. For the revisionnists, for which modules have to be understood at a functional level, the acceptance of the joint-in-nature hypothesis depends on the metaphysical status of the functional description of the mind. On the one hand, the notion of function could be understood as the biological function of the components of the mind, shaped by natural processes of selection. Therefore, the distinction between functions is a joint-in-nature. On the other hand, the notion of function can be considered as a scientific artifact, therefore their distinction is not a joint-in-nature, but at the same time, the functional description of the mind is also not a realist description. So this position is harmless for the equivalence between realist positions and the joint-in-nature hypothesis. On the explanatory section, both encapsulationists and revisionists argue that modules, whether architectural or functional, are privileged constituents of explanation of mental behavior, as they are isolated (either in terms of architecture or in terms of function). They work as “building blocks” in a project of exploring complex mental mechanisms. Encapsulationists globally

reject intermediary cases and interactions between perception and cognition qua modules (Firestone and Scholl, 2016). Some proponents of encapsulated perception accept some diachronic interaction though, associated with learning (Mandelbaum, 2017). In the functional understanding of modularity, interaction and intermediary cases are not a problem.

As evoked before, this illustration does not exhaust the way realist positions comply with the joint-in-nature hypothesis, but shows how the hypothesis is flexible enough to accommodate a large range of positions.

The characterization of joint-in-nature in terms of fundamentality and explanatory power refines the notion by clearly putting aside 'natural kind talk', i.e. the use of natural kinds as a theoretical tool to understand realist position. Natural kind realism is only one kind of realism, and is actually not interesting when the issue at stake is the distinctness of perception. We should shift from natural kind talk to joint-in-nature in order to tackle questions such as the perception/cognition divide. Disagreements among realists are infighting, the real challenge in the debate is the existence of new theories of the mind that put in doubt the existence of a distinction. This is here a fight of paradigms. In a nutshell, the notion of joint-in-nature allows us to articulate a disagreement between paradigms, while natural kinds do not make the difference between proponents of non-natural kind realists and eliminativists. According to a natural kind framework, they are all non-believer positions.

One may fear that the notion of joint-in-nature is actually the same as natural kind but twisted in order to answer objections of intermediate cases and interactions. The idea is the following: as a joint-in-nature is fundamental, grounded in the "architecture of reality" and explanatory, it is supposed to capture something like "essences". Talks about "pure perception" could be seen as these "essences". Block says:

Perception should be restricted to pure perception, perception that does not occur as part of a judgment or as part of a working memory representation. Pure perception is perception without any cognitive envelope. (Block, ms, p.7)

I reject this idea of essence for the characterization of joint-in-nature. I do not need that joints-in-nature separate *pure things*. I also reject Block's idea that the notion of joint-in-nature is not expected to be scrutinized:

Although the joint between liquids and solids is explanatorily important, its role in explanations in physics and chemistry is background not foreground in current disputes and I expect that the same is true of the joint between perception and cognition (Block, ms, p.9)

On the contrary, the idea of "joint-in-nature" as the commitment to realism about the distinctness of perception is informative on the metaphysical import that subsequent empirical debates drag around.

7.3 Three joints

My initial project in this chapter was to find the ontological assumptions behind the notion of joint-in-nature, in order to inform empirical debates such as the cognitive penetrability

of perception. However, there is actually not a unique way to meet the minimal characterization of joint-in-nature presented in the previous section. The notion is therefore ambiguous. This would not be a problem if this ambiguity was not transferred to all questions in which the boundary of perception is recruited.

The three understandings are:

1. taxonomical joint-in-nature
2. mereological joint-in-nature
3. focal joint-in-nature

In a first approximation, a taxonomical joint-in-nature is a distinction between two kinds. A mereological joint-in-nature is a distinction between two parts of a whole. Finally, a focal joint-in-nature is a distinction between two parts of a partial description of a complex system.

I show in what follows how these three types of joints-in-nature can be applied to the perceptual boundaries.

7.3.1 Taxonomical joint-in-nature

A taxonomical joint-in-nature is a joint-in-nature between two kinds of entities. These kinds of entities should be mind-independent, i.e. groups of things that are grouped together for mind-independent reasons. This requires that the grouping is made non-arbitrarily and based on criteria that can be observed or measured by different means, so that they can be considered as objective. It does not require however that the kinds are natural kinds (mind-independence alone is a weaker requirement than the three conditions of naturalkindness I presented in chapter 2). As a consequence, Boyd's and Dupré's kinds are also separated by taxonomical joints-in-nature.

I used the adjective *taxonomical* because these joints-in-nature are highlighted in scientific taxonomies, and taxonomical joints-in-nature are actually what is tracked by scientists when they are taxonomizing a domain. Scientific taxonomies are where we can find best examples of taxonomical joints in nature. For example, the periodic table is a good way to find taxonomical joints in nature between kinds of chemical elements. In biological classifications, we find taxonomical joints-in-nature between taxa (at every taxonomical rank). There is a taxonomical joint-in-nature between 'animal kingdom' and 'plant kingdom', between the classes 'mammals' and 'birds', and all the way down to species, e.g. between 'wolves' and 'coyotes'.

Scientific classification gives us a hint on what it takes to find a taxonomical joint-in-nature: the kinds must be distinct *within the same superordinate kind*. There is no sense to say that there is a joint-in-nature between mammals and plants, because they are not on the same rank in biological taxonomies.

Definition Given the previous comments, the definition of a taxonomical joint-in-nature is the following:

Taxonomical joint-in-nature: In a hierarchical taxonomy T, for A and B, two mind-independent kinds of things at the same rank, and S a superordinate mind-independent kind at a higher rank, there is a taxonomical joint-in-nature between A and B if and only if (i) A and B are both subkinds of S, (ii) the distinction between A and B meet the minimal characterization of a joint-in-nature.

Application to the perceptual boundary The boundary between perception and cognition has been thought in the literature as a taxonomical joint-in-nature. For example, Block compares the joint-in-nature between perception and cognition to the distinction between living and non-living, between animals and plants and between hydrogen and helium. These are all taxonomical joints-in-nature.

Saying that there is a taxonomical joint-in-nature between perception and cognition is saying that perception and cognition are two mind-independent kinds of mental states, both belonging to a same superordinate mind-independent kind, and distinct in a fundamental way, such that this distinction participates in explanation in cognitive science. They must not be sharply divided, intermediary or hybrid mental states may exist, and they may be combined in ways that allow interactions between them in mental activity.

This superordinate mind-independent kind may be the ensemble of mental states, or restricted to a subkind of mental states such as representational mental states. It does not matter which superordinate kind is chosen, what matters is its existence, in order to secure the idea that a taxonomical joint-in-nature is a fundamental distinction between resembling things.

In order to defend the thesis according to which there is a taxonomical joint-in-nature between perception and cognition, one has to defend two claims: (i) the distinction between the kind of perceptual states and the kind of cognitive states are belonging to the same superordinate kind of things, and (ii) within this superordinate kind, their differences are fundamental and explanatory.

The goal of this chapter is not to argue in favor of the existence of a taxonomical joint-in-nature between perception and cognition, but to show that this thesis hosts a specific metaphysical view about the mind and its boundaries (which will have consequence on other debates).

Here, proponents of a taxonomical joint-in-nature between perception and the rest of the mind (or more restrictively between perception and cognition) are assuming that individual mental states are *entities of the mind, constituents of mental activity*, and that their classification is part of cognitive science research⁵.

These assumptions are shared by the whole camp of taxonomical joint-in-nature, independently of their disagreements on where and how to draw the perceptual boundaries. Likewise, biological taxonomists are not putting in doubt that their entities in need of classification are individual organisms. Chemical taxonomists are not putting in doubt either that their classified entities are atoms. Here, the commitments are: (i) taxonomizing

⁵Whether this classification is quite sharp (like the classification of chemical elements) or debatable and pluralistic (like biological classifications) is another debate.

individual mental states would provide a fundamental and explanatory tool to cognitive science and (ii) this (mental) taxonomy includes perception and cognition (qua kinds of mental states) as taxa.

In this perspective, a taxonomical joint-in-nature is fundamental for cognitive science in the sense that it organizes *fundamental constituents* of mental life. These constituents (like building blocks) combine to give rise to complex mental activity and behavior. Explaining mental activity and behavior by finding their fundamental constituents and their combinatorial rules or mechanisms seems to be a good explanatory strategy. This could be compared to explaining a molecule by identifying the types of atoms it contains and how these atoms are combined.

Being fundamental and explanatory does not prevent from allowing interaction between states and hybrid states. The distinction between perceptual states and cognitive states may be a difference in degree, but sufficient to ground fundamental homogeneity within a kind, and fundamental difference with the other. Furthermore, combinatorial rules may be such that perceptual and cognitive states may combine in hybrid states, with influence from perceptual to cognitive content and vice-versa.

7.3.2 Mereological joint-in-nature

A mereological joint-in-nature is a distinction between two parts of a whole. Paradigmatic mereological joints-in-nature could be the distinction between body parts, or the distinction between the different parts of a cell. For example, in my body, there is a mereological joint-in-nature between my left arm and my torso. By putting together in the right way all the parts that compose my body, it makes my body as a whole. In a cell, very roughly, one may distinguish the membrane, the nucleus and the cytoplasm. Putting them together in the right way will give you a cell as a whole.

The distinction between the parts may be based on spatial, structural or functional criteria. The crucial point is that the decomposition of the whole in parts must be fundamental and explanatory. The fundamentality of a decomposition is the idea that the decomposition literally "carve nature at its joints" like a butcher. Fundamentality of a decomposition can be tracked with the two following criteria: (1) the parts must be descriptively independent from the whole, and (2) the arrangement of the parts to constitute the whole must be simple.

As an upshot, the whole entity can be described in terms of these parts and their arrangements (functional, structural or spatial). The independence of the parts is not (necessarily) a strong notion of independence (e.g. being functional in isolation), but one of description, such that the full description of the whole is equivalent to the full description of the parts in addition with their arrangement.

The explanatory power of the decomposition comes from the fact that the complexity of the decomposed whole is broke down into smaller parts, arranged in simple relations. Here is an easy example of the explanatory power of such a decomposition. In order to find the area of a trapezoid, it is easier to break down the shape into simpler shapes, namely a rectangle and two right triangles, arranged side by side. Finding the area of a rectangle and right triangles is easier because they are simpler shapes, and their arrangement is simple (side by side) so that to find the area of the trapezoid, one has to only add the area of the

simpler shapes. I do not argue that this example displays mereological joints-in-nature, but it is an illustration of how decomposing a complex whole into simply arranged simpler parts participates in explanation about the whole.

Definition *Mereological joint-in-nature*: for A and B, two parts of a whole W, there is a mereological joint-in-nature between A and B if and only if (i) A and B are parts in a fundamental decomposition of W, (ii) A and B are adjacent.

Here the notion of "adjacent" depends on the kind of arrangement is required for the composition of the parts into the whole. I choose to understand it vaguely as a kind of neighborhood in order to be flexible enough to accommodate a range of different arrangements. If this notion is more problematic than what I expect, I propose this following formulation of a mereological joint-in-nature:

Mereological joint-in-nature v.2 : for A, a part of a whole W, the boundary of A is a mereological joint-in-nature (between A and non-A in W) if and only if A belongs to a fundamental decomposition of W.

In terms of the requirements of the minimal characterization, they follow from this definition of a mereological joint-in-nature. The fundamentality and explanatory power of the distinction between A and B is derived from the fundamentality of the decomposition to which they belong.

Application to the perceptual boundary Applying the notion of mereological joint-in-nature to the perceptual boundary means that perception (and cognition if we focus of the perception-cognition border) is a part of the mind in a fundamental decomposition of it.

This decomposition should describe the whole mind through independent descriptions of its parts and their relation. It means that there is a decomposition of the mind, such that perception and cognition, in addition with other parts, (i) can be independently fully described, and (ii) whose relations are simple.

A problem comes from the notion of "simplicity" here. In the previous geometrical example, the simplicity of relations is intuitive, as the parts were put side by side in a spatial arrangement. However, a spatial model of decomposition does not work well for the mind. A spatial decomposition of the mind with "perception" and "cognition" as parts in terms of brain area seems old fashioned today and recruited erroneous images of the mind like Gall's phrenology in which each ability of the mind has a specific associated brain area. Best candidates for decomposing the mind with perception and cognition as parts are instead *functional* or *structural* decompositions. The mind is here understood as a system composed of processes or mechanisms. Yet, what is a simple arrangement of sub-processes or sub-mechanisms? It is less intuitive than in spatial cases.

I suggest to use the notion of decomposability of complex systems (Simon, 1969). Complex systems are more or less decomposable into subsystems aggregated together. If you take an industrial production line of balls of wool, the whole system is decomposable into subsystems that are simply aggregated in a series: the product of a subsystem is the input of the next one. In a non-decomposable systems (or integrated system), all the subparts

of the system are interacting with other parts such that it is impossible to describe the system as an aggregation of independent subsystems.

Perfect aggregation is the simplest way to arrange subsystems. Subsystems are aggregated (versus integrated) into a complex system when the behavior of each subsystem is independent from the behavior of the other subsystems. Simon (1969) develops the idea that systems that evolved naturally (e.g. minds) are near-decomposable.

At least some kinds of hierarchic systems can be approximated successfully as nearly decomposable systems. The main theoretical findings from the approach can be summed up in two propositions: (1) in a nearly decomposable system the short-run behaviour of each of the component subsystems is approximately independent of the short-run behaviour of the other components; (2) in the long run the behaviour of any one of the components depends in only an aggregate way on the behaviour of the other components. (Simon, 1969, p.198)

If the mind is a near-decomposable system, the interaction among processes or states belonging to the same subsystem are more numerous and significant than the interaction between processes or states belonging to different subsystems, so that the interactions across subsystems are explanatorily negligible. Consequently, even if subsystems cannot be entirely studied in isolation, describing the system as a set of weakly interacting subsystems is a good approximation.

To sum up, I consider that the arrangement of the parts of the mind is simple enough (for the decomposition to be fundamental) if the interactions across subsystems are negligible in terms of explanatory power. Note that the interactions are possible, so that mereological joint in nature still meet the requirement of enabling interactions and intermediary cases. The important commitment of this metaphysical view on the mind is the near-decomposability of the mind at the level of macro-abilities such as perception. Whether the mind is actually near-decomposable is debatable (Schierwagen, 2009, 2012), but this is beyond the scope of this chapter.

The existence of a mereological joint-in-nature between perception and cognition (or the rest of the mind) is well represented in the literature by modularists, and more generally by the proponents of the architectural strategy to draw the borders of perception (Firestone and Scholl, 2016; Fodor, 1983; Mandelbaum, 2017; Quilty-Dunn, 2019)). Most of them argue in favor of a structural decomposition of the mind: modularity of perception is hard-wired in the architecture of the mind. A version of functional modularism is also possible (Ogilvie and Carruthers, 2016). Within the proponents of a mereological joints-in-nature in the mind, there are still disagreements, e.g. on the processes included within the perceptual module, on the interactions with the others modules, etc.

Here, what interest me is that the decomposition into modules is thought as fundamental and explanatory, while leaving room for overlaps, gaps and interactions because the "sharpness" of the decomposition is an approximation.

Partial conclusion Taxonomical and mereological joints-in-nature are sometimes combined, especially in Block's view about perception (2014, 2016, ms). It is appealing indeed to think that the organization into distinct mental processes matches the organization into kinds of mental states, with the idea that a same kind of mental state is produced by

the same subpart of the complex system. However, this association of taxonomical and mereological joint-in-nature for perception is an additional commitment, and these types of distinction are in principle independent.

The distinction could be a mereological joint-in-nature without any commitment towards the existence of a taxonomical joint-in-nature. For example, Mandelbaum (2017) argues that perception should be distinguished as a module, but some mental states included within the module are conceptual like cognitive states (at a specific level of abstraction). Similarly, Quilty-Dunn (2019) argues that perceptual states do not share any specific format of representation. Consequently, there may be no fundamental similarity between states included within the perceptual module, so that the mereological joint-in-nature does not correspond to any taxonomical joint-in-nature. The other way around, the perceptual joint-in-nature can be taxonomical without being mereological. Imagine a taxonomy of mental states based on phenomenological properties (supposedly objective, like in Nagel's project of "objective phenomenology" (1974)), in which perceptual phenomenology is the criteria of membership into the class of perceptual states. Depending on the characterization of this phenomenology, one may accept that different mental processes, different subparts of the complex mental system, can produce perceptual states.

The independence of taxonomical and mereological joints-in-nature shows how an intuitive idea such as the correspondence between subparts of the mind and kinds of mental states is much stronger than expected in terms of metaphysical commitments. Block's view is actually (and counter-intuitively) stronger than Mandelbaum's or Quilty-Dunn's ones.

