Perceptual engagement

a sensorimotor approach to phenomenal experience

Jan Degenaar

Jan Degenaar 2012
Perceptual engagement
a sensorimotor approach to phenomenal experience

Cover artwork: Josien Buiter

Printed by Wöhrmann Print Service, Zutphen, the Netherlands

ISBN: 978-90-367-5670-9
ISBN: 978-90-367-5671-6 (electronic version)

RIJKSUNIVERSITEIT GRONINGEN

Perceptual engagement

a sensorimotor approach to phenomenal experience

Proefschrift

ter verkrijging van het doctoraat in de
Wijsbegeerte
aan de Rijksuniversiteit Groningen
op gezag van de
Rector Magnificus, dr. E. Sterken,
in het openbaar te verdedigen op
donderdag 13 september 2012
om 11:00 uur

door

Jan Degenaar

geboren op 22 mei 1981
te Groningen

Promotores:	Prof. dr. A.J.M. Peijnenburg Prof. dr. E.W.E. Myin
Copromotor:	Dr. F.A. Keijzer
Beoordelingscommissie:	Prof. dr. T.A.F. Kuipers Prof. dr. J.K. O'Regan Prof. dr. M.V.P. Slors

UNIVERSITEIT ANTWERPEN

Perceptual engagement

a sensorimotor approach to phenomenal experience

Proefschrift voorgelegd tot het behalen van de graad van
doctor in de Wijsbegeerte aan de Universiteit van Antwerpen,
Faculteit Letteren en Wijsbegeerte

Jan Degenaar

Prof. dr. Jeanne Peijnenburg (promotor)
Prof. dr. Erik Myin (promotor)
Dr. Fred Keijzer (copromotor)

Antwerpen, 2012

Contents

Preface and acknowledgements

"The subject of consciousness is proving to be almost like a black hole to those who draw close to it. Once seduced inside the event horizon, they appear lost to normal scientific activity but follow a trajectory towards an explanation of the phenomenon which others, standing well away from the fateful edge, shout out to them is impossible to follow scientifically or is of little interest to those in the scientific mainstream. Those already lost to the black hole hear the cries of their companions but cannot escape the fatal attraction exerted on them. And as more and more fall into the black hole it expands, so swallowing more and more into its ever roomier interior. Will information ever get out about what the intrepid explorers have discovered? Or will there just be a gradual separation into those who have disappeared and those who resolutely turn their eyes from the glorious sight of infalling colleagues." (John G. Taylor, The Enchanting Subject of Consciousness.)

For quite some time now I have been intrigued by consciousness. And I am grateful to the people introducing me to the topic. If it were not for the great attractions of the topic of conscious experience, I might by now be enjoying the scenery of some tropical island, as an ecologist, as I used to think I would be. The fact that instead I focused my studies in Biology on the behavioral and cognitive neurosciences shows the mesmerizing pull of the subject of consciousness, and it reflects the satisfaction it gives me to study experience – and besides there's still plenty of time to enjoy the scenery.

People have worried, as it is stated in the quote above, that consciousness may be impossible to approach scientifically, or that the topic is of little interest to mainstream science. While I am now inclined to think that both worries reflect a hopelessly confused notion of consciousness, this is of course not to say that the study of consciousness is easy. I experienced the difficulty myself during a study that I undertook as part of my master program in biology.

The study concerned binocular rivalry: when you simultaneously present one picture to the one eye, and another to the other, you may experience the pictures alternatingly; there is perceptual rivalry between the pictures. This phenomenon is excellently suited for studying the neural processes involved in conscious experience, for the stimuli remain the same, while experience changes. Given that differences in the brain are not simply due to the presentation of a new stimulus, they seem particularly associated with consciousness. One may then ask, as indeed many studies

do, which changes in neural activity are correlated with the differences in conscious experience. What I tried to do was to find out to what extent these changes in neural activity are truly depending on conscious experience of the stimuli.[1]

One difficulty for a study like this is simply that one needs to get a method working. Before I could take a new step compared to what I found in the literature, I first had to replicate findings from previous studies (I tried to use so-called steady state visually evoked potentials in an EEG study). This turned out to be far from trivial. But I came to realize that research like this is facing a much more fundamental difficulty as well. The reason is that within such studies, a crucial issue remains untouched. Despite all the interesting findings I was hoping for, based on my study nothing substantial could be said about the *differences* between the experiences of the pictures. Even if my study would have gone smoothly (which it didn't), I simply could not have explained why one picture is experienced, say, as a house, the other as a face. The challenge remains to relate the differences in experience to differences in the correlating neural activity. We may establish correlations, but the trouble is to interpret these correlations. We may even ask ourselves: if we fail to explain how the differences in neural processes relate to the differences in experience, how can we claim to be studying conscious experience at all? The specific character of experiences seems to present crucial difficulties for understanding consciousness.