In the next subsection, I present a third realist possibility, i.e. focal joints-in-nature. My goal is to show that taxonomical and mereological joints-in-nature do not exhaust the possible understandings of the minimal requirements for a realist position. The existence of a taxonomical joint-in-nature between perception and cognition is roughly the metaphysical commitment of the proponents of the representational strategy (see chap.1) and the existence of a mereological joint-in-nature is the metaphysical commitment of the proponents of the architectural strategy. These two strategies are often thought as exhausting the realist possibilities, in absence of other alternatives. For example, Quilty-Dunn (2019) used the failure of the representational strategy as an indirect argument in favor of the architectural strategy. He clarifies in a footnote that:

Perceptual pluralism [in terms of format] is compatible with non-architectural approaches (whatever they might look like), and even with eliminativist approaches to the perception–cognition border (Shea, 2014; Lupyan, 2015). However, the arguments in this paper presuppose a border and use experimental evidence to distinguish perceptual processes from cognitive ones. (Quilty-Dunn, 2019, p.3)

The next section is thought to shed at least a dim light on what might look like a non-architectural and non-representational approach of the perception-cognition border.

7.3.3 Focal joint-in-nature

There is a focal joint-in-nature between two objects when they are distinguished within a *flattened* description of a system or a group, i.e. a description taken into account only a set of relevant variables. The term "focal" comes from the image of the focus of the eye lens,

that makes us see clearly at a specific distance. A metaphoric lens could make us see some structures more clearly when some properties within a system are not taken into account. Some structure, hidden in a complex system, can be put in light by ignoring some variables.

Ishihara's test for color deficiency gives us a good example of structures revealed by a *flattened* description. The test is composed of 38 patches of colored points, organized such that some digits are visible within the patch under some specific color vision conditions (i.e. good color vision), and another digit (or nothing) visible under a different color vision condition (i.e. red-green color blindness). In some cases, the design is made to reveal a digit to color-blind people and nothing (or only a path) to people with normal color vision. This is the case of the following plate (see figure 6.1) in which only red-green color-blind people can perceive the digit 5.




Figure 7.1: Ishihara's test – Plate 14 (source : colour-blindness.com)

Red-green color-blind people hardly distinguish between reds, greens, browns and oranges, so that their ability to perceive the digit 5 is due to a lack of discrimination between these hues. In terms of description of the patch, the differences between these hues are not taken into account in the characterization of the cloud of point. There is therefore a structure within this cloud of points, only salient if some variables are turned off.

Ishihara's test is designed for nesting this hidden structure. The idea of focal joint-in-nature is that similarly hidden structures may be naturally present in complex systems like organisms, minds, or groups of individuals. These structures provide ways to organize an object into parts (like the distinction "digit 5/ ground" in Ishihara's plate), but this organization is tied to the set of variables chosen to reveal this structure. The set of variables constitutes the lens through which the object is contemplated, resulting in a *flattened* description. This is what I call a section of an object.

A *section* of an object is an incomplete description of it, involving a set of variables used to characterize and organize it into parts by ignoring the differences that are not tied to the

chosen set of variables. The description is *partial*, as some properties are not represented, but *systematic* in covering the whole object. The chosen set of variables is not arbitrary but must secure the fact that at least one available partition of the object through the section is fundamental⁶ and explanatory.

The fundamentality of a partition can be tracked by the same criteria than the fundamentality of a decomposition in the case of mereological joint-in-nature, i.e. (1) the parts must be descriptively independent from the whole section, and (2) the arrangement of the parts to constitute the section must be simple. The partition must be explanatory in the sense that the parts and their interaction provide a better understanding of the behavior of the whole studied object.

The term 'section' is taken from Wimsatt (1994), and is similar to his concept of 'theoretical perspective'. Wimsatt's perspectives are also systematic and partial account of a domain of phenomena, meant to solve a class of problem (Wimsatt, 1994, p.259-260). I chose the name *section* as a tribute to one of his comment in his 1994's article in which he says regarding perspectives:

If I were to rename them now (as I probably should), I would call them sections—short for cross sections [...]—views chosen by architects, engineers, and anatomists to give particularly revealing views of aspects of their complex structures, views which can cross-cut one another in various ways, and at various angles, views which are individually recognized as incomplete, views which may be specialized for or better for representing or for solving different problems, and views which, like perspectives, contain information not only individually, but also in how they articulate. (Wimsatt, 1994, p.265).

Sections and perspectives are really close concepts. However, contrary to Wimsatt, I do not include the concept of section into a global ontology of complex system.

Let's take a biological example of focal joint-in-nature. In comparative morphology, most properties taken into account to describe the head of an organism are tissue-level and organ-level static structures such as muscles and bones. As a result, the skull is a specific part of the head, partitioned again into separate bones distinguished by the sutures. One may say that there is a joint-in-nature between the bone in the facial region, and the posterior cranial bone. Within another view, in functional morphology, some head muscles and bones can be integrated into the same process, when the partition into processes is focused on explaining feeding behavior (Winther, 2006). One may also say that there is a joint-in-nature between these processes. Joints-in-nature in comparative morphology and functional morphology are good candidates for being focal joints-in-nature. The partitions to which they belong are produced through partial description of the head of an organism, i.e. through sections of the head. In comparative morphology, dynamical aspects of the organisms are not taken into account, while in functional morphology, the distinction in static structure like bones and muscles is not taken into account, to the benefit of a processual structure linked to a specific behavior of the whole organism. Both sections and associated partitions are explanatory in the sense that they participate in a better understanding of an organism. While comparative morphology may be seen as participating in determining the origin of organisms (through ontogenetic history or ecological adaptation

⁶The "fundamentality" of a partition is relative to the section, which is maybe a twisted notion of fundamentality. I prefer keeping it to make explicit the analogy between a fundamental partition of a section and a fundamental decomposition of an object in the case of mereological joint-in-nature.

scenarii), functional morphology is explicitly organized to explain well-identified behaviors of organisms such as feeding, swimming, walking, etc.

Determining which requirements a section should meet to secure the fundamentality and explanatory power of its partition is certainly a complex matter and need careful theorizing. I take for granted here that sections are proper scientific theoretical tools, used in scientific practice and able to generate fundamental and explanatory partitions of objects.

Definition *Focal joint-in-nature*: for an object O, a section S of O and A, a part of S, the boundary of S is a focal joint-in-nature if and only if A is a part of a fundamental partition of S.

Application to the perceptual boundary The boundary of perception is a focal joint-in-nature if perception is part of a fundamental partition of a section of the mind. The parts may be processes, mechanisms or physical structures.

To my knowledge, this kind of view is not endorsed explicitly by any researcher in the perception/cognition divide. However, in the broader domain of philosophy of mind and psychology, both Maar (1982) and Chalmers (2011a, 2012) may be seen as proponents of the use of sections in the study of the mind.

Maar's three levels of description of (mental) information-processing systems can be said to be three different sections of the mind. He says:

Almost never can a complex system of any kind be understood as a simple extrapolation from the properties of its elementary components. [...] If one hopes to achieve a full understanding of a system as complicated as a nervous system, a developing embryo, a set of metabolic pathways, a bottle of gas, or even a large computer program, then one must be prepared to contemplate different kinds of explanation at different levels of description that are linked, at least in principle, into a cohesive whole, even if linking the levels in complete detail is impractical. (Maar, 1982, p.245)

He proposes the three levels of description of information-processing systems like vision:

1. the computational level: the goals of the computation displayed by the system (or what the system do).
2. the algorithm level: the inputs, outputs and transformations performed by the computation (or how the system do what it does).
3. the implementation level: the physical realization of the computation and algorithm

All levels of descriptions are systematic but partial descriptions of the studied object. Each level takes into account some variables of the complex system, but leaves some aside. For example, the kind of algorithms or physical structures used for the computations are not specified at the computational level. The other way around, the implementation level will not make the computation salient in its description – the computation would be "hidden"

like the 5 in Ishihara's test.

Each level of description brings a different level of explanation. Some problems can be tackled within a level but the full understanding of a system is possible only through multiple levels and their articulation. The three levels are required for understanding the information-processing systems of the mind, but the computational level is crucial to the understanding of vision, so that it would probably be the one used to distinguished vision from the rest of the mind. When vision is characterized, the computational description is favored:

Vision is a process that produces from images of the external world a description that is useful to the viewer and not cluttered with irrelevant information.
(Maar, 1982, p.254)

One could say that the distinction between vision and other processes of the mind is a focal joint-in-nature (in a "computational section" of the mind).

The specificity of Maar's levels compared to my general view on sections is that they are hierarchical in the sense that a description at a higher-level is multiple realizable in lower-levels (e.g. descriptions at the computational level are multiply realizable at the algorithm and implementation levels, and descriptions at the algorithm level is multiply realizable at the implementation level).

Chalmers (2011, 2012) could also be seen as a proponent of focal joint-in-nature between perception and the rest of the mind. He argues that mental properties (perception included) are *organizational invariants* that should be understood at the level of computations. An organizational invariant is invariant with respect to the causal organization of the system. It means that they do not change as long as the causal relation between them remains intact. All the other changes (such as changes in the environment, in the physical realization of the system, etc.) do not impact them. This causal organization of mental systems mirrors the formal state-transitional structure of a computation. The computation is then a partial description of the mind because it does not take into account variables outside of the causal organization. Furthermore, it produces a division of the system into parts. This computation is therefore a partial abstract description of the system and produces a partition of it. I can therefore be considered as a section of the mind.

Like sections, there are numerous computations available for a studied system:

The right way to think about things is to conceive of physical systems as having multiple causal structure, corresponding to different ways of grouping states of the system into state-types.(Chalmers, 2012, p.16)

In Chalmers' view, distinction between psychological properties such as perception and beliefs can be seen as a focal joint-in-nature.

Maar and Chalmers gives us two examples of sections in which perception (vision for Maar) is one part of the mind. Their two sections involve a computational partition of the mind, but one may imagine other kinds of sections with different criteria of individuation (at the neural level for example).

The particularity of a section of the mind is its *systematic but incomplete* description of it, such that it explains some mental behaviors, without any ambition of being exhaustive. Contrary to the mereological understanding, the model is incomplete and can cohabit with other partial models of the mind, describing other patterns of the complex system. Some problems are solved from within a section, and others require resources from other sections. Explanations may be also found in the articulation between sections. Multiple sections can be applied to the system, helping understand different aspects of the system. Especially different sections can cut up systems differently, and these partitions are not easily comparable.

For example, the traditional hierarchical understanding of perception and cognition could cohabit with a bayesian model of it, if both approaches are accepted as proper sections of the mind.

One may object that, in these conditions, a focal joint-in-nature is not really a joint-in-nature at the same degree than the preceding two versions (taxonomical and mereological) because the sections are not mapped onto *the* fundamental structure of the system (contrary to mereological joints-in-nature). Yet, fundamentality is one of the requirements of the minimal characterization of joints-in-nature. I maintain that both requirements (fundamentality and explanatory power) are met by focal joints-in-nature. They also give a nice understanding of interactions between parts and fuzziness of boundaries.

The fact that descriptions in a section are incomplete does not prevent from describing real structures of the system. Let's take again the intuitive example of Ishihara's test: the shape of the digit 5 is really within the cloud of points. Similarly, in Maar's and Chalmers' views, computations are really performed by the system. At best, one may accept that a focal joint-in-nature is fundamental *in an incomplete manner*, while a mereological joint-in-nature is *all that is needed* for a fundamental description of the system. However, it is not a question of degrees, but of completeness.

Sections are explicitly made to explain a specific range of problems. By narrowing a system to a section, a complex system is simplified, nested patterns appear and explain specific behaviors or properties of the system. However, some problems fall outside a section, and others are only partly explained with resources of a specific section. In these cases, a section and its partition only reveal partial causes, or partial mechanisms. Furthermore, understanding how several different section cross-cut each other may help unified the understanding of the mind.

Focal joints-in-nature accept strong interactions, contrary to mereological joints-in-nature that accept only interactions negligible for the near-decomposability of the system. Interactions across a focal joints-in-nature are those that are not taken into account within the specific section to which the focal joint-in-nature belongs. Focal joints-in-nature also nicely accommodate intermediary cases. The variables used to produce a section may vary in degrees, and some parts of the system (or states) may be differentiated sharply within a section (with a specific range of properties) but they may still resemble each other strongly with respect to other properties.

7.4 Conclusion

In this chapter, I suggested a new conceptual framework to organize the diversity of realist positions in favor of a perceptual boundary: they all hold the joint-in-nature hypothesis, i.e. there is a joint-in-nature between perception and the rest of the mind.

I showed then that there are at least three different ways for the perceptual boundary to be a joint-in-nature. These different kinds of joints-in-nature are actually bound to different conceptions of the mind as object of research. Is it a world of states that we should taxonomize like animal species? Is it an organ or a system that we should decompose into simpler subparts? Or is it a system that we can break up into many partial sections in order to understand its complexity at many different levels? In each of these views, the perceptual boundary is understood differently.

In the next chapter, I will show the influence that the metaphysical conception of the perceptual boundary can have on an empirical debate such as the cognitive penetrability of perception. This will show how metaphysical assumptions about a studied object – here, the mind – shape broader debates, and especially debates for which empirical data is crucial.

Here, I merely presented the three kinds of joints-in-nature applied to perceptual boundaries and do not take side. I actually favor the third position - the focal framework - but a proper defense of this view against the others will wait another work.

Chapter 8

Consequences on penetrability

The goal of this chapter is to show that the metaphysical conception about the perceptual border, i.e. whether it is a taxonomical, mereological or focal joint-in-nature, shapes the way empirical debates are understood and interpreted. An empirical debate consists in a scientific question for which empirical data are crucial, and appear as the ultimate arbitrator between the opposite positions in the debate. This chapter is built as a case study, focused on the *cognitive penetrability of perception* (CPP).

The cognitive penetrability of perception is roughly the thesis according to which what we think affects how we perceive the world. More precisely, the thesis holds that

cognitive states such as beliefs, desires, and possibly other states can causally influence perceptual processing in such a way that they end up determining subjects' perceptual contents or experiences. (Raftopoulos and Zeimbekis, 2015, p.1)

The opposite position in the debate supports that perception is impenetrable, so that there is no *direct* influence from cognitive states on perception (Fodor, 1983; Pylyshyn, 1999; Firestone and Scholl, 2016).¹

Whether or not perception is penetrable is a matter of empirical inquiry: either scientists will discover direct influence from cognition to perception or they will discover mechanisms that secure the absence of influence. Empirical studies should however take the boundary question into consideration: the scope of perception influences the interpretation of experimental data. For example, experiments on perceptual categorization show that we are able to categorize objects immediately when presented with them (Potter et al., 2014, a picture of flower is categorized as a flower in 13 ms). This could be interpreted as a case in which our conceptual knowledge on objects influences perception, but this could also be understood as the manifestation of a local perceptual process of categorization (Mandelbaum, 2017). In this latter case, the results do not support cognitive penetrability of perception, but inform on a specific process within the perceptual module. Therefore, the interpretation of Potter's results depends on whether perception includes categorization processes. This shows that CPP experimental data depends on the *location* of the boundaries of perception.

Here, I want to show that the debate is also influenced by the *nature* of the perceptual boundary (and not only its location)². Depending on whether the perceptual border is

¹For a more detailed introduction to the debate, see the introduction of the dissertation, section 1.2

thought as a taxonomical, mereological or focal joint-in-nature, the positions in CPP (both impenetrability and penetrability) are understood differently. It influences how experimental data are interpreted, and changes the kind of revision required to resist the CPP thesis.

CPP is also interesting as a case study here because it has been argued that the multiplication of results in favor of CPP puts pressure on the distinction between perception and cognition itself. If so-called cognitive and perceptual processes interact at all levels, there would be no good reason anymore to distinguish them as two distinct parts of the mind. This is roughly the reasoning of eliminativists. They consider PERCEPTION and COGNITION to be misconceptions of folk psychology, which do not refer to any actual difference in types of mental states, mental processes or mental subsystems (Clark, 2013; Lupyan, 2015; Shea, 2014). I will show in this chapter that depending on its nature, the joint-in-nature between perception and cognition is more or less resistant to eliminativism: when taxonomical and focal joints-in-nature do not easily lead to eliminativism, the mereological border struggles with it.

This chapter is composed of three parts. I begin by presenting the positions of CPP according to each metaphysical view about the perceptual border. I will show that these positions do not interpret in the same way experimental data in favor of CPP. In a second part, I explore how the different metaphysical views can accommodate CPP data to maintain impenetrability. Finally, the last part is dedicated to explore the resistance of each metaphysical view to eliminativism, given CPP data.

Caveat: Some authors reject data in favor of CPP as being methodologically flawed and therefore unreliable (Firestone and Scholl, 2016). For my purpose here, let's suppose that at least a part of the data is experimentally reliable.

8.1 Penetrability of joints-in-nature

In a *taxonomical framework*, perception and cognition consist in two kinds of mental states. They are two kinds of "building blocks" of mental life. These building blocks are associated with combinatory rules (or production rules), which describe how these blocks work together in the mind. Traditionally, the way perceptual states and cognitive states are associated is the following: perceptual states generate perceptual judgments (cognitive states based on the content of perceptual states), which can enter the global economy of cognitive states in cognitive reasoning.

This traditional view holds the existence of some production rules involving perceptual and cognitive states:

1. Perceptual states participate in the production of some cognitive states (they produce perceptual judgments)³
2. Cognitive states participate in the production of other cognitive states (they can be combined together in reasoning and produce new beliefs and knowledge)

²Here, the perceptual boundary can be restricted to the perception/cognition border, as the CPP debate tackles the alleged influence of *cognition* to perception.

3. Some perceptual states participate in the production of other perceptual states. This rule is more controversial, but harmless with respect to CPP.