In more recent years I tried to take up this issue of the specific character of experiences. (Rather than focusing on experiences of houses or faces, however, I have primarily focused on more simple experiences, like the visual experience of a straight line or the distance between an object and its background – matters are already complicated enough here, and one has to start somewhere.) I came to believe that a crucial part of the challenge lies in finding a good conception of experience, and a good way to bring neural activity in view. It is not just a lack of detailed knowledge that gets us puzzled about conscious experience. To a serious extent the problem is how to make sense of the scientific facts.

As you may see in the chapters to come, there are radically different frameworks for thinking about experience. This is a PhD thesis in

[1] The reasoning was as follows: even if a subject is unaware of any rivalry between two simultaneously presented pictures, because you keep him or her occupied with some task, there may stil be 'neural rivalry'. If so, these changes in neural activity are certainly not sufficient for changes in experience, and they are not dependent on rivalry arising at the level of conscious experience either.

Philosophy, and my main concern here is to develop a fruitful way to think about experience. For this purpose I have explored a particular framework – exemplified in the sensorimotor approach – and the advantages it may have for explaining our perceptual experiences.

I learned a lot during the previous years, and there are many people to thank. Let me start by mentioning the people introducing me to the topic of consciousness, first especially Daniel Dennett through his books and later Fred Keijzer and Hans Dooremalen through their lectures in philosophy. This thesis has also been deeply influenced by the work of the late Susan Hurley, and I am grateful for the few stimulating conversations we had when this was still possible. Dennett's teacher, Gilbert Ryle, exerted his influence partly through the enthusing lectures of Charlotte Vrijen. Although the style of philosophy practiced in this thesis is very different from Ryle's, I hope I have remained critical towards some of the problems with which philosophers and scientists are concerned. Some questions should be dissolved by critical analysis, rather than solved by empirical conjectures (for example, the question how consciousness and matter may influence each other arguably presupposes an untenable notion of consciousness as a non-material entity).

Also the neuroscientific work during my studies in biology has greatly influenced my thought. I am grateful for the opportunity to choose my own path, supported by Jaap Koolhaas, as well as for the excellent supervision I have received: Ritske de Jong and Paolo Toffanin have taught me a lot during the study of visual consciousness mentioned above. At that time I collaborated with Marten Harbers, who used the same technique for different purposes; I learned a lot from him, both then and later. Also during my next study, an fMRI study on social perception, I have been fortunate with the excellent supervision by Valeria Gazzola. I hope that this thesis will be read as not just an exploration in the philosophy of mind, but as a contribution to the interpretation of neuroscience as well.

I am most grateful to the supervisors of my PhD project, Jeanne Peijnenburg, Fred Keijzer, and Erik Myin, for all their help; for their patience, confidence, criticism and encouragements. For the discussions as well as for the chat. They have greatly helped me to clarify my ideas and to improve my texts, and especially Jeanne enabled me to get rid of quite a few ambiguities that otherwise would have completely escaped my notice. I haven't been a fluent writer but I at least learned to eliminate many sketchy thoughts from my drafts and to work out the rest into more readable texts (although I'm afraid I tended to burden my supervisors with all too early versions). I also learned a lot from Fred and Erik, from all the comments they gave and ideas they expressed, and from working with them on two

different papers. I have been fortunate to have these supervisors, so many of whose interests I share. Where my outlook differed from Fred's, he encouraged me to find my own path, and also with Erik I could always work from my own fascinations. I wish I could have made even more use of their comments to improve the present work, but unfortunately the time for writing a PhD thesis is limited. I am grateful for all the help and support my dear supervisors gave me in the last four years. Without Jeanne Peijnenburg, Fred Keijzer, and Erik Myin, I might have accumulated a curious bunch of notes, but certainly I wouldn't have a book finished today.

Kevin O'Regan, Marc Slors, and Theo Kuipers have been so kind as to act as members of the reading committee. I thank them warmly for their willingness to assess my work. Kevin O'Regan I would also like to thank for discussions that helped me focus my thoughts, and for his advise on the color paper.