The important point with the traditional view is that one production rule is banned, namely the rule in which *some cognitive states participate in the production of some perceptual states*. Experimental data in favor of CPP supports the existence of this new production rule. Therefore, proponents of a taxonomical joint-in-nature between perception and cognition understand CPP data as evidence for an additional production rule in the global economy of mental states.

This understanding does not apply in a mereological framework. If the perception/cognition divide is a *mereological joint-in-nature*, it means that perception and cognition are two quasi-independent parts in a decomposition of the mind. The interaction between the two parts is supposed to be negligible for the understanding of mental activity as a whole, and their arrangement to compose the complex system of the mind should be simple.

The arrangement is usually considered simple between perception and cognition as they are supposed to work more or less in a series: the output of perception is an input of cognition.

Accepting penetrability within a mereological framework means that the cognitive subsystem influences the perceptual subsystem in a way that preserves the *simplicity of the arrangement*. It is really important to keep this simplicity, because if the arrangement between the parts is too complicated, the decomposition loses explanatory power (as the goal of a decomposition is to break down a complex system into smaller systems simply arranged). A metaphor that works well for mereological joints-in-nature is the distinction between machines in a factory production line. Imagine that someone wants to describe the production line to a newcomer in the factory. She will try to produce a description involving a mereological decomposition of the line. If two machines are not associated in a series, but share two or three contact points (like a re-entrant loop for quality checking), it is still simple enough to describe the two machines in isolation with their contact points. But if their contacts are too complex or numerous, it would probably be easier (and more easy to understand) to describe the two machines as a single unit. Similarly, if the two subsystems perception and cognition are arranged in a complex way, there is no epistemological advantage anymore to consider them as distinct (or in other words, they are not near-decomposable anymore).

Data in favor of CPP show that the interface between perception and cognition is more complex than expected. One solution, in order to keep the simplicity of arrangement, would be to find a unified account of all penetration cases in terms of a general influence of cognition on perception. One could imagine that all cases of cognitive penetration on vision are made through integration of mental imagery (influenced by other cognitive states) to the output of early perceptual processes, in order to generate perceptual experience (Macpherson, 2012, for a detailed account of a similar view of penetration applied to color vision.). It would secure the decomposition and the simplicity of arrangement, while accepting the CPP experimental results.

³The process that leads to the formation of a perceptual judgment is debated. Some authors argue that the content of the perceptual state is merely endorsed to produce a cognitive states with the same content. Others involve a process of conceptualization.

Finally, if the distinction between perception and cognition is a *focal joint-in-nature*, it means that perception and cognition are two parts of the mind within a specific section of the mind. A section is built through a set of variables (voluntarily excluding some dimensions of variability). The understanding of the penetrability of perception depends on whether CPP is *salient* within the section in which perception and cognition are parts.

Saying that CPP is salient in a section means that (1) the variables used to build the section are also used to account for cognitive penetrations, and that (2) the problems tackled by the section are directly connected with cognitive influence (CPP modifies the explanations given to these problems within the section). Let's take an example. Consider the section corresponding to the computational level in Maar's theory. If cognitive penetrations change the computations performed by the mind, then CPP is salient at the computational level. It means also that the computational explanation given to some mental behaviors should be revised in order to take into account CPP data. In this case, proponents of focal joint-in-nature should explore the same solutions as proponents of mereological joint-in-nature: they should find a way to modify the arrangement between the parts of the section to give space for more interactions, while keeping a relative simplicity.

However, contrary to the mereological framework, there is another possibility for the focal one, which is that CPP is not salient within the section in which perception and cognition are parts (let's call it S1). I continue with Maar's example. CPP can be salient in the algorithmic section or implementation section in Maar's theory (corresponding respectively to the algorithmic and implementation levels), and does not change anything in terms of computations. In this case, CPP data leave untouched the computational section. In other words, CPP is an extra-sectional problem.

Rolfs and Dambacher (2016) advocate a similar explanation, in which two different sections account for the distinction between perception and cognition on one hand, and CPP on the other hand:

The divide between perception and cognition is hard to maintain at the physiological level. [...] In fact, anatomy tells us that the only substrates of visual processing that are not targeted by top-down feedback are in the retina, leaving little room to distinguish vision and cognition at this level of description. We propose instead that perception is separated from cognition by its function. [...] At this functional level, we argue, the distinction between perception and cognition works. (Rolfs and Dambacher, 2016, p.46)

For Rolfs and Dambacher (2016), there are (at least) two sections of the mind:

1. A physiological section in which interaction between visual and high-level processes prevents any clear demarcation between perception and cognition, but accounts nicely for CPP.
2. A functional section in which perception and cognition are distinct (individuated with their distinct functional roles).

If CPP does not influence the functional role of cognition and perception, then CPP is not part of the section and CPP data do not require any accommodation from the proponents

of this focal joint-in-nature between perception and cognition.

A possible consequence is that S1 loses some importance in the general understanding of the mind, pointing toward additional sections required to account for new experimental data. It may also reduce the explanatory scope of S1.

For example, if S1 was thought to exhaustively explain interactions between the content of perceptual and cognitive states, CPP (accounted for in another section) shows that S1 leaves something out in terms of explanation. Explaining content production will in this case be a cross-sectional problem. S1 was wrongly thought to explain more than it actually could. Recognizing that a specific problem cannot be tackled by a single section is not enough to dismiss a section as illegitimate, as long as other phenomena (e.g. belief acquisition, epistemic justification, etc.) are well explained with it.

The focal framework, contrary to the others, accepts that CPP should be explained (at least partly) through another section of the mind. Examples of such sections could be a physiological section (like in Rolfs and Dambacher, 2016), or a bayesian section, in which perception and cognition does not exist as such, but are replaced by bottom-up (sensory information) and top-down processes (expectation, prior).

It is unfortunate in these cases (and confusing) that the debate is called the *cognitive penetrability of perception*, as the penetration does not occur at the level of salience of cognition and perception.

Sections are incomplete descriptions of a complex system, it is therefore possible – and expected – that some data cannot be explained well within a single section. The power of the focal framework comes from the idea that several sections, that do not map well onto each other (like a traditional hierarchical organization of perception and cognition on one side, and a bayesian model of perception on the other side), are not competing against each other, but are complementary.

8.2 Maintaining Impenetrability

There are reasons why people want to maintain impenetrability between perception and cognition. For example, impenetrability can be said to be indispensable for keeping mental computations tractable or for explaining why some perceptual illusions remain despite knowledge of their illusory nature.

Some reasons are specific to a metaphysical view. In the mereological framework, impenetrability can be considered as the only way to keep the arrangement between parts simple enough to allow scientific success in psychology. The following quote of Fodor can be interpreted this way:

The condition for successful science (in physics, by the way, as well as psychology) is that nature should have joints to carve it at: relatively simple subsystems which can be artificially isolated and which behave, in isolation, in something like the way that they behave in situ. (Fodor, 1983, p.128)

In the taxonomical framework, penetrability put some traditional epistemic relations in danger. The production/combinatory rules (associated with additional considerations) could be said to partly ground some epistemic relations:

- (i) perceptual states justify perceptual judgments (that they produce).
- (ii) perceptual judgments can work as additional evidence in reasoning.

I have said that accepting CPP data in a taxonomical framework means accepting an additional production rule (i.e. some cognitive states participate in the production of some perceptual states). This rule puts pressure on the previous epistemic relations. Is it still true that perceptual states justify beliefs? Is it still acceptable to add evidence based on a perceptual judgment in a reasoning, if the belief apparatus of the perceiver influences perception in the first place?

Let's take an example of a perceiver put in front of a yellow banana. In traditional models, sensory information produces a perceptual state P with a content roughly of the form: "yellow banana" and this perceptual state may produce a belief B1 (a perceptual judgment) with the content: "there is a yellow banana here". Independently, we may have knowledge about what is the typical color of bananas, so we may have another belief B2 with the content: "bananas are yellow". In this example, P justifies B1 and B1 is a supplementary evidence for B2. However, according to CPP, if knowledge about the typical colors of fruits influences perception (Hansen et al., 2006), then B2 participates in the production of P. It seems therefore less acceptable to ground knowledge about typical colors (B2) on P and B1 (indirectly, B2 justifies itself).⁴

For this reason, as well as others, one could prefer resisting CPP in order to maintain traditional epistemic relations. Each metaphysical view can be adequately revised to interpret CPP data as not being in favor cognitive penetrability of perception, but instead as describing intra-perceptual effects (or intra-cognitive effects). The revisions imply to move the boundaries of perception. However, whereas taxonomical and mereological framework are forced to reduce or expand the scope of perception, the focal framework is more flexible and allows to move the boundary in another section, excluding the variables associated with CPP. In what follows, I describe these revisions.

Taxonomical and mereological joints For accommodating CPP in the taxonomical framework, a new production rule is added to the traditional ones: some cognitive states (let's call them Cx) participate in the production of some perceptual states (let's call them Px). In order to avoid postulating this rule (while accepting the reliability of CPP data), it is possible to move the boundary of perception to make CPP effects intra-perceptual or intra-cognitive. For this purpose, one should claim that either *Cx are perceptual*, or *Px are cognitive*.

⁴Some researchers accept to tackle this new challenge, in which traditional epistemic relations should be modified to take into account CPP. One brilliant solution consists in acknowledging rationality (and irrationality) to perceptual experiences and their processes of formation, so that the traditional relations (i) and (ii) depend on the epistemic status of P. This is the thesis called the rationality of perception (Siegel, 2017).

If the effects are intra-perceptual, they are accounted by the traditional production rule: some perceptual states participate in the production of some perceptual states. If the effects are intra-cognitive, they are accounted by the traditional production rule: some cognitive states participate in the production of some cognitive states. The upshot (among others) is that the epistemic relations mentioned previously are safe.

Concerning perception, the suggested revisions consist in moving the boundaries to include more mental states as perceptual, or to exclude some perceptual states as being genuinely perceptual. When it comes to the concept PERCEPTION, the former expands its extension, when the latter shrinks it.

Likewise, in the mereological framework, one may revise the boundaries of the decomposition to interpret CPP effects as being intra-perceptual or intra-cognitive. In moving the boundaries of perception, the interactions between parts will be reduced, which allow to keep a simple arrangement (the subsystems remain aggregated instead of integrated). This is Mandelbaum's strategy when he incorporates categorization into the perceptual module (Mandelbaum, 2017). What appeared to be cognitive influence on perception in the traditional decomposition of the mind is shown to be interaction *within perception* in his modular decomposition.

In order to avoid impenetrability (i.e. in order to avoid supplementary interaction between classes or systems), both the taxonomical and mereological frameworks can change the scope of perception. However, changing the perceptual boundaries impact the homogeneity of classes or systems (if one adds a new kind of mental states in the class of perceptual states, the class of perceptual states become more heterogeneous). Yet, the homogeneity is an important ingredient for the explanatory power of a classification or a decomposition (see chapter 6). That is why the alternative solution provided by the focal framework is interesting: it allows to avoid penetrability without putting the homogeneity of its partition in danger.

Focal joint or the cure to Fodor's pessimism In a focal framework, it is also possible to change the boundaries of perception within a section. This solution is similar to the revisions suggested above in the mereological framework. Another solution is possible here, and not available within the other metaphysical views: if CPP is salient in the section in which perception and cognition are parts, it could be possible to build a new section with a different set of variables, in which CPP is *not* salient. In this case, the boundaries of perception are moved *from a section to another one*.

In this new section (let's call it S-new), perception and cognition can be distinct without CPP effects. Obviously, S-new should justify its existence in explaining some scientific problems. Let's take for example Rolfs and Dambacher's functional section (2016) in which perception and cognition are individuated by their functional roles. According to them, in this functional section, perception and cognition can be considered distinct without CPP effects. If some problems (e.g. epistemological issues like justification of beliefs) can be tackled in terms of functions, then CPP is not a problem anymore for them. It does not mean that CPP effects does not exist, but they belong to another explanatory network than the one required for problems involving perception and cognition as such.

This solution provides additional freedom in explanatory strategies set up by scientists

and philosophers. It is an answer to Fodor's pessimism when he claims that scientific investigation about a complex system is possible only if the system is breakable into quasi-independent parts (like in the mereological framework). The focal framework provides an explanatory strategy for entering into integrated complex systems like the mind (a complex system that is not near-decomposable), in orienting a decomposition with a chosen set of variables considered important for the problem at hand.

8.3 Resistance to Eliminativism

CPP effects have been said to put the distinctness of perception in danger. This is roughly the position of eliminativists, for which perception (as well as cognition) should not be considered as genuine distinct parts of the mind.

The idea of eliminativists is the following: the more permeable the border, the less persuasive the evidence of the separation between two distinct parts. In this line of reasoning, Lupyan says:

By stressing the interactive and distributed nature of neural processing, I am committing to a collapse between perception and cognition (Lupyan, 2015, p.38)

In order to illustrate the eliminativist reasoning, here is a short fiction. A scientific society commissioned a group of ethnologists to do research on a population living in a very isolated region of the world. Before the study, they had heard of an old legend talking about two different populations living in this region, the Pecs and the Pens. The ethnologists studied cultural, physical and linguistic traits of some of these individuals. The data were messy, but they managed to group the individuals into two groups, even if some traits were present in both groups. They concluded that two groups of people inhabit the area, and that there is a population flow going only from the Pecs area to the Pens area, as they found some traits of the Pecs in the Pens, but not the other way around. After several discussions, the scientific society approved the following explanation: the region is actually constituted by a valley, in which the Pens lived, and a plateau, in which the Pecs lived, with a big cliff between them. The Pecs mastered the technology of toboggan and built a giant toboggan between the plateau of the Pecs and the valley of the Pens. As a result, the Pecs were able to reach the valley, but the Pens were unable to go to the plateau. Several years later, an ethnologist found new evidence that not only some Pens' traits were found in the Pens, but some Pens' traits were found in the Pecs too. The new scientific challenge was then to explain how the Pens could have climbed the toboggan, until a young ethnologist suggested that the Pens know maybe how to build a ladder. This last suggestion was adopted, especially to silence another disturbing idea, that maybe there was only a single group of people living in a cliff-free region, with only individual differences.

When it comes to the debate about the perception/cognition border, some researchers have a similar reasoning. CPP effects put into question the mere existence of the distinction between perception and the rest of cognition. As the old legend on Pecs and Pens, the label of perception would be a misconception of folk psychology and does not refer to an actual difference in the mind. It does not mean that the cognitive penetrability of

perception is incompatible with the existence of a perceptual border (after all, mutual influence between Pecs and Pens is compatible with the existence of two groups, a toboggan and a ladder). The cognitive penetrability of perception does not imply eliminativism, but it gives some reasons to reconsider the distinction between perception and the rest of cognition.

I show in this last section that the nature of the joint-in-nature between perception and cognition influences the resistance of the perceptual border to eliminativism. A mereological joint-in-nature is especially vulnerable to eliminativist pressure. On the contrary, a taxonomical or focal joint-in-nature is more resistant, but for different reasons.

In a *mereological framework*, CPP effects put in doubt the *decomposability* of the system. I have shown before that a mereological model can attempt to accommodate CPP in changing the arrangement between its parts, or avoid it in changing the decomposition. However, the change of arrangement or decomposition has to satisfy the conditions of near-decomposability, i.e. the arrangement should remain simple. If the revisions fail, the (near-)decomposability of the system is in danger, which means that there is no mereological joints-in-nature between parts of the mind. The upshot is that the border between perception and cognition is not a (mereological) joint-in-nature. This is an eliminativist conclusion. In more details, the argument is the following:

- (1.) The revisions fail, namely there is no fundamental decomposition of the mind (in which perception and cognition are parts)
- (2.) A system with no fundamental decomposition is an integrated system (vs aggregated)
- (3.) An integrated system is not (near-)decomposable
- (C.) Therefore, the mind is not (near-)decomposable

However, this argument only says something about *mereological* joints-in-nature. If the mind is an integrated complex system, it is still possible that there is a taxonomical or a focal joint-in-nature between perception and cognition.

In a taxonomical framework, the non-decomposability of the mental system does not lead to eliminativism. Nothing prevents indeed interaction between kinds of states or processes, so that CPP is not a direct argument in favor of eliminativism. There is however an abductive argument for it, based on some limitation that CPP imposes on a taxonomical joint-in-nature. CPP limits the way kindhood membership is defined, as the informational isolation is not available anymore as a ground for membership.

In a taxonomical framework, a mental state belongs to a kind. Whether a mental state is cognitive or perceptual is grounded on something – either a fundamental resemblance with the members of the kind, the same source (common mechanism) or the same fate (common functional role). CPP limits taxonomical framework in the sense that the kindhood of perception cannot be grounded anymore in an isolated production mechanism.

However, a lot of other criteria are untouched by CPP and can still ground mental kinds like perception and cognition in a fundamental way. In these cases, the taxonomical joint-in-nature between perception and cognition is maintained. The list of available candidates

is still well-stocked: all representational criteria (type of content, format), phenomenological criteria, functional criteria, or even production mechanisms if they are not defined by their isolation.

This limit can however be used as an indirect argument in favor of the elimination of the distinction between the perceptual and cognitive kinds of mental states. The argument is the following:

- (1.) The taxonomy in terms of perceptual/cognitive mental states is not based in their respective isolated production mechanism.
- (2.) Isolated production processes are the only possible ground for efficiently categorizing mental states.
- (3.) Non-efficient taxonomies should be eliminated from science.
- (C.) Therefore, the taxonomy in terms of perceptual/cognitive states is not an efficient taxonomy of mental states and should be eliminated.