Most of the work was done at the Faculty of Philosophy of the University of Groningen, where I have been member of the Department of Theoretical Philosophy. This has been a great place to be. With the shared lunches and coffee breaks it has been a good social environment, hosting stimulating academic events, from work in progress sessions to workshops on embodied cognition. I enjoyed the company of my roommate Peter Timmerman and the frequent visits of Han Thomas Adriaenssen. Among the many dear people in the Faculty, Peter and Han Tomas have been particularly important to me for trying out ideas, for allowing me to voice the occasional frustrations, and for discussions on a wide range of topics. I may also mention Allard, Jan-Willem, Barteld, Frank, and others with whom I went to have so many lunches or dinners.

I profited much from the work in progress sessions of the PCCP and the Graduate School of Philosophy in Groningen. I thank my colleagues for the many helpful comments on earlier versions of my papers and talks. Also from the scientists and philosophers I met at conferences I learned a lot. Among them I particularly thank Sanneke de Haan, who gave valuable comments on parts of this thesis. It has been stimulating to find that so many people share my interests, to be able to learn from them, and to find out that many people actually know exactly what I am talking about. (It is true that the topic of consciousness may meet with the occasional glazed looks – the attempt to really understand experience is not everybody's cup of tea. But few people are contented with knowing which parts of the brain are active when you see, or with the assertion that you experience what you do because of the things that are out there.)

At present I work at the Center for Philosophical Psychology of the Department of Philosophy at the University of Antwerp, where part of this PhD thesis has been written. I thank Erik for having me here. I have been

enjoying our discussions over lunch, with Karim Zahidi and others, and I look forward to the coming period.

Finally, let me warmly thank my friends and family. My parents, Annelies van Gijsen and Joost Degenaar deserve special mention, as do the other family members, my sister Froukje, Reindert, Wiebe, Miep, Guido, Henriëtte, and my brother Joris. Others I like to thank are Maarten, Jorrit, Josien, Marten and Lisette, Jan, Maja, Gerda, and Manfred, Yasmijn, and Roberto. I consider myself fortunate knowing that there are places where I am always welcome, to enjoy the company, the conversation, the food, the scenery, together with these great people who appreciate a good experience.

Parts of this thesis have been published or submitted for publication elsewhere:

Degenaar, Jan; Keijzer, Fred (2009) 'Workspace and sensorimotor theories: complementary approaches to experience', *Journal of Consciousness Studies*, 16 (9), pp. 77-102.

Degenaar, Jan (*forthcoming*) 'Perception from the phenomenal stance', *Logique et Analyse*.

Degenaar, Jan; Myin, Erik (*submitted*) 'The structure of color experience and the existence of surface colors'.

Introduction

From our own experience we know what it is like to see colors, to feel things, or to hear sounds. Compare for example the experience of a red object with the experience of a green object: for someone with normal color vision, these experiences will clearly be different. The experiences have their own specific character. One of the fundamental aims for the sciences and philosophy of mind is to understand experience as part of the natural world. But there are various ways to consider the facts, and the place of experience in a scientific view of the world is far from obvious.

In an attempt to explain the differences in the character of experiences, we could consider the different underlying processes. In the case of color experience, we may then consider how objects reflect light, how the light stimulates photoreceptors in the eyes, and how this results in changes in the activity of the brain. But not all details of these processes will be equally important. For example, the details of the workings of the photosensitive pigments could presumably have been different without making any difference at the level of experience. Also the location of the relevant neural processes might have been different, and perhaps even the specific type of neurons involved does not matter. The question then is what the differences are that make a difference. How can we explain the specific character of our experiences?

In this Ph.D. thesis I aim to contribute to the understanding of perceptual experience, and more in particular, the *phenomenal character* of perceptual experiences. This means that emphasis lays on what the experience is like for the person, independent of the question whether the experience is veridical or not. I shall further set aside affective aspects of experiences as well as cognitive associations: whether you prefer red or green, whether looking at these colors influences your mood or not, and whatever associations you may have with the colors, I assume that there are basic differences between the experience of red and the experience of green. The question is how we can explain such basic differences.

I shall argue that a *sensorimotor approach* helps to understand the phenomenal character of perceptual experiences. To understand what it is like to see that an object is located at your right side, on this approach, we should not focus exclusively on the facts of sensory stimulation, such as the fact that the left side of your retinae is stimulated. Under normal circumstances these facts may of course correlate with the character of experience. But they do not suffice to provide an explanation. This can be demonstrated by altering the sensory stimulation, for example by wearing

glasses that invert the light falling on your eyes – in Chapter 6 I report on my experiences with left/right inverting glasses. As one adapts to the new situation, one may once again see objects that are at the right side as being at the right side, despite the fact that the retinal stimulation is inverted. The reason for this is that experiences depend on bodily action as well as on the stimulation of the sense organs. Sensorimotor approaches take this seriously and therefore focus on the ways in which sensory stimulation depends on the bodily movements of a perceiver in an environment (e.g. Kevin O'Regan & Alva Noë 2001). In this thesis I aim to explicate how perceptual experiences may be characterized and explained in terms of such sensorimotor patterns.

A recurrent theme throughout this thesis concerns the role of *action* in perception. As I shall explain in Chapter 2, there are several ways in which perceptual experience is crucially action-dependent, and Chapter 6 shows in detail how an accurate description of spatial vision must appeal to bodily action. But let me stress at the outset that I do not propose to reduce perceptual experience to action: being dependent on action is not being constituted by action. Perception and action are deeply interdependent, as Susan Hurley (1998) proposed, and neither reduces to the other. The reason why I stress this here is that some of the remarks of sensorimotor theorists are easily read as promoting an account in which perception is reduced to action. For example, Noë (2004) starts by stating: "The main idea of this book is that perceiving is a way of acting. Perception is not something that happens to us, or in us. It is something we do" (Noë 2004, p. 1). In this thesis I shall explicate, develop, and defend the idea that the phenomenal character of perceptual experience can be understood in terms of sensorimotor engagement with the environment. But, as I explain in Chapter 2, for this sensorimotor engagement no overt bodily movement is required.[1]

A second recurrent theme throughout this thesis concerns the relation of the phenomenal character of experience to underlying processes at the so-called *subpersonal* level of description. The type of account of perceptual experience in which I am interested must do more than describing our

[1] While perceivers may deliberately engage in active perceptual exploration, there often isn't much that we have to do in order to perceive. While I shall consider perceptual experience as exercising a skill, I shall not be explicitly concerned with the age-old dispute on the 'active' or 'passive' nature of perception. Perhaps to some extent this dispute may be sidestepped by acknowledging the 'active' contribution of organismic factors to experience, while recognizing the sometimes 'passive' nature of perception at the level of the person.

perceptual experiences. It must be able to link descriptions of our experiences to descriptions of the processes underlying the experience. The challenge we are facing here is nicely captured in the words of Gilbert Ryle: it is the challenge to avoid both the mistake of 'Nothing But' and the mistake of 'Something Else as Well'.[2] An account consisting of an enumeration of nothing but bodily movements or low-level neural or physical processes simply will not do, for it would fail to even address the phenomenal character of our experience; it leaves out what it is like (Levine 1993). At the same time, an account postulating something else as well will not do in as far as this 'something else' introduces a mysterious or flawed notion of experience, which fails to square with a scientific view of the world (Dennett 1991). An account of the phenomenal character of experience should show, at least in outline, how experience fits into the natural world. In Chapter 1 I shall introduce a strategy for avoiding the nonstarter of proposing that there is 'nothing but' a subpersonal story to be told, as well as the equally problematic proposal that there must be 'something else as well' parallel to the subpersonal story.

I think that the above-mentioned recurrent themes, while certainly sufficiently important to warrant emphasis, should not be too controversial among cognitive scientists and scientifically-oriented philosophers of mind. It is true that, when casually reflecting on perceptual experience, the role of action may easily be overlooked, and that one may too easily conflate personal and subpersonal levels of description. But any scientific account of perceptual experience should acknowledge the crucial role of action, and any satisfactory explanation of perceptual experience should relate subpersonal processes to the personal-level phenomena of interest. What will be probably more controversial in the account I defend is the specific way in which subpersonal processes are brought into view. As I shall make more precise in the chapters to come, the sensorimotor approach conceives of perceptual experience as a skillful *mode of engagement* with the environment, which does not consist in the possession of *inner models* of the environment.

This latter claim goes against a rich and perhaps intuitively appealing tradition of thinking about experience. According to this traditional view,

---

[2] Ryle used these terms in a somewhat different context, namely in explaining his long-range objective to "talk sense about the thinking that *Le Penseur* is occupied in doing" without committing the mistake of behaviorism and the mistake of Cartesianism (Ryle 1979, p. 17). Ryle argued that both the behaviorist/reductionist and the Cartesian/duplicationist commit a *category mistake*. Category mistakes and a way to avoid them are discussed in Chapter 1 below.