The second premise of the argument is strong and would require a lot of defense. This is why I do not think that CPP leads easily to eliminativism of the perception/cognition border in a taxonomical framework. It is too quick to jump from accepting CPP to eliminate the perceptual border when the perception/cognition divide is a taxonomical joint-in-nature.

Finally, the *focal framework* gives a nice resistance to eliminativism as it does not claim that a section should explain everything. In order to eliminate any focal joint-in-nature between perception and cognition, one should show that all possible sets of variable building sections with perception and cognition as parts generate bad sections, i.e. sections that explain nothing at all. This would really be a change of paradigm, in which all mentions of perceptual and cognitive phenomena are swept away from cognitive science in favor of other scientific concepts.

8.4 Conclusion

I have shown in this chapter how a metaphysical view on the nature of the joint-in-nature between perception and cognition influences the way experimental data are accommodated and call for revisions of traditional models. I have also shown how these data put pressure on the divide between perception and cognition. There are two striking conclusions to this chapter. First, depending on the metaphysical framework, integrating CPP data requires different revisions. Secondly, CPP is not a strong argument in favor of eliminating the divide. There are many ways in which the divide can make room for CPP, and only few revisions that actually lead to eliminativism.

The three eliminativist threats are the following:

1. In a *mereological* framework, there is a *direct* eliminativist threat if the fundamental decomposition of the mind cannot accommodate CPP effects.

2. In a *taxonomical* framework, there is an *indirect* eliminativist threat if isolated production mechanisms is the only legitimate ground for kinds of mental states.
3. In a *focal* framework, there is an *indirect* eliminativist threat if all sections built with perception and cognition as parts as proved to be explanatorily useless.

Note that all eliminativist threats only count for their type of joint-in-nature. Therefore, in order to argue that there is no joint-in-nature at all between perception and cognition, it requires that the three threats are decisive.

What I want to emphasize here is how much the metaphysical view about the mind modifies the interpretation and integration of new experimental data in cognitive science. When one says that cognitive penetrability of perception is an empirical matter, it is incomplete. Yes, of course, we will not solve the question with 'armchair' philosophy alone, but ignoring the import of metaphysical views of the mind is also a mistake. I hope this chapter can help clarify the metaphysical possibilities offered to the theorists, in order to ground a better interpretation of empirical data when it comes to disentangle perceptual phenomena from others (cognitive, emotional, imaginary, etc.).

In this third part of the dissertation (chapters 7 and 8), I focused on the *ontological* border of perception, and I showed that this border is actually ambiguous between three kinds of joints-in-nature: taxonomical, mereological and focal. These distinctions are orthogonal to the one I explored in the second part, during conceptual engineering. Each *x*-PERCEPTION can actually endorse each of the metaphysical views I presented here⁵. After the second and third parts, the object of scientific study 'perception' is conceptually fragmented and metaphysically ambiguous. Should we really continue to work with this slippery notion? Eliminativism could find here a new breach: we should abandon PERCEPTION as scientific notion and work with new and more specific notions, such as BOTTOM-UP / TOP-DOWN PROCESSES, or even *x*-PERCEPTIONS. This is a kind of *conceptual* eliminativism, based on a pragmatic argument. The idea is roughly that PERCEPTION is a confusing notion, and science should avoid confusing notions. I will show in the last part of the dissertation that we should resist this argument and the temptation to abandon PERCEPTION.

⁵The focal framework gives a nice integrated picture of the four *x*-PERCEPTIONS: they are parts of different sections of the mind, built through variables corresponding to the problems they respectively tackle.

Part IV

Against Eliminativism

Chapter 9

The scientific role of perception

This is the last chapter of the dissertation and it is time to draw the consequences of the analyses carried out during this work. I have shown that perception, as an object studied by science, is difficult to circumscribe, and therefore difficult to study. The concept of perception is used in scientific and philosophical research programs with different aims, ascribing it different meanings. When it comes to the nature of the distinction between perception and the rest of the mind, the border is metaphysically ambiguous. The final question to be asked is what the role of perception in science can be. Does it still have a role to play? Or must we accept that the notion is outdated by more specific scientific notions, to which it should give way.

If, in a scientific context, perception is no more than an umbrella term, i.e. a term used to talk about a group of heterogeneous things with no epistemic contribution, then it is legitimate to question perception as a genuine scientific object in the first place, and to call for an elimination of it. This conclusion, to which my work could lead, must however be treated carefully. If terms have been eliminated in the history of science (e.g. phlogiston), they have often been removed along with the theories that lent credibility to them. Here we are not in this situation. The theories using the notion of perception are far from being dead, and even on the contrary, they abound. The idea pursued by the advocates of the elimination of perception is not to dismiss bad theories, but to refine scientific instruments to make them more effective.

In this last chapter, I propose an alternative. I suggest that the general notion of perception still has a role to play in science, and this role is organizational. I call it an *organizational concept*. Unlike an umbrella term, an organizational concept contributes to research and scientific success. In other words, it provides a genuine epistemic contribution.

In claiming that perception is an organizational concept, my goal is to recognize the usefulness of a concept in a role that is not traditionally attributed to it. I do not have the ambition to build a general theory on the structure of cognitive science, in which such concepts would have a place. It would be interesting, though, to explore the links between the notion I sketch here, and the theories of scientific structures in general. I leave these investigations for a future work.

The identification of this organizational role of PERCEPTION allows me, in this work, to counter effectively arguments in favor of eliminativism. Moreover, it allows me to make better recommendations for the study of perceptual phenomena in general. I do not merely

say *what not to do* with PERCEPTION, but also *what can be done* with it to help research.

In order to appreciate the value of an organizational concept, it is required to have a dynamic vision of science. Science should be conceived as a human activity instead of a final product of scientific research. Doing science in this sense requires a conceptual structure to guide and carry scientific progress. In a metaphorical image, doing science is like building a house. You need bricks, but also a framework, plans and scaffolding. Some of these tools will disappear once the final product is delivered (e.g. scaffolding and plans), and some will remain (e.g. framework). Organizational concepts are part of the framework of science. They are not sufficient for scientific knowledge, but they are necessary.

The idea that science is a human activity (and should be studied as such) is advocated by the proponents of a cognitive approach of science, which rejects the idea that science should be an interpreted, axiomatic system, but is instead a complex interrelated network of models, built with tools such as metaphors, analogies, etc. to produce new hypotheses and theories (Giere, 2017; Nersessian, 2008). I consider (without defending it here) organizational concepts as one of these tools used by scientists to organize their knowledge and beliefs, both individually and collectively.

The challenge of this last chapter is to show that the general concept of perception, despite its conceptual fragmentation and metaphysical ambiguity, still has a role to play in science. This role must be strong enough to resist arguments aimed at eliminating the concept (either because it is redundant, useless or confusing), but at the same time should take into account the analysis I produced in the previous parts of the dissertation. The ultimate goal is to rightfully place the notion of perception into the interrelated networks of scientific production. Once this place is identified, we should be able to use the concept in the best conditions for it to contribute to scientific success.

The chapter will unfold as follow: I first show why PERCEPTION is an organizational concept and its contribution to science (section 1), and then I confront PERCEPTION to arguments in favor of eliminativism (section 2). I will show that the organizational role of PERCEPTION protects it against elimination.

9.1 Organizational Concept

In this section, I argue that PERCEPTION is an organizational concept. It contains a research agenda and provides epistemic standards. Organizational concepts are scientific tools with a genuine epistemic contribution, both instrumental and intrinsic. In arguing that PERCEPTION is an organizational concept, I point out a specific role PERCEPTION plays in scientific research. This role justifies to keep it in the scientist toolbox (against eliminativism), but also designates the scientific roles PERCEPTION does not play, and should not play: PERCEPTION is neither a mere umbrella term, nor a scientific concept referring to a single and well-demarcated entity in the world.

9.1.1 PERCEPTION is an organizational concept

An organizational concept, in a scientific context, is a concept that embodies a general epistemic goal, which organizes a complex network of more specific research investiga-

tions. It does not have well-defined extension and intension, either because the concept has been fragmented through its history, or because it already refers to a general inquiry. It contributes to science in many ways that could be described together as a general role of organization: it orients scientific questions, evaluates discoveries, models and theories, coordinates knowledge coming from different scientific perspectives and integrates these perspectives into a global understanding of its object of research.

Well-known organizational concepts are famous disciplines of science, like PHYSICS or BIOLOGY. They refer to complex domain to be studied, with specific methods, known success and unresolved problems. We can build artificially another extension for PHYSICS as applying to all physical objects, or BIOLOGY to all biological entities. But this is not what means PHYSICS or BIOLOGY. They do not refer to a class of object, but to a specific human activity of problem-solving. This seems obvious for these concepts, but less for the numerous other organizational concepts that have not been (yet) institutionalized as constituting 'scientific disciplines' and do not give their label to departments in universities. For example, this is the case of the concept EVOLUTIONARY NOVELTY (Brigandt, 2012). This concept is used in evolutionary biology, more particularly in evolutionary developmental biology. There are disagreements on how to define NOVELTY, so that the concept is not really apt in identifying novel structural changes. It has consequently no well-defined extension or intension. However, according to Brigandt, this concept is not a classificatory tool.

[Its] primary function [...] is to set a problem agenda, i.e., to point to a phenomenon in need of explanation. [...] In this case the problem is the explanation of the evolutionary origin of novelty, and given the nature of this particular problem, it is clear that knowledge from different biological disciplines is required – developmental biology, paleontology, phylogeny, and evolutionary genetics, among others. (Brigandt, 2012, p.8)

EVOLUTIONARY NOVELTY is what I call an organizational concept. It does not refer to a clear object or mechanism of the world. However, this is not a failure, as it is not its vocation to do it. Instead, this concept works like a question, a scientific problem that creates and organizes scientific research and knowledge.

Likewise, PERCEPTION is an organizational concept. It sets a problem agenda for several disciplines, namely *explaining how and to what extend the external world is brought into our mental world, for us to learn, decide, act, know, survive*. This broad problem is actually so complex that it should be accessed by different specific questions, approaches and methods. In this perspective, PERCEPTION do not refer to a single mechanism or a well-defined class of states or processes.

In this dissertation, I gathered some evidence that this organizational role of PERCEPTION is at work in cognitive science. I have shown that perception is currently subject to conceptual fragmentation (chapter 6, section 6.3). This process is the sign that PERCEPTION is used with different meanings in different scientific investigations. These different meanings are related, though (chapter 5, section 5.3). There is a sense in which people are still talking about *the same things* – not the same mental mechanism or entity, but *the same questions*.

As an organizational concept, PERCEPTION is almost institutionalized. It has been rightfully noted that the distinction between perception and cognition is

woven so deeply into cognitive science as to structure introductory courses and textbooks, differentiate scholarly journals, and organize academic departments. (Firestone and Scholl, 2016, p.1).

Interestingly, this fact is taken to be an additional evidence for the distinction between two (natural) kinds of mental states or processes (see also introductions of Beck, 2018; Phillips, 2019). However, this pervasive distinction could also be the manifestation that PERCEPTION (and probably COGNITION) is an organizational concept. This institutional distinction does not say much about how the world is (institutionalizing a concept does not secure the clarity of its reference), but it says something about scientific goals, interests and organization.

The main alternative to the idea that PERCEPTION is an organizational concept is the view that perception is a mere umbrella term, i.e. a term used to talk about a group of concepts with no epistemic contribution. Umbrella terms could at best be "useful shorthand for a descriptive phrase" (Machery, 2009, p.239).

However, I show in the next subsection that, as an organizational concept, PERCEPTION contributes genuinely to scientific research and knowledge. PERCEPTION guides, coordinates and integrates interdisciplinary knowledge. It also provides epistemic standards, without which scientists won't be able to compare different theories. In other words, as an organizational concept, PERCEPTION genuinely contributes to the epistemic value of scientific knowledge.

9.1.2 The epistemic contribution of organizational concepts

The usual role granted to umbrella terms is their usefulness in communication, and I can easily find an example in which using a general and vague notion of 'perception' has been useful in this sense. For example, I usually begin the short presentation of my work to non-academic friends by something along the following lines: "Perception, through vision or audition, does not work like a camera or a voice recorder". But even if 'perception' is a useful term for me to keep my non-academic friends on board, this is not its only value. PERCEPTION has also a genuine epistemic value as an organizational concept.

The problem agenda carried by PERCEPTION works as a general epistemic goal and produces epistemic standards, which guides, coordinates and integrates research about perception.

By stating which phenomena are in need of explanation, what problems are unresolved, the problem agenda sets *standards of adequacy*, i.e. what counts as a good explanation (that meets the epistemic goal) and what are the suitable evidence and methods of investigation to reach them (Brigandt, 2012). Similar epistemic standards are also set by the epistemic goals of the four subconcepts of perception I identified in chapter 5 (phenomenal, epistemological, algorithmic and biological concepts of perception, see section 5.3). However, the problem agenda carried by PERCEPTION produces epistemic standards for these different subconcepts to work together, to feed each other, in order to foster scientific creativity and knowledge integration. Here are at least three epistemic contributions of PERCEPTION for the whole field of perception science:

1. **A guidance role:** PERCEPTION guides hypotheses in a more flexible way than its subconcepts can do. It may produce new subconcepts (conceptual fragmentation does not have to be over by now, there may be more approaches to find).
2. **A coordination role:** I argued that the four subconcepts are in principle independent but may actually work as constraints for each other (see chapter 4). This is actually the case because they are all subsumed under the organizational concept PERCEPTION. The implicit recognition of PERCEPTION as a single complex network of problems is what makes scientists across disciplines talk to each other and learn from each other. Even if there is a risk of misunderstanding because everybody is not yet clear about which subconcepts they are using, the possibility of coordination, even with a fallible tools, is indispensable for scientific creativity and progress (Nersessian, 2008).
3. **An integration role:** PERCEPTION also organizes the results across disciplines and different epistemic goals into a comprehensive whole. Brigandt says for biology that "integration [...] is not the stable theoretical unification of different fields, but the dynamic coordination of various epistemic units (explanations, models, concepts, methods) across several fields" (Brigandt, 2012, p.10). I think the same applies to cognitive science, at least in the case of perception studies. The questioning about perception is complex enough for its ramification to produce epistemic units sufficiently different for them to not be unifiable into a single theoretical apparatus. Instead, our best hope is that our scientific knowledge is structured enough for us to build representational bridges between the different windows opened on reality by different approaches.

9.1.3 Instrumental and intrinsic epistemic value

This epistemic value can be said to be merely instrumental. Guidance, coordination and integration improve the chance that theories produced by a domain are true, but they are not part of theories, not part of knowledge.

Let's take again the metaphor of house building: One may think that organizational concepts are like scaffolding for building scientific knowledge, i.e. they help build the house, but they are removed at the end, when the building holds by itself. I do not think this is the case for PERCEPTION. On the contrary, PERCEPTION will still be there in mature theories of the mind, because it contributes to the general understanding of a broad and complex subject. In this sense, PERCEPTION belongs to the *framework* of cognitive science, instead of its scaffolding. It remains within the building at the end, because it helps holding the structure. There is therefore an intrinsic epistemic value to this organizational concept.

This intrinsic epistemic value is actually associated with the integration role. They are two sides of the same coin. An organizational concept does not only participate in the integration of knowledge, it also keep it together for scientists to have access to it, and grasp it as a whole. It participates therefore in the general understanding of the phenomena under scrutiny, which is arguably a genuine part of what it is to have full-fledged knowledge of something. In ascribing a similar value to characterizations (i.e. open-ended intuitive conceptions of a subject), Camp says:

On the one hand, in many domains we credit an agent with understanding the subject when they display a flexible ability to navigate among, draw

connections between, and imagine appropriate counterfactual modifications to particular facts, even if they are unable to explicitly articulate those connections in themselves, let alone form higher-order explanations for why those connections and counterfactual modifications obtain. And on the other hand, if an agent cites a particular higher-order structural propositions in response to a particular "why"-question, without displaying any more general ability to explain why that proposition is true, then this will call into question our ascription of both understanding and knowledge-why. (Camp, 2019, p.40)

Understanding a complex subject requires grasping knowledge coming from different approaches into a cohesive whole. In this enterprise, organizational concepts like PERCEPTION are keystones that hold irregular bricks of knowledge together. To push the metaphor a little further, if we eliminate PERCEPTION from our body of scientific concepts, we will end up with a pile of bricks, and a mere pile of bricks is not a house.

The notion of organizational concept is largely inspired by Camp's work on characterizations. In Camp's vocabulary, a characterization is

a set of intuitive beliefs about an individual or a kind, which need not be extension-determining, or constitutive of conceptual competence, or even reflectively endorsed; but which are easily evoked and provide the standard mental setting for thinking about a subject (Camp, 2019, p.19).

She ascribes a similar role to characterizations:

Characterizations facilitate smooth interaction within contexts by guiding attention and response, and by synthesizing rich bodies of information and explanation into intuitive wholes (Camp, 2019, p.28)

Organizational concepts are different from characterizations. They must not be intuitive and they are scientific tools (characterizations are general cognitive tools). However, their similar role can be associated with their shared properties, namely their lack of clear extension and their power to provide a framework in which people can understand each other even if they do not share exactly the same view about the discussed object.

To summarize, the general notion of perception still has a genuine role to play in scientific investigation and in the general understanding of the mind. In the following section, I show that this is sufficient to resist eliminativist arguments. The upshot is that we should keep PERCEPTION as a legitimate scientific concept, but should recognize its role of organizational concept, and use it as such.

9.2 Eliminativist arguments

In this section, I present three arguments that a proponent of eliminativism could put forward. These arguments do not defend a radical elimination à la Churchland (1981), in which propositional attitudes like perception are supposed to be eliminated because they do not exist. Here, the case of PERCEPTION asks for replacement more than elimination. When a theory is revealed to be wrong and involves non-existing entities (e.g. phlogiston or ether), these entities are completely eliminated (without replacement). This kind of radical elimination calls for different arguments and is not what is at stake here. I nonetheless

call eliminativism the thesis according to which PERCEPTION should be replaced, because it is how it is called in the literature.

The first argument, i.e. the risk of confusion, is the more threatening one for PERCEPTION. Thanks to its organizational role, it is however possible to resist it. I will then show that the two other eliminativist arguments (conceptual parsimony and the necessity of representing natural kinds) are flawed arguments for eliminating scientific concepts in general.

The conclusion is that PERCEPTION should not be eliminated from science. However, the argument of the risk of confusion highlights a genuine pitfall in scientific research today when it comes to perception, and this risk will be taken into consideration for the recommendations I will make in the conclusion of the dissertation.

9.2.1 The risk of confusion

Roughly, the argument of the risk of confusion says that a concept subject to conceptual fragmentation is both *ambiguous* and *superfluous*, which generate a high risk of confusion in scientific debates. Consequently, eliminating it will improve scientific practice.

I showed in chapter 5 (section 5.3) that PERCEPTION is subject to conceptual fragmentation. This process happens either because a group of phenomena is discovered to be heterogeneous and actually consists in several more specific classes of homogeneous phenomena (e.g. the case of the concept CONCEPT (Machery, 2009)), or because an object of research is approached with different epistemic goals or theoretical apparatus (e.g. the case of SPECIES (Ereshefsky, 1992)).

In each of these cases, the new concepts that result from the conceptual fragmentation are considered as alternative concepts to the original one¹. What I call alternative concepts are concepts that can replace another one in each of its instance in theoretical generalizations, without changing the scope of the generalizations.

The idea is that each time the concept X, say SPECIES, is used in a scientific investigation, it is always used with a specification corresponding to one of its alternative concepts. If a scientist needs the concept SPECIES, she specifies the approach according to which she individuates them (ecological, interbreeding or phylogenetic approaches). She does so because generalizations will be different given the approaches.

Once conceptual fragmentation gave birth to alternative concepts, there is a pragmatic argument that urges to *get rid of* the original concept. The idea is that the original concept, not needed anymore, can generate confusion in debates.

In more details, the argument is the following: for a concept X, subject to conceptual fragmentation,

(P1.) X presents a high risk of confusion, as scientists will mean different things by X

(P2.) All things being equal, such confusion should be avoided

¹Conceptual fragmentation is not the only scientific process that can give birth to alternative concepts. Consider scientific reduction of a phenomenon to another one. Concepts from one scientific domain can be imported as alternatives within another scientific field, e.g. temperature is the average kinetic energy of molecules.

(P3.) Eliminating X avoids a possible confusion, without any loss of epistemic value

(C.) Therefore, we should eliminate X.

(P1) and (P3) are derived from conceptual fragmentation and its consequences. First, conceptual fragmentation entails that X is ambiguous between different definitions. Scientists mean different things by X, which may generate confusion and increases the possibility of verbal disputes and metalinguistic negotiations. Even if professional philosophers and scientists are expected to take special care of how terms and concepts are used, it is still not enough to secure the clarity of discourse. Taylor and Vickers (2017) illustrate this idea with the debate about the inconsistency of 'classical electrodynamics' in which Frisch was taken to claim that classical electrodynamics was inconsistent, when his view was actually not about classical electrodynamics in general but about a more specific notion. (Frisch, 2005). Taylor and Vickers comment this dispute:

The problem here is that philosophers sometimes have very fixed views on what a word can be taken to mean, such that either (i) they misinterpret the definition that is given, or (ii) they simply think the definition is wrong, so that any discussion based on that definition is at best misleading. [...] The debate would have been much more effective if 'classical electrodynamics' had been eliminated, and claims stated in alternative ways. (Taylor and Vickers, 2017, p.30)

This comment brings us to (P3), in which it is suggested that eliminating X prevents confusion while keeping the possibility to express all valuable scientific discourse. This is based on the idea that conceptual fragmentation of X makes X superfluous in creating alternative concepts (in the technical meaning I presented previously, i.e. a group of concept that can replace another one in each of its instance in theoretical generalizations, without changing the scope of the generalizations.)

Consequently, because of conceptual fragmentation, X is both ambiguous and superfluous. Conceptual fragmentation is here the producer of problems (i.e. ambiguity and confusion) and the provider of solutions (i.e. replacement by alternative concepts).

(P2) is a pragmatic claim: it seems to be more efficient for conducting scientific research to eliminate a concept that is source of confusion if it has additionally no useful theoretical role.

Now I want to nuance the scope of the argument. (P1) states that X presents a high risk of confusion. It means that the scope of the argument can be reduced to a contextual argument. In certain context, an ambiguous term or concept X will indeed produce confusion, but in some context, it will not. What is supposed to be avoided is not the risk of confusion, but confusion itself. A concept should therefore be eliminated only when it presents a *high risk* of confusion. This suggestion, called *selective eliminativism* (Taylor and Vickers, 2017) is more appealing than *eliminative pluralism* (Ereshefsky, 1992) in which ambiguity and superfluity are sufficient to justify eliminativism. Let's illustrate it with a non-scientific toy example. The term 'bank' is ambiguous between the financial institution and the land at river's edge. One could say that the term 'bank' is superfluous in the sense that we can paraphrase the two meanings easily without using the term. According to eliminative pluralism, 'bank' should be eliminated. However, for selective

eliminativism, as there is no risk (or very low risk) of confusion between the two terms used in very different contexts, there is no need to eliminate the term 'bank' to replace it by more specific concepts or paraphrases.

However, selective eliminativism faces an additional difficulty, which is to be able to evaluate the risk of confusion associated with a concept. Taylor and Vickers (2017) suggest that the risk of confusion associated with a concept X increases with: (i) the number of different theoretical roles associated with X and (ii) the centrality of X in arguments in which X is recruited.

The idea in (i) is that the more theoretical roles performed by a concept, the greater the chances to mistakenly understand the wrong meaning. Furthermore, as the different meanings are the product of conceptual fragmentation, the different meanings may be only subtly different (this is not the case of 'bank', in my previous toy example, in which the two meanings are not produced by conceptual fragmentation of the same concept, and belong to different lexicon).

(ii) estimates how dramatic the consequences may be for scientific debates if a concept is mistakenly taken to have the wrong meaning. If an argument relies heavily on a term or concept, confusion about this term can transfer to the conclusion of the argument (that may be different depending on the meaning of X), and be spread over a whole scientific field, resulting in a loss of time and energy for what is actually a verbal dispute.

The two 'heuristics' (i) and (ii) (as Taylors and Vickers call them, 2017, p. 29) are not supposed to reveal all the causes that can lead to confusion and verbal disputes. Furthermore, they are vague. One may ask how many different theoretical roles are necessary for reaching the threshold of high risk, or how we can measure quantitatively the centrality of a concept in a scientific domain. This vagueness is acceptable for the application of the argument to PERCEPTION, as my previous analyses tend to show that according to (i) and (ii), PERCEPTION actually presents a high risk of confusion in scientific discussions.

Concerning the first heuristic, I showed previously the high number of different meanings PERCEPTION can take. I presented four different epistemic goals that organize an even more higher number of different conceptions generated in the literature (chapter 4, section 4.3). Furthermore, the existence of these different meanings generate representation-level oriented issue, such as metalinguistic negotiation and verbal dispute (chapter 3).²PERCEPTION is therefore used with a sufficient number of different theoretical meanings for generating confusion.

When it comes to the second heuristic (i.e. the centrality of PERCEPTION in debates), the borders of perception play a central role in the cognitive penetrability of perception, the content of perception and the classification of peculiar mental events (see introduction, section 1.2). If the borders of perception are important for these debates, it means that the meaning of PERCEPTION matters for these issues. Depending on the intension and extension of PERCEPTION, the whole debate about the possible content of perception will differ. More dramatically, depending on the kind of metaphysical entity PERCEPTION

²One may add three other concepts of perception: a taxonomical concept of perception, a mereological concept of perception and a focal concept of perception, derived respectively from the taxonomical, mereological and focal conceptions of the perceptual borders.

refers to, empirical data can be understood and accommodated differently in a debate such as the cognitive penetrability of perception (chapter 8). PERCEPTION is therefore sufficiently central in these debates to generate a different interpretation of empirical data and to inspire different conclusions. A confusion about the meaning of PERCEPTION can have consequences on the understanding of a broad part of debates in philosophy and psychology of perception.

Following (i) and (ii), PERCEPTION presents therefore a high risk of confusion in scientific discussions. Therefore, the argument about the risk of confusion applies to PERCEPTION:

- (P1.) PERCEPTION presents a high risk of confusion
- (P2.) All things being equal, such confusion should be avoided
- (P3.) Eliminating PERCEPTION avoids a possible confusion, without any loss of epistemic value
- (C.) Therefore, we should eliminate PERCEPTION

This argument is sound, however (P3) can be resisted with the organizational role of PERCEPTION. I accept that each specific meaning of PERCEPTION can be replaced by an alternative concept (e.g. *x*-PERCEPTION) within their respective epistemic goals and associated research. However, none of these alternative concepts can play the role of an organizational concept. In other words, conceptual fragmentation of PERCEPTION does not provide alternatives for the organizational concept PERCEPTION.

I showed in the previous section that PERCEPTION, as an organizational concept, provides genuine epistemic contribution. Consequently, even if it is true that eliminating PERCEPTION avoids confusion, it requires an epistemic loss in return. This exchange, in a scientific context, is not acceptable, especially as confusion can be tackled with other tools (see recommendations in the conclusion). As (P3) is false, the conclusion fails and eliminativists lose a warrant for eliminating PERCEPTION.

The argument of the risk of confusion rests on a pragmatic claim, i.e. confusion should be avoided in order to improve scientific practice. One may think that pragmatic reasons are less powerful than ontological or epistemological reasons for eliminativism, that is why I explore two other non-pragmatic arguments. The first one is epistemological and based on the ideal of the unity of science. The second one, metaphysical, states that we should map our scientific concepts on natural kinds. I show now that the former is unconvincing and the latter rests ultimately on the same kind of pragmatic claim than the argument of the risk of confusion.

9.2.2 You (damn) conceptual luxuriance!

Unifying all scientific knowledge into a single huge and coherent system is arguably the ultimate goal of science. Conceptual fragmentation, when it leads to pluralism, may seem to be contradictory with this ideal of unification. One can build the following argument:

- (P1.) The pursuit of scientific knowledge requires its unification
- (P2.) Unification of scientific knowledge requires its (conceptual) reduction

(C.) Therefore, the pursuit of scientific knowledge requires (conceptual) reduction

We should therefore favor eliminativism of PERCEPTION (which allows conceptual reduction) towards pluralism about PERCEPTION (which multiplies the concepts) when a concept is subject to conceptual fragmentation. In this spirit, Hull warns against the "danger of pluralism" in the debate about SPECIES. He says that pluralism "provides no means or even motivation for reducing conceptual luxuriance. Without such pruning, the integration of scientific knowledge is impossible" (Hull, 1987, p.178).

Contrary to the argument of the risk of confusion, the argument of unity is epistemological. It states that conceptual reduction is a necessary scientific step towards unity of scientific knowledge. Unifying scientific knowledge will make us understand the world better than our current disparate, unrelated, even incompatible systems that describe and explain only partial portions of the world. Whether unity of science is desirable, or even possible, is debatable (Dupré, 1995), but I concede it for this argument. I focus my objection on the second premise.

It is false to affirm that conceptual reduction is required for unification of scientific knowledge. Conceptual reduction is here thought as a corollary thesis to scientific reductionism, the thesis according to which all scientific domains can ultimately be reduced to a single one, usually physics. The success of reduction theories is probably due to the existence of famous reductions in the history of Science, like the reduction of Newtonian mechanics to relativity theory, or the reduction of chemistry to atomic physics (van Riel and Van Gulick, 2019). This scientific reductionism is associated with a conceptual reductionism, in the sense that the conceptual apparatus of the reductive base science is supposed to be sufficient to explain all the reduced phenomena. It is true that reductionism is one way to reach unity of science. If a single domain of science like physics encompasses all scientific explanations, then science is unified. Yet, reductionism is not necessary for the unity of science. Fodor is a good example of a physicalist (someone who believes that everything is ultimately made of physical stuff), believing in the unity of science, and being anti-reductionist. Fodor's anti-reductionist idea is that kinds of entities that appear in the generalizations of special sciences are not reducible to kinds of entities to which physical laws apply. Yet, special sciences earn their stripes in identifying and explaining regularities, patterns and mechanisms at their own scale (with their specific kinds). There is nonetheless reduction-relation at the level of tokens : each token of special science phenomenon can be reduced to a physical phenomenon. However, it is false that the generalizations made by special sciences can be reduced necessarily to physical generalizations.

In Fodor's words:

Physics develops the taxonomy of its subject-matter which best suits its purposes[...]. But this is not the only taxonomy which may be required if the purposes of science in general are to be served [...]. So, there are special sciences, with their specialized taxonomies, in the business of stating some [...] generalizations. If science is to be unified, then all such taxonomies must apply to the same things. If physics is to be basic science, then each of these things had better be a physical thing. But it is not further required that the taxonomies which the special sciences employ must themselves reduce to the taxonomy of physics. It is not required, and it is probably not true. (Fodor, 1974, p.114)

As Fodor suggests, the unification of scientific knowledge can be reached through physicalism, but physicalism is weaker than scientific reductionism to physics. Some concepts of metaphysics have been developed to express anti-reductionist relations in a unified body of knowledge. One may cite, among others, the concepts of multiple realizability, supervenience, emergence and downward causation (Cat, 2017, for an overview).

I do not deny that scientific reduction, complete or partial, is possible and is a progress for scientific investigations. I deny that reduction is required to lead to the unification of knowledge. Unification of scientific knowledge is much weaker than any kind of reductionism.

Once it has been argued that reductionism, a fortiori conceptual reduction, is not required for the unity of science, then the epistemological argument to prefer eliminativism over pluralism is weaker. Furthermore, I argued previously that the general concept PERCEPTION can help integrating knowledge, and therefore help unifying science. If one maintains the first premise, and argues that unification of knowledge is a goal science should pursue, then PERCEPTION should be especially preserved because it provides a tool for reaching this goal.

9.2.3 Natural kinds... again?

Recently, (Machery, 2009) produced an argument in favor of the elimination of the concept CONCEPT in cognitive science. He analyzed the notion of concept as a mental representation and argues that it constitutes an heterogeneous group of three kinds of mental representations that do not have many properties in common: theories, exemplars and prototypes. For the sake of clarity (as talking about the concept CONCEPT can easily bring confusion), it is useful to remind the distinction between concepts-as-theoretical-tools and concepts-as-mental-representations. The former are bodies of knowledge used in theories, when the latter are bodies of knowledge used in cognitive processes. What Machery argues is that the theoretical tool CONCEPT-AS-MENTAL-REPRESENTATION should be eliminated from cognitive science.

His argument rests on the idea that we should eliminate concepts that do not refer to natural kinds in favor of concepts that refer to natural kinds. He shows indeed that CONCEPT does not refer to a natural kind (as he shown that this kind is an heterogeneous group of three kinds of mental representations), but that THEORY, EXEMPLAR, and PROTOTYPE refers respectively to natural kinds of mental representations. He concludes therefore that we should replace CONCEPT by these natural kind concepts (Machery, 2009).

The argument is interesting here because at first sight, this is a metaphysical argument. Furthermore, all the evidence I gathered through my work tend to show that PERCEPTION does not refer to a natural kind. Based on Machery's argument, it is therefore possible to build a metaphysical argument in favor of the elimination of PERCEPTION :

- (P1.) PERCEPTION does not refer to a natural kind
- (P2.) Alternative concepts to PERCEPTION refer to natural kinds
- (P3.) We should eliminate concepts that do not refer to natural kinds in favor of concepts that refer to natural kinds.

(C.) Therefore, we should eliminate PERCEPTION in favor of its alternative concepts.

(P3) underlies an idea of science in which the only legitimate scientific concepts are natural kind concepts. In defending the organizational role of PERCEPTION and its epistemic contribution to the field of perception studies, I reject this view of science and give reasons to believe that a concept that do not refer to a natural kind can fruitfully contribute to science. Furthermore, it is not clear that (P2) applies to perception, so that the argument may fail anyway.

However, I have a deeper problem with this argument. Actually, Machery uses a non-classical notion of natural kind, close to Boydian kinds, that he calls the causal notion of natural kind. He defines it as such:

A class C of entities is a natural kind if and only if there is a large set of scientifically relevant properties such that C is the maximal class whose members tend to share these properties because of some causal mechanism. (Machery, 2009, pp.232-233)

He explains:

A natural kind is a class about which many generalizations can be formulated: its members tend to have many properties in common. These generalizations are not accidental: there is at least one causal mechanism that explains why its members tend to have those properties. Finally, this class is not a subset of a larger class about which the same generalizations could be formulated. (ibid.)

In his formulation, a concept that does not refer to a natural kind is a concept that is not involved in any useful generalization. It seems therefore that the reason why we should eliminate concepts that do not refer to natural kinds is because these concepts are *not useful* in science, and not because only natural kinds should be the targets of scientific concepts. This is after all a *pragmatic* argument, and not a metaphysical one. Defining natural kinds by their *scientific relevance* in a context in which the question pertains to the scientific relevance of a concept artificially incorporates ontological issues where there is none.

To come back to the general argument, I do not think that the notion of naturalkindness, either in its classical form or not, is a good ally for the proponents of conceptual eliminativism. Ontological arguments would certainly be welcome in cases in which a concept do not refer at all (e.g. phlogiston), but not in cases of replacement due to conceptual fragmentation.

I have shown in this last chapter that the general concept PERCEPTION still has an important role to play in cognitive science as an organizational concept. This role provides it a genuine epistemic value and allows to resist arguments in favor of its elimination. We should keep PERCEPTION in our scientific conceptual toolbox, even if it may be an important source of confusion. PERCEPTION guides, coordinates and integrates knowledge about the mind along its general epistemic goal, i.e. explaining how and to what extent the external environment is grasped by our mental world. This question organizes both research and understanding of the mind, even in a purportedly end of inquiry.

Chapter 10

Conclusion

The original question that motivated this work was the boundary question, i.e. what the boundaries of perception are. At the end of the day, there isn't any simple and unique answer to this question. I showed through this dissertation that the boundary question calls actually for two other interrogations: (1) 'What are the epistemic purpose for which the notion of perception is used as a theoretical tool ?' and (2) 'Which metaphysical entity is taken to be perception?'. Once epistemic goals and metaphysical assumptions have been determined, it is then possible to draw the contours of a contextualized notion of perception. I chose to apply a method of conceptual engineering for a notion of perception (more precisely, vision) within the knowledge epistemic goal, with the assumption that perception is a set of mental states. I end up with visual perception circumscribed as follow: the ensemble of visual states contains the states produced by retinal information as soon as the information is transformed into a 3D model of the environment, and until it reaches the abstraction of basic-level categories. All the states whose content comes from the transformation of basic-level categories are post-perceptual and fall within the extension of another kind of states.

This result is not really innovative (it is close to Fodor, Burge and Block's views). However, I provide a framework here to understand the ground, the scope and the consequences of this result. This view can be grounded on a principle of epistemic utility within a specific (and limited) research goal. The consequences are: (1) this is not the only legitimate view, (2) this view of perception is especially adapted for answering questions about knowledge when they involve the notion of perception and, (3) this view provides a notion of perception that could be unsuitable to answer questions outside of its original epistemic goal.

In my work, I sketched a new framework to understand representation-level oriented issues in science and philosophy of perception. In taking into account the level of theoretical representations of the scientific domain 'perception', I showed that the field is rich in theoretical orientations. Being aware of these discrepancies between theories could help avoiding fruitless disputes and help taking advantage of all the discoveries of the domain.

This view on perception studies can be rejected based on the idea that a scientific object should not be defined by its research context (but is instead an objective object that science should reveal). To the advocates of a distinct object 'perception' with general boundaries, I answer that they make a category mistake. PERCEPTION, in its general form, is an organizational concept. If it has abstract boundaries, they inform on human knowledge and questioning, not on the architecture of the mind.

I am aware that I did not produce in this work any fatal objection to the idea that perception is a natural kind. The identification (a la Machery) of sub-classes of perceptual states that do not share the same properties would be a more appropriate defense against it. This approach, which I hope will be followed by others, will however need a solid foundation for grounding the notion of perception without the help of natural kinds. In this perspective, my dissertation can be considered as a tool provider with the method of conceptual engineering and metaphysical tools such as taxonomical, mereological and focal joints-in-nature.

While waiting for this work to be done, one can always say that it is possible that the different angles of study of perception may end up converging and discovering the existence of a single mechanism, a single set of processes or states that allow us to meet all the epistemic goals of perception at the same time. Indeed, this possibility is not excluded. However, by showing that this idea should come from discoveries - and is not an acceptable assumption- I have shifted the burden of proof. Anyone who wishes to defend a unified approach to perception has to provide reasons for believing that the unification of questions into a single object of research has a chance of success.

Following my analysis, here are some recommendations in order to clarify and help debates in cognitive science, and more specifically in philosophy of perception. The idea is to avoid as much confusion as possible around a notion that is currently under a process of fragmentation, while taking advantage of the creativity produced by the friction of different perspectives around the same scientific question.

1. We should not use PERCEPTION (alone) to refer to an entity or a class of entities, because it is ambiguous. It is better to specify the extension, or to use an alternative concept (e.g. early vision, the phenomenal concept of perception, the class of states with non-conceptual content, the informational-encapsulated part of the mind, etc.).
2. When faced with a debate involving the boundaries of perception, we should refine the debate in order to identify the epistemic goal associated with perception. It becomes then possible to define a clear extension or use an alternative concept.
3. Generalization about perception simpliciter should be scrutinized with care: either they concern a subconcept of perception (and should be reformulated with alternative concepts or specifications), or they are probably false.
4. We should stop asking questions such as "is X perceptual?" for a mental phenomenon X. PERCEPTION cannot be derived as a classificatory concept, cannot be used for ascribing properties. At best, this kind of question can be interpreted as a question about the epistemic value of X in the research program defined by PERCEPTION. In this perspective, saying that X is perceptual means that its study should be included in the research program about PERCEPTION. But I doubt that this meaning is salient. It is however legitimate to ask whether a mental phenomenon X is the product of early vision, or is encapsulated from high-level content: classificatory purposes should be guided by more specific concepts.
5. We should not eliminate PERCEPTION, or be subject to a loss of scientific structure.
6. We should institutionalized PERCEPTION in order to bring the different research teams and perspectives closer to each other.

Bibliography

- Frederick Adams and Kenneth Aizawa. *The Bounds of Cognition*. Wiley-Blackwell, Oxford, 2010. ISBN 978-1-4051-4914-3 978-1-4443-3270-4.
- Kathleen Akins. Of Sensory Systems and the "Aboutness" of Mental States. *The Journal of Philosophy*, 93(7):337–372, 1996. ISSN 0022-362X. doi: 10.2307/2941125.
- Colin Allen. On (not) defining cognition. *Synthese*, 194(11):4233–4249, November 2017. ISSN 0039-7857, 1573-0964. doi: 10.1007/s11229-017-1454-4.
- D. Ariely. Seeing sets: Representation by statistical properties. *Psychological Science*, 12(2):157–162, March 2001. ISSN 0956-7976. doi: 10.1111/1467-9280.00327.
- D. M. Armstrong. *A Materialist Theory of the Mind*. Routledge, September 2002. ISBN 978-1-134-85635-0.
- Robert Audi. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. Routledge Contemporary Introductions to Philosophy. Routledge, New York, 3rd ed edition, 2011. ISBN 978-0-415-87922-4 978-0-415-87923-1 978-0-203-84646-9.
- Emily Balcetis and David Dunning. Wishful seeing: More desired objects are seen as closer. *Psychological Science*, 21(1):147–152, January 2010. ISSN 1467-9280. doi: 10.1177/0956797609356283.
- H. Clark Barrett and Robert Kurzban. Modularity in cognition: Framing the debate. *Psychological Review*, 113(3):628–647, 2006. ISSN 1939-1471, 0033-295X. doi: 10.1037/0033-295X.113.3.628.
- Howard Bartley. *Principles of Perception*. Harper and Row, New York, second edition edition, 1969.
- Jacob Beck. Marking the Perception–Cognition Boundary: The Criterion of Stimulus-Dependence. *Australasian Journal of Philosophy*, 96(2):319–334, April 2018. ISSN 0004-8402, 1471-6828. doi: 10.1080/00048402.2017.1329329.
- Jean-Charles Billaut, Denis Bouyssou, and Philippe Vincke. Should you believe in the Shanghai ranking?: An MCDM view. *Scientometrics*, 84(1):237–263, July 2010. ISSN 0138-9130, 1588-2861. doi: 10.1007/s11192-009-0115-x.
- Ned Block. Attention and mental paint. *Philosophical Issues*, 20:23–63, October 2010. doi: 10.1111/j.1533-6077.2010.00177.x.
- Ned Block. Seeing-As in the Light of Vision Science. *Philosophy and Phenomenological Research*, 89(3):560–572, November 2014. ISSN 00318205. doi: 10.1111/phpr.12135.

- Ned Block. Is there a joint in Nature between Perception and Cognition?, 2015.
- Ned Block. The Anna Karenina Principle and Skepticism about Unconscious Perception. *Philosophy and Phenomenological Research*, 93(2):452–459, September 2016. ISSN 00318205. doi: 10.1111/phpr.12258.
- David Bourget and David J. Chalmers. What Do Philosophers Believe? *Philosophical Studies*, 170(3):465–500, 2014. doi: 10.1007/s11098-013-0259-7.
- Richard Boyd. Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 61(1/2):127–148, 1991. ISSN 0031-8116.
- Richard Boyd. Homeostasis, Species, and Higher Taxa. In R. A. Wilson, editor, *Species: New Interdisciplinary Essays*, pages 141–85. MIT Press, 1999.
- Ben Bradley. *Well-Being and Death*. OUP Oxford, March 2009. ISBN 978-0-19-156787-2.
- Franz Brentano. *Psychology From an Empirical Standpoint*. Routledge, 1874.
- Bill Brewer. *Perception and Reason*. Oxford University Press, March 2002. ISBN 978-0-19-159711-4.
- Ingo Brigandt. The epistemic goal of a concept: Accounting for the rationality of semantic change and variation. *Synthese*, 177(1):19–40, November 2010. ISSN 0039-7857, 1573-0964. doi: 10.1007/s11229-009-9623-8.
- Ingo Brigandt. The Dynamics of Scientific Concepts: The Relevance of Epistemic Aims and Values. In Uljana Feest and Friedrich Steinle, editors, *Scientific Concepts and Investigative Practice*. DE GRUYTER, Berlin, Boston, January 2012. ISBN 978-3-11-025361-0. doi: 10.1515/9783110253610.75.
- Robert Eamon Briscoe. Mental imagery and the varieties of amodal perception. *Pacific Philosophical Quarterly*, 92(2):153–173, June 2011. ISSN 02790750. doi: 10.1111/j.1468-0114.2011.01393.x.
- Tyler Burge. *Origins of Objectivity*. Oxford University Press, Oxford ; New York, 2010. ISBN 978-0-19-958140-5 978-0-19-958139-9.
- Tyler Burge. Reply to Block: Adaptation and the Upper Border of Perception. *Philosophy and Phenomenological Research*, 89(3):573–583, November 2014a. ISSN 1933-1592. doi: 10.1111/phpr.12136.
- Tyler Burge. Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints. *Philosophy and Phenomenological Research*, 88(2):485–501, 2014b. ISSN 1933-1592. doi: 10.1111/phpr.12093.
- Elisabeth Camp. Thinking with Maps. *Philosophical Perspectives*, 21(1):145–182, 2007. ISSN 1520-8583. doi: 10.1111/j.1520-8583.2007.00124.x.
- Elisabeth Camp. Perspectives and Frames in Pursuit of Ultimate Understanding. In Stephen R Grimm, editor, *Varieties of Understanding*, pages 17–46. Oxford University Press, October 2019. ISBN 978-0-19-086097-4 978-0-19-086100-1. doi: 10.1093/oso/9780190860974.003.0002.

- Joseph Keim Campbell, Michael O'Rourke, and Matthew H. Slater, editors. *Carving Nature at Its Joints: Natural Kinds in Metaphysics and Science*. Topics in Contemporary Philosophy. MIT Press, Cambridge, Mass, 2011. ISBN 978-0-262-01593-6 978-0-262-51626-6.
- Herman Cappelen. *Introduction to Conceptual Engineering*. Oxford University Press, March 2018. ISBN 978-0-19-185240-4.
- Susan Carey. *The Origin of Concepts*. Oxford University Press, May 2009. ISBN 978-0-19-988791-0.
- Peter Carruthers. *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Clarendon Press ; Oxford University Press, Oxford : Oxford ; New York, 2006. ISBN 978-0-19-920707-7 978-0-19-920708-4.
- Jordi Cat. The Unity of Science. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2017 edition, 2017.
- Chalmers. A Computational Foundation for the Study of Cognition. *Journal of Cognitive Science*, 12(4):325–359, December 2011a. ISSN 1598-2327. doi: 10.17791/jcs.2011.12.4.325.
- David J. Chalmers. Verbal Disputes. *Philosophical Review*, 2011b.
- David J Chalmers. The Varieties of Computation: A Reply. *Journal of Cognitive Science*, 13(3):211–248, September 2012. ISSN 1598-2327. doi: 10.17791/jcs.2012.13.3.211.
- David J Chalmers. What is Conceptual Engineering and What Should it Be? September 2018.
- Rocco Chiou and Anina N. Rich. The role of conceptual knowledge in understanding synaesthesia: Evaluating contemporary findings from a “hub-and-spokes” perspective. *Frontiers in Psychology*, 5, 2014. ISSN 1664-1078. doi: 10.3389/fpsyg.2014.00105.
- Paul M. Churchland. Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy*, 78(2):67, February 1981. ISSN 0022362X. doi: 10.2307/2025900.
- Andy Clark. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03):181–204, June 2013. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X12000477.
- Austen Clark. *A Theory of Sentience*. Oxford University Press, 2000. ISBN 978-0-19-823851-5.
- David Colaço, Wesley Buckwalter, Stephen Stich, and Edouard Machery. Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(2):199–212, June 2014. ISSN 1742-3600, 1750-0117. doi: 10.1017/epi.2014.7.
- Tim Crane. The nonconceptual content of experience. In Tim Crane, editor, *The Contents of Experience: Essays on Perception*, pages 136–157. Cambridge University Press, March 1992. doi: 10.1017/CBO9780511554582.007.
- Jonathan Dancy. *An Introduction to Contemporary Epistemology*. B. Blackwell, Oxford, UK ; New York, NY, USA, 1985. ISBN 978-0-631-13622-4 978-0-631-13621-7.

- Donald Davidson. *Inquiries Into Truth And Interpretation*. Oxford University Press, 1984.
- Michael Devitt. Against incommensurability. *Australasian Journal of Philosophy*, 57(1): 29–50, March 1979. ISSN 0004-8402. doi: 10.1080/00048407912341021.
- Jérôme Dokic and Jean-Rémy Martin. Felt Reality and the Opacity of Perception. *Topoi*, 36(2):299–309, June 2017. ISSN 1572-8749. doi: 10.1007/s11245-015-9327-2.
- Fabian Dorsch and Fiona Macpherson, editors. *Phenomenal Presence*. Oxford University Press, Oxford, New York, June 2018. ISBN 978-0-19-966641-6.
- Fred Dretske. *Seeing And Knowing*. Chicago: University Of Chicago Press, 1969.
- Fred Dretske. *Knowledge and the Flow of Information*. MIT Press, 1981a.
- Fred Dretske. Sensation and perception. In York H. Gunther, editor, *Essays on Nonconceptual Content*. The MIT Press, 1981b.
- Fred Dretske. Perception without Awareness. In Tamar Szabo Gendler and John Hawthorne, editors, *Perceptual Experience*, pages 147–181. Oxford University Press, January 2006. ISBN 978-0-19-171104-6.
- Fred I. Dretske. *Naturalizing the Mind*. MIT Press, 1995. ISBN 978-0-262-54089-6.
- John Dupré. Natural Kinds and Biological Taxa. *The Philosophical Review*, 90(1):66–90, 1981. ISSN 0031-8108. doi: 10.2307/2184373.
- John Dupré. *The Disorder of Things, Metaphysical Foundations of the Disunity of Science*. Harvard University Press, 1995. ISBN 978-0-674-21261-9.
- Paul Egré and Cathal O'Madagain. Concept Utility. *The Journal of Philosophy*, 116(10): 525–554, November 2019. doi: 10.5840/jphil20191161034.
- Brian Ellis. *Scientific Essentialism*. Cambridge University Press, 2001.
- Marc Ereshefsky. Eliminative Pluralism. *Philosophy of Science*, 59(4):671–690, 1992.
- Gareth Evans. *The Varieties of Reference*. Oxford University Press, Oxford, New York, October 1982. ISBN 978-0-19-824686-2.
- Paul Feyerabend. *Against Method: Outline of an Anarchistic Theory of Knowledge*. Verso, London : Atlantic Highlands, January 1975. ISBN 978-0-902308-91-6.
- Chaz Firestone and Brian J. Scholl. Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and Brain Sciences*, 39, 2016. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X15000965.
- William Fish. *Perception, Hallucination, and Illusion*. Oxford University Press, October 2009. ISBN 978-0-19-986918-3.
- William Fish. *Philosophy of Perception: A Contemporary Introduction*. Routledge Contemporary Introductions to Philosophy. Routledge, New York, 2010. ISBN 978-0-415-99911-3 978-0-415-99912-0 978-0-203-88058-6.
- J. A. Fodor. Special Sciences (Or: The Disunity of Science as a Working Hypothesis). *Synthese*, 28(2):97–115, 1974. ISSN 0039-7857.

- Jerry A. Fodor. *The Modularity of Mind: An Essay on Faculty Psychology*. MIT Press, Cambridge, Mass, 1983. ISBN 978-0-262-06084-4 978-0-262-56025-2.
- Jerry A. Fodor. *Concepts: Where Cognitive Science Went Wrong*. Oxford Cognitive Science Series. Clarendon Press ; Oxford University Press, Oxford : New York, 1998. ISBN 978-0-19-823637-5 978-0-19-823636-8.
- Jerry A. Fodor. The Revenge of the Given. In Brian P. McLaughlin and Jonathan D. Cohen, editors, *Contemporary Debates in Philosophy of Mind*, pages 105–116. Blackwell, 2007.
- Ronald H. Forgas and Lawrence E. Melamed. *Perception: A Cognitive-Stage Approach*. McGraw-Hill Series in Psychology. McGraw-Hill, New York, 2d ed edition, 1976. ISBN 978-0-07-021620-4.
- L. R. Franklin-Hall. Natural kinds as categorical bottlenecks. *Philosophical Studies*, 172 (4):925–948, April 2015. ISSN 0031-8116, 1573-0883. doi: 10.1007/s11098-014-0326-8.
- Ori Friedman and John Turri. Is Probabilistic Evidence a Source of Knowledge? *Cognitive Science*, 39(5):1062–1080, 2015. ISSN 1551-6709. doi: 10.1111/cogs.12182.
- Mathias Frisch. *Inconsistency, Asymmetry, and Non-Locality: A Philosophical Investigation of Classical Electrodynamics*. Oxford University Press, Oxford ; New York, 1 edition edition, March 2005. ISBN 978-0-19-517215-7.
- Michael Ghiselin. Species Concepts, Individuality, and Objectivity. *Biology and Philosophy*, 2(2):127–43, 1987.
- Ronald N. Giere. Cognitive Approaches to Science. In W. H. Newton-Smith, editor, *A Companion to the Philosophy of Science*, pages 41–43. Blackwell Publishers Ltd, Oxford, UK, September 2017. ISBN 978-1-4051-6448-1 978-0-631-23020-5. doi: 10.1002/9781405164481.ch6.
- Peter Godfrey-Smith. A Modern History Theory of Functions. *Noûs*, 28(3):344–362, 1994. ISSN 0029-4624. doi: 10.2307/2216063.
- Robert L. Goldstone, Joshua R. de Leeuw, and David H. Landy. Carving nature at its joints or cutting its effective loops? On the dangers of trying to disentangle intertwined mental processes. *Behavioral and Brain Sciences*, 39, 2016. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X1500271X.
- Nelson Goodman. *Fact, Fiction, and Forecast*. Harvard University Press, 1955. ISBN 978-0-674-29071-6.
- Peter J. Graham. The Function of Perception. In Abrol Fairweather, editor, *Virtue Epistemology Naturalized*, pages 13–31. Springer International Publishing, Cham, 2014. ISBN 978-3-319-04671-6 978-3-319-04672-3. doi: 10.1007/978-3-319-04672-3_2.
- Verne Grant. *Plant Speciation*. Columbia University Press, New York, 1981. ISBN 978-0-231-05112-5 978-0-231-05113-2.
- Paul E. Griffiths. *What Emotions Really Are: The Problem of Psychological Categories*. What Emotions Really Are: The Problem of Psychological Categories. University of Chicago Press, Chicago, IL, US, 1997. ISBN 978-0-226-30871-5. doi: 10.7208/chicago/9780226308760.001.0001.

- Thomas L. Griffiths, Charles Kemp, and Joshua B. Tenenbaum. Bayesian Models of Cognition. In Ron Sun, editor, *The Cambridge Handbook of Computational Psychology*, pages 59–100. Cambridge University Press, Cambridge, 2008. ISBN 978-0-511-81677-2. doi: 10.1017/CBO9780511816772.006.
- Peter G. Grossenbacher and Christopher T. Lovelace. Mechanisms of synesthesia: Cognitive and physiological constraints. *Trends in Cognitive Sciences*, 5(1):36–41, January 2001. ISSN 13646613. doi: 10.1016/S1364-6613(00)01571-0.
- Thorsten Hansen, Maria Olkkonen, Sebastian Walter, and Karl R Gegenfurtner. Memory modulates color appearance. *Nature Neuroscience*, 9(11):1367–1368, November 2006. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn1794.
- Sally Haslanger. Gender and Race: (What) Are They? (What) Do We Want Them To Be? *Noûs*, 34(1):31–55, March 2000. ISSN 0029-4624, 1468-0068. doi: 10.1111/0029-4624.00201.
- Richard G. Heck. Nonconceptual Content and the "Space of Reasons". *The Philosophical Review*, 109(4):483–523, 2000. ISSN 0031-8108. doi: 10.2307/2693622.
- Franz Huber. The Plausibility-Informativeness Theory. In V.F. Hendricks and D.H. Pritchard, editors, *New Waves in Epistemology*, pages 164–191. Basingstoke, palgrave macmillan edition, 2008.
- David L. Hull. Genealogical Actors in Ecological Roles. *Biology and Philosophy*, 2(2):168–184, 1987.
- Zoe Jenkin. The Epistemic Role of Core Cognition. *The Philosophical Review*, 129(2): 251–298, April 2020. ISSN 0031-8108, 1558-1470. doi: 10.1215/00318108-8012850.
- Mark Johnston. The Obscure Object of Hallucination. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 120(1/3):113–183, 2004. ISSN 0031-8116.
- Daniel Kaiser, Genevieve L. Quek, Radoslaw M. Cichy, and Marius V. Peelen. Object Vision in a Structured World. *Trends in Cognitive Sciences*, 23(8):672–685, August 2019. ISSN 1364-6613. doi: 10.1016/j.tics.2019.04.013.
- Philip Kitcher. Does 'Race' Have a Future? *Philosophy & Public Affairs*, 35(4):293–317, 2007. ISSN 1088-4963. doi: 10.1111/j.1088-4963.2007.00115.x.
- Joshua Knobe and Shaun Nichols. Experimental Philosophy. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2017 edition, 2017.
- Uriah Kriegel. The Perception/Cognition Divide: One More Time, with Feeling. In C Limbeck-Lilienau and F Stadler, editors, *The Philosophy of Perception and Observation*. De Gruyter, 2019.
- Saul A. Kripke. *Naming and Necessity*. Harvard University Press, 1980.
- Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago, 1962.

- Daniel T. Levin and Mahzarin R. Banaji. Distortions in the perceived lightness of faces: The role of race categories. *Journal of Experimental Psychology: General*, 135(4):501–512, 2006. ISSN 1939-2222, 0096-3445. doi: 10.1037/0096-3445.135.4.501.
- John Locke. *An Essay Concerning Human Understanding*, volume III, bk 3. Oxford ed.1894 edition, 1690.
- Peter Ludlow. Contextualism and the New Linguistic Turn in Epistemology. January 2005.
- Gary Lupyan. Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems are Penetrable Systems. *Review of Philosophy and Psychology*, 6(4):547–569, December 2015. ISSN 1878-5158, 1878-5166. doi: 10.1007/s13164-015-0253-4.
- William G. Lycan. *Philosophy of Language: A Contemporary Introduction*. Routledge Contemporary Introductions to Philosophy. Routledge, New York, 2nd ed edition, 2008. ISBN 978-0-415-95751-9 978-0-415-95752-6 978-0-203-93000-7.
- Michael P. Lynch. Alethic Functionalism and Our Folk Theory of Truth. *Synthese*, 145(1): 29–43, 2005. ISSN 0039-7857.
- Maar. *Vision*. 1982.
- Edouard Machery. Concepts Are Not a Natural Kind*. *Philosophy of Science*, 72(3):444–467, July 2005. ISSN 0031-8248, 1539-767X. doi: 10.1086/498473.
- Edouard Machery. Concept Eliminativism. In *Doing without Concepts*, chapter Doing without Concepts. Oxford University Press, 2009. ISBN 978-0-19-986795-0.
- Edouard Machery. Cognitive Penetrability: A No-Progress Report. In Athanassios Raftopoulos and John Zeimbekis, editors, *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, chapter The Cognitive Penetrability of Perception. Oxford University Press, 2015. ISBN 978-0-19-180210-2.
- Fiona Macpherson. Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism: COGNITIVE PENETRATION OF COLOUR EXPERIENCE. *Philosophy and Phenomenological Research*, 84(1):24–62, January 2012. ISSN 00318205. doi: 10.1111/j.1933-1592.2010.00481.x.
- Fiona Macpherson. The relationship between cognitive penetration and predictive coding. *Consciousness and Cognition*, 47:6–16, January 2017. ISSN 1053-8100. doi: 10.1016/j.concog.2016.04.001.
- Eric Mandelbaum. Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception. *Philosophy and Phenomenological Research*, February 2017. ISSN 00318205. doi: 10.1111/phpr.12368.
- M. G. F. Martin. The Limits of Self-Awareness. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 120(1/3):37–89, 2004. ISSN 0031-8116.
- Joseph McCaffrey and Edouard Machery. Philosophical issues about concepts: Philosophical issues about concepts. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(2):265–279, March 2012. ISSN 19395078. doi: 10.1002/wcs.1166.
- John Henry McDowell. *Mind and World: With a New Introduction*. Harvard University Press, Cambridge, Mass, 1st harvard university press paperback ed edition, 1996. ISBN 978-0-674-57610-0.

- P. M. Merikle, D. Smilek, and J. D. Eastwood. Perception without awareness: Perspectives from cognitive psychology. *Cognition*, 79(1-2):115–134, April 2001. ISSN 0010-0277. doi: 10.1016/s0010-0277(00)00126-8.
- G. E. Moore. The Refutation of Idealism. *Mind*, 12(48):433–453, 1903. ISSN 0026-4423.
- George Edward Moore. A Defence of Common Sense. In J. H. Muirhead, editor, *Contemporary British Philosophy, Second Series*. George Allen and Unwin, 1925.
- Blake Myers-Schulz and Eric Schwitzgebel. Knowing That P without Believing That P. *Noûs*, (47):371–384, April 2013. doi: 10.1111/nous.12022.
- Thomas Nagel. What Is It Like to Be a Bat? *The Philosophical Review*, 83(4):435–450, October 1974. ISSN 00318108. doi: 10.2307/2183914.
- Nancy J. Nersessian. *Creating Scientific Concepts*. MIT Press, Cambridge, Mass, 2008. ISBN 978-0-262-14105-5.
- Casey O’Callaghan. *A Multisensory Philosophy of Perception*. Oxford University Press, Oxford, United Kingdom ; New York, NY, United States of America, first edition edition, 2019. ISBN 978-0-19-883370-3.
- Ryan Ogilvie and Peter Carruthers. Opening Up Vision: The Case Against Encapsulation. *Review of Philosophy and Psychology*, 7(4):721–742, December 2016. ISSN 1878-5158, 1878-5166. doi: 10.1007/s13164-015-0294-8.
- Christopher Peacocke. *A Study of Concepts*. MIT Press, 1992.
- Christopher Peacocke. Phenomenology and Nonconceptual Content. *Philosophy and Phenomenological Research*, 62(3):609–615, 2001. doi: 10.1111/j.1933-1592.2001.tb00077.x.
- Ben Phillips. The Shifting Border Between Perception and Cognition. *Noûs*, 53(2):316–346, June 2019. ISSN 0029-4624, 1468-0068. doi: 10.1111/nous.12218.
- Ian Phillips. Unconscious Perception Reconsidered. *Analytic Philosophy*, December 2018. ISSN 2153-9596, 2153-960X. doi: 10.1111/phib.12135.
- Jean Piaget. *The Construction of Reality in the Child*. The Construction of Reality in the Child. Basic Books, New York, NY, US, 1954. doi: 10.1037/11168-000.
- David Plunkett. Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy. *Inquiry*, 58(7-8):828–874, November 2015. ISSN 0020-174X, 1502-3923. doi: 10.1080/0020174X.2015.1080184.
- Karl R. Popper. *The Open Universe: An Argument for Indeterminism*. Psychology Press, 1995. ISBN 978-0-415-07865-8.
- Mary C. Potter, Brad Wyble, Carl Erick Haggmann, and Emily S. McCourt. Detecting meaning in RSVP at 13 ms per picture. *Attention, Perception, & Psychophysics*, 76(2): 270–279, February 2014. ISSN 1943-3921, 1943-393X. doi: 10.3758/s13414-013-0605-z.
- Jesse J. Prinz. Is the Mind Really Modular? In Robert J. Stainton, editor, *Contemporary Debates in Cognitive Science*, pages 22–36. Blackwell, 2006.

- Jesse J. Prinz. When Is Perception Conscious? In Bence Nanay, editor, *Perceiving the World*, pages 310–332. Oxford University Press, September 2010. ISBN 978-0-19-538619-6. doi: 10.1093/acprof:oso/9780195386196.003.0011.
- Daria Proklova, Daniel Kaiser, and Marius V. Peelen. Disentangling Representations of Object Shape and Object Category in Human Visual Cortex: The Animate–Inanimate Distinction. *Journal of Cognitive Neuroscience*, 28(5):680–692, January 2016. ISSN 0898-929X. doi: 10.1162/jocn_a00924.
- Philip Pullman. *His Dark Materials*. Scholastic, London, first edition, October 2012. ISBN 978-1-4071-3559-5.
- Hilary Putnam. The Meaning of "Meaning". 1975a.
- Hilary Putnam. *Mind, Language, and Reality*. Cambridge University Press, 1975b.
- Zenon Pylyshyn. Is vision continuous with cognition?: The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3):341–365, June 1999. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X99002022.
- Jake Quilty-Dunn. Perceptual Pluralism. *Noûs*, 2019. ISSN 1468-0068. doi: 10.1111/nous.12285.
- W. V. Quine. Two Dogmas of Empiricism. *The Philosophical Review*, 60(1):20–43, 1951. ISSN 0031-8108. doi: 10.2307/2181906.
- W. V. Quine. Natural Kinds. In Nicholas Rescher, editor, *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of His Sixty-Fifth Birthday*, Synthese Library, pages 5–23. Springer Netherlands, Dordrecht, 1969. ISBN 978-94-017-1466-2. doi: 10.1007/978-94-017-1466-2-2.
- Athanassios Raftopoulos. Reentrant Neural Pathways and the Theory-Ladenness of Perception. *Philosophy of Science*, 68(S3):S187–S199, September 2001. ISSN 0031-8248, 1539-767X. doi: 10.1086/392908.
- Athanassios Raftopoulos. *Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy?* MIT Press, Cambridge, UNITED STATES, 2009. ISBN 978-0-262-25914-9.
- Athanassios Raftopoulos. The cognitive impenetrability of the content of early vision is a necessary and sufficient condition for purely nonconceptual content. *Philosophical Psychology*, 27(5):601–620, 2014.
- Athanassios Raftopoulos. Abductive Inference in Late Vision. In Lorenzo Magnani, Ping Li, and Woosuk Park, editors, *Philosophy and Cognitive Science II*, Studies in Applied Philosophy, Epistemology and Rational Ethics, pages 155–176. Springer International Publishing, 2015. ISBN 978-3-319-18479-1.
- Athanassios Raftopoulos and John Zeimbekis. The Cognitive Penetrability of Perception: An Overview. In *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, chapter The Cognitive Penetrability of Perception. Oxford University Press, 2015. ISBN 978-0-19-180210-2.

- V. S. Ramachandran and E. M. Hubbard. Psychophysical investigations into the neural basis of synaesthesia. *Proceedings. Biological Sciences*, 268(1470):979–983, May 2001. ISSN 0962-8452. doi: 10.1098/rspb.2000.1576.
- Antti Revonsuo. Binding and the Phenomenal Unity of Consciousness. *Consciousness and Cognition*, 8(2):173–185, June 1999. ISSN 10538100. doi: 10.1006/ccog.1999.0384.
- Mark Richard. *Meanings as Species*. Oxford University Press, Oxford, New York, July 2019. ISBN 978-0-19-884281-1.
- Martin Rolfs and Michael Dambacher. What draws the line between perception and cognition? *Behavioral and Brain Sciences*, 39, 2016. ISSN 0140-525X, 1469-1825. doi: 10.1017/S0140525X15002617.
- Eleanor Rosch. *Principles of Categorization*. Foundations of Cognitive Psychology: Core Readings. MIT Press, Cambridge, MA, US, 2002. ISBN 978-0-262-12247-4.
- Alexander Rosenberg. *Instrumental Biology, or, The Disunity of Science*. Science and Its Conceptual Foundations. University of Chicago Press, Chicago, 1994. ISBN 978-0-226-72725-7 978-0-226-72726-4.
- Susanna Schellenberg. *The Unity of Perception*. Oxford University Press, New York, NY, new product edition edition, 2018. ISBN 978-0-19-882770-2.
- Andreas Schierwagen. Brain Complexity: Analysis, Models and Limits of Understanding. In José Mira, José Manuel Ferrández, José R. Álvarez, Félix de la Paz, and F. Javier Toledo, editors, *Methods and Models in Artificial and Natural Computation. A Homage to Professor Mira's Scientific Legacy*, volume 5601, pages 195–204. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-642-02263-0 978-3-642-02264-7. doi: 10.1007/978-3-642-02264-7_21.
- Andreas Schierwagen. On Reverse Engineering in the Cognitive and Brain Sciences. *arXiv:1201.4896 [nlin, q-bio]*, January 2012.
- Brian J. Scholl and Tao Gao. Perceiving animacy and intentionality: Visual processing or higher-level judgment? In *Social Perception: Detection and Interpretation of Animacy, Agency, and Intention*, pages 197–229. MIT Press, Cambridge, MA, US, 2013. ISBN 978-0-262-01927-9. doi: 10.7551/mitpress/9780262019279.003.0009.
- Nicholas Shea. Distinguishing Top-Down from Bottom-Up Effects. In Dustin Stokes, Mohan Matthen, and Stephen Biggs, editors, *Perception and Its Modalities*, pages 73–92. Oxford University Press, October 2014. ISBN 978-0-19-983279-8. doi: 10.1093/acprof:oso/9780199832798.003.0004.
- Theodore Sider. *Writing the Book of the World*. Oxford University Press, Oxford : New York, November 2011. ISBN 978-0-19-173209-6.
- Theodore Sider. *The Tools of Metaphysics and the Metaphysics of Science*. Oxford University Press, Oxford, New York, February 2020. ISBN 978-0-19-881156-5.
- Susanna Siegel. Which Properties are Represented in Perception? In Tamar Szabo Gendler and John Hawthorne, editors, *Perceptual Experience*, pages 481–503. Oxford University Press, January 2006. ISBN 978-0-19-928976-9. doi: 10.1093/acprof:oso/9780199289769.003.0015.

- Susanna Siegel. *The Rationality of Perception*. Oxford University Press, Oxford, New York, January 2017. ISBN 978-0-19-879708-1.
- Susanna Siegel and Alex Byrne. Rich or Thin? In Bence Nanay, editor, *Current Controversies in Philosophy of Perception*. New York, USA: Routledge, 2017.
- Herbert Alexander Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, Mass., 3. ed., [nachdr.] edition, 1969. ISBN 978-0-262-19374-0 978-0-262-69191-8.
- Matthew H. Slater and Andrea Borghini. Introduction: Lessons from the Scientific Butchery. In Joseph Keim Campbell, Michael O'Rourke, and Matthew H. Slater, editors, *Carving Nature at Its Joints*, pages 1–32. The MIT Press, October 2011. ISBN 978-0-262-01593-6. doi: 10.7551/mitpress/9780262015936.003.0001.
- Joulia Smortchkova and Michael Murez. Representational Kinds. In *Mental Representations: The Foundations of Cognitive Science?* forthcoming.
- Jeff Speaks. Is There a Problem About Nonconceptual Content? *Philosophical Review*, 114 (3):359–98, 2005.
- Elizabeth S. Spelke. Where Perceiving Ends and Thinking Begins: The Apprehension of Objects in Infancy. In A Yonas, editor, *Perceptual Development in Infancy. The Minnesota Symposia on Child Psychology*, volume 20, pages 187–233. Lawrence Erlbaum, Assoc., Hillsdale, NJ, 1988. doi: 10.4324/9780203771143-9.
- Elizabeth S. Spelke. Core knowledge. *American Psychologist*, 55(11):1233–1243, 2000. ISSN 1935-990X(Electronic),0003-066X(Print). doi: 10.1037/0003-066X.55.11.1233.
- George Sperling. The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11):1–29, 1960. ISSN 0096-9753(Print). doi: 10.1037/h0093759.
- Peter F. Strawson. Carnap's Views on Conceptual Systems Versus Natural Languages in Analytic Philosophy. In Paul Arthur Schilpp, editor, *The Philosophy of Rudolf Carnap*, pages 503–518. Open Court: La Salle, 1963.
- Brent Strickland and Brian J. Scholl. Visual perception involves event-type representations: The case of containment versus occlusion. *Journal of Experimental Psychology: General*, 144(3):570–580, 2015. ISSN 1939-2222(Electronic),0096-3445(Print). doi: 10.1037/a0037750.
- Henry Taylor and Peter Vickers. Conceptual fragmentation and the rise of eliminativism. *European Journal for Philosophy of Science*, 7(1):17–40, January 2017. ISSN 1879-4912, 1879-4920. doi: 10.1007/s13194-016-0136-2.
- Simon Thorpe, Denis Fize, and Catherine Marlot. Speed of processing in the human visual system. *Nature*, 381(6582):520–522, June 1996. ISSN 0028-0836, 1476-4687. doi: 10.1038/381520a0.
- A. Treisman. Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 353(1373):1295–1306, August 1998. ISSN 0962-8436. doi: 10.1098/rstb.1998.0284.

- Michael Tye. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Representation and Mind. MIT Press, Cambridge, Mass., 1995. ISBN 978-0-262-70064-1 978-0-262-20103-2.
- Michael Tye. Representationalism: The Theory and Its Motivations. In *Consciousness, Color, and Content*, pages 45–68. MIT Press, 2000. ISBN 978-0-262-70088-7.
- Raphael van Riel and Robert Van Gulick. Scientific Reduction. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2019 edition, 2019.
- John Paul Wilson, Kurt Hugenberg, and Nicholas O. Rule. Racial bias in judgments of physical size and formidability: From size to threat. *Journal of Personality and Social Psychology*, 113(1):59–80, July 2017. ISSN 1939-1315, 0022-3514. doi: 10.1037/pspi0000092.
- William C. Wimsatt. The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets. *Canadian Journal of Philosophy Supplementary Volume*, 20: 207–274, 1994. ISSN 0229-7051, 2633-0490. doi: 10.1080/00455091.1994.10717400.
- William C. Wimsatt. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press, Cambridge, Mass, 2007. ISBN 978-0-674-01545-6.
- Rasmus Grønfeldt Winther. Parts and Theories in Compositional Biology. *Biology and Philosophy*, 21(4):471–499, September 2006. ISSN 1572-8404. doi: 10.1007/s10539-005-9002-x.

Appendices

Evaluation P_i with 10 items

The domain of evaluation is constituted by 10 items i_1, \dots, i_{10} , that are respectively retinal information, 2 1/2 sketch, 3D model, perceptual object representation, scene perception, event perception, representation of basic-level categories, representation of expert categories, perceptual demonstrative thoughts and paradigmatic beliefs.

The table of description of the domain is the following (reproduced from chapter 6, section 6.1.2):

	F_1	F_2	F_3	F_4	F_5
Objects	Conceptual	Iconic	Representational	Stimulus-driven	Hallucination
1. Retinal information	0	1	0	1	0
2. 2 1/2 Sketch	0	1	0	1	0
3. 3D Model	0	1	1	1	0
4. Perceptual Object Rep.	0	0	1	1	1
5. Scene perception	0	0	1	1	1
6. Event perception	0	0	1	1	1
7. Basic-level cat.	1	0	1	1	1
8. Expert categories	1	0	1	1	1
9. Dem. thoughts	1	0	1	1	1
10. Beliefs	1	0	1	0	1

Figure 1: table of description k-PERCEPTION

This domain of evaluation is classified into 3 categories SENSATION, PERCEPTION and COGNITION.

I evaluate 4 partitions:

- $P_1 = \text{SENSATION} = i_1, i_2$; $\text{PERCEPTION} = i_3, \dots, i_7$; $\text{COGNITION} = i_8, i_9, i_{10}$
- $P_2 = \text{SENSATION} = i_1$; $\text{PERCEPTION} = i_2, i_3$; $\text{COGNITION} = i_4, \dots, i_{10}$
- $P_3 = \text{SENSATION} = i_1$; $\text{PERCEPTION} = i_2, \dots, i_6$; $\text{COGNITION} = i_7, \dots, i_{10}$
- $P_4 = \text{SENSATION} = i_1$; $\text{PERCEPTION} = i_2, \dots, i_8$; $\text{COGNITION} = i_9, i_{10}$

The measure of inclusiveness of a concept, noted $Incl(C)$, is the proportion of items of the domain included into a concept.

The homogeneity of a concept, noted $Hom(C)$, is the average score of its homogeneity with respect to each feature F_i such that $Hom(C) = \sum_1^n (Hom(C, F_i)) / n$

$Hom(C, F_i)$ is the highest score between either the proportion of items contained in C that possess F_i or the proportion of items contained in C that lack F_i . In order to scale this measure of homogeneity between 0 and 1, the result is multiplied by 2 and we take 1 away.

The concept utility of C , noted $U(C)$ is given by: $U(C) = Incl(C) \times zHom(C)$ and the utility of a partition is the average utility of its concepts:

$$U(P_i) = (U(\text{SENSATION}) + U(\text{PERCEPTION}) + U(\text{COGNITION})) / 3$$

The results are the following:

P_1	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/5	1/5
PERCEPTION	3/5	3/5	1	1	3/5	19/25	1/2	19/50
COGNITION	1	1	1	1/3	1	13/15	3/10	13/50

Figure 2: Utility of concepts in P_1

P_2	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/10	1/10
PERCEPTION	1	1	0	1	1	4/5	1/5	4/25
COGNITION	1/7	1	1	5/7	1	27/35	7/10	27/50

Figure 3: Utility of concepts in P_2

P_3	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/10	1/10
PERCEPTION	1	1/5	3/5	1	1/5	3/5	1/2	3/10
COGNITION	1	1	1	1/2	1	9/10	2/5	9/25

Figure 4: Utility of concepts in P_3

P_4	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/10	1/10
PERCEPTION	3/7	3/7	5/7	1	3/7	3/5	7/10	21/50
COGNITION	1	1	1	0	1	4/5	1/5	4/25

Figure 5: Utility of concepts in P_4

	P_1	P_2	P_3	P_4
U	$7/25$	$4/15$	$19/75$	$17/75$

Figure 6: Partition Utility of P_1 , P_2 , P_3 and P_4

Evaluation P_i with 15 items

The domain of evaluation is extended from 10 to 15 items ("2 ½ sketch" is replaced by "orientation perception, edge perception, luminance perception, motion perception, shape perception and color perception"). The matrix of description is the following:

	F_1	F_2	F_3	F_4	F_5
Objects	Conceptual	Iconic	Representational	Stimulus-driven	Hallucination
1. Retinal information	0	1	0	1	0
2. Orientation perception	0	1	0	1	0
3. Edge perception	0	1	0	1	0
4. Luminance perception	0	1	0	1	0
5. Motion perception	0	1	0	1	0
6. Shape perception	0	1	0	1	0
7. Color perception	0	1	0	1	0
8. 3D Model	0	1	1	1	0
9. Perceptual Object Rep.	0	0	1	1	1
10. Scene perception	0	0	1	1	1
11. Event perception	0	0	1	1	1
12. Basic-level cat.	1	0	1	1	1
13. Expert categories	1	0	1	1	1
14. Dem. thoughts	1	0	1	1	1
15. Beliefs	1	0	1	0	1

Figure 7: table of description k-PERCEPTION with 15 items

Here are the results of partition utility for P_1 , P_2 , P_3 and P_4 , as defined in Appendix A:

P_1	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	7/15	7/15
PERCEPTION	3/5	3/5	1	1	3/5	19/25	1/3	19/75
COGNITION	1	1	1	1/3	1	13/15	1/5	13/75

Figure 8: Utility of concepts in P_1 with 15 items

P_2	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/15	1/15
PERCEPTION	1	1	0	1	1	4/5	7/15	28/75
COGNITION	1/7	1	1	5/7	1	27/35	7/15	9/25

Figure 9: Utility of concepts in P_2 with 15 items

P_3	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/15	1/15
PERCEPTION	1	1/5	3/5	1	1/5	3/5	2/3	2/5
COGNITION	1	1	1	1/2	1	9/10	4/15	6/25

Figure 10: Utility of concepts in P_3 with 15 items

P_4	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1	1	1	1	1	1	1/15	1/15
PERCEPTION	3/7	3/7	5/7	1	3/7	3/5	4/5	12/25
COGNITION	1	1	1	0	1	4/5	2/15	8/75

Figure 11: Utility of concepts in P_4 with 15 items

	P_1	P_2	P_3	P_4
U	14/47	4/15	4/17	17/78

Figure 12: Partition Utility of P_1 , P_2 , P_3 and P_4

Utility random partition

I calculate the partition utility of the following random partition of the domain i_1, \dots, i_{10} :
 $P_{random} = \{\text{SENSATION} = \{i_3, i_5, i_6, i_8\}; \text{PERCEPTION} = \{i_2, i_7\}; \text{COGNITION} = \{i_1, i_4, i_9, i_{10}\}\}$

The results are the following:

P_{random}	Hom/ F_1	Hom/ F_2	Hom/ F_3	Hom/ F_4	Hom/ F_5	Hom	Incl	U
SENSATION	1/2	1/2	1	1	1/2	7/10	2/5	7/25
PERCEPTION	1/5	0	0	0	1	0	1/5	1/25
COGNITION	0	1/2	1/2	1/2	1/2	2/5	2/5	4/25

Figure 13: Utility of concepts in P_{random}

$$U(P_{random}) = 4/25$$

RÉSUMÉ

Les scientifiques tentent de découvrir les mécanismes mentaux qui nous permettent de voir, de sentir, de toucher, etc. Les théories sur la perception abondent et pourtant, la perception, en tant qu'objet scientifique, reste vague. Comment pouvons-nous savoir qu'un phénomène spécifique est un phénomène perceptuel ? Comment les scientifiques savent-ils qu'ils étudient la perception et non une autre activité de l'esprit ? Quelles sont les limites de la perception ?

Ces questions ont des implications directes dans la pratique scientifique en laboratoire, en particulier dans l'interprétation des résultats expérimentaux. Elles sont également au cœur de plusieurs débats philosophiques tels que la pénétrabilité cognitive de la perception. L'objectif de cette thèse est de contribuer au débat sur les frontières de la perception par une étude du concept de perception au sein des sciences cognitives.

J'explore en premier lieu la stratégie utilisée habituellement pour déterminer les limites de la perception, qui repose sur le présupposé que la perception constitue une espèce naturelle d'états mentaux. Dans un second temps, j'élabore une stratégie appelée « ingénierie conceptuelle ». J'y étudie la perception en tant que *concept scientifique*. Je montre qu'il existe quatre concepts différents de perception. Je défends que ces quatre concepts sont légitimes. Ce pluralisme conceptuel s'inscrit dans un processus de « fragmentation conceptuelle », processus courant lors de la maturation d'un domaine scientifique.

Dans une troisième partie, j'étudie les fondements métaphysiques des limites de la perception. Je montre qu'il existe plusieurs manières de concevoir métaphysiquement les frontières de la perception, et que ce choix a une influence sur la manière dont les résultats empiriques sont interprétés.

Enfin, dans une dernière partie, je défends que le concept de perception joue aujourd'hui un rôle de « concept organisationnel », dont la fonction principale est de guider, coordonner et intégrer la recherche interdisciplinaire sur la perception. Reconnaître la place spécifique du concept de perception en science permet d'éviter les débats infructueux et d'enrichir les discussions scientifiques en présentant certaines théories comme complémentaires plutôt que rivales.

MOTS CLÉS

perception, concept scientifique, pluralisme, espèce naturelle, évaluation conceptuelle

ABSTRACT

Scientists try to uncover the mental mechanisms that allow us to see, feel, touch, etc. They are helped by philosophers, who elaborate characterizations of perception, of its function and nature. Yet, perception as a scientific object, is quite vague. In particular, perception is hard to circumscribe in the architecture of the mind.

The aim of this dissertation is to shed light on the kind of scientific object perception is, through a philosophical work based on analytical philosophy of mind, philosophy of science and on current philosophical developments concerning perception. How could we know that a specific phenomenon is a perceptual one? How does the scientist know that she is studying perception and not another activity of the mind? What are the boundaries of perception?

These questions have direct implications in everyday scientific practice, especially in the interpretation of experimental results. They also pertain to philosophical debates such as the cognitive penetrability of perception.

In this dissertation, I first explore the usual method for determining the boundaries of perception, based on the idea that perception is a natural kind of mental states. I then elaborate another strategy called "conceptual engineering". In this strategy, I study perception as a scientific concept. I show that there are at least four different concepts of perception. I argue in favor of conceptual pluralism, i.e. that these concepts of perception are all legitimate. This pluralism is part of a classical process commonly undergone by scientific concepts - conceptual fragmentation – in which a scientific concept is fragmented into several sub-concepts.

In a third part, I explore metaphysical foundations for the boundaries of perceptions. I show that there are several alternative ways to metaphysically ground the boundaries of perception. The choice between them influence how empirical results are interpreted.

Finally, I argue that the concept of perception should be today considered an organizational concept in cognitive science, whose main function is to guide, coordinate and integrate interdisciplinary research about perception. Recognizing this specific place of the concept perception in scientific and philosophical investigations about the mind would contribute to enrich discussions, as well as to avoid ill-posed questions and fruitless debates.

KEYWORDS

perception, conceptual engineering, concept utility, natural kind, scientific concept, mental boundaries