for example, vision and mental imagery are thought of in terms of the possession of something like an 'image' or representation in the head, or a 'movie in the brain'. If the account defended in this thesis is right, there need not be such an inner model. Perceptual systems do not have to construct inner models, and a whole layer can be eliminated from the explanation of perceptual experience. Thus the sensorimotor account offers not only an alternative way of thinking about perceptual experience; it offers an alternative framework for thinking about its underlying processes as well. The sensorimotor account is certainly not the only account aiming to find alternatives for 'inner model'-based views. In fact, it has been argued that much of present-day cognitive science is already turning away from inner models, often tacitly and unrecognized (see Ramsey 2007), sometimes explicitly (e.g. Chemero 2009). I think that the main contribution of the sensorimotor approach consists in drawing the consequences of the rejection of model-based accounts for our understanding of the phenomenal character of experience. In this thesis I shall argue that the skill-oriented perspective, as fleshed out in the sensorimotor account, has important advantages for aligning descriptions of experience with descriptions of its underlying processes.

The first two chapters introduce and explain the fundamentals of the sensorimotor account of the phenomenal character of perceptual experience, part of which is worked out in further detail in the third chapter. Chapter 1 introduces the explanatory challenge posed by the phenomenal character of experience, and it presents a line of reasoning leading to a skill-oriented approach. Following Dennett I advocate a formulation of the challenge as one of relating personal and subpersonal levels of description. Difficulties for inner model-oriented accounts are pointed out, which may be avoided by a skill-oriented perspective. Chapter 2 presents the skill-oriented view as developed in the sensorimotor account. This chapter discusses the ways in which perception is action-dependent, supporting a focus on sensorimotor patterns for explaining the character of experience. The structure of the sensorimotor account is explicated, with special attention to the way in which the account relates to subpersonal levels of description, and neural processes in particular. Chapter 3 discusses the role of neural processes within an account focused on dynamic patterns of engagement with the environment. For this, the sensorimotor account is investigated in relation to influential 'neural workspace' theories. It is argued that sensorimotor accounts and neural workspace accounts have complementary strengths and weaknesses, and that difficulties of the individual accounts can be overcome by taking them in combination.

In the next three chapters, the developed perspective is applied to various discussions in the philosophy of mind, and further support is presented for a sensorimotor account. Chapter 4 explores two ways to analyze perceptual experience: one focused at the phenomenal character of experience, the other focused at the (real or apparent) objects of experience. It is argued that the phenomenal character of experience should not be thought of in terms of the objects of perception, but rather in terms of the perceiver's mode of engagement, as proposed by the defended sensorimotor account. Crucially, this chapter argues that, while the phenomenal character of experience does not reduce to the objects of perceptual experience, this does not imply that 'something else as well' needs to be added to the account of the processes underlying perceptual experience. Chapter 5 applies the sensorimotor perspective to a philosophical discussion of color vision. Based on an interpretation of vision science, it has been argued that color experiences are neural phenomena without a genuine environmental counterpart. This argument is reconsidered and rejected on the basis of an analysis in which the interaction with the environment acquires a larger role than the argument has recognized. Chapter 6 offers an analysis of spatial vision, facilitated by my experience with wearing left/right inverting glasses. It is shown that pictorial descriptions of visual experience fail, and that a sensorimotor perspective can help to articulate the phenomenal character of experience. Also the distinction (discussed in Chapter 4) between the phenomenal character of experience and the object of experience is vindicated.

The last part concludes and provides a synthesis of the previous chapters. What is it that we are reflecting on, when we reflect on the phenomenal character of experience? And how can our experiences be explained? It is concluded that a non-trivial characterization of experience must link to non-experiential descriptions of the processes involved. For this purpose the sensorimotor account has several crucial advantages.

This thesis has been written as a collection of essays, so that the different chapters can be read independently. As a result, there is some overlap between the chapters. If you just wish to get a quick impression of the action-dependence of perception, or of the sensorimotor account, you can turn to Chapter 2. If you are interested in the 'explanatory gaps' of experience, or in the relation of the sensorimotor account to neural workspace accounts, you can read Chapter 3. Chapter 4 discusses the notion of a 'phenomenal stance', and it may be of particular interest for those who are intrigued by debates in the philosophy of mind regarding the 'qualitative character' of experience. If you are interested in philosophical ideas on color, in explanations of the structure of experience, or in the

explanatory status of neural correlates, you can pick out Chapter 5. For a report of perceptual adaptation to left/right inverting glasses, and for personal-level descriptions of spatial vision and visual imagery, the reader may turn to Chapter 6. The background of my formulation of the explanatory challenge is provided in Chapter 1.

Chapter 1
The phenomenal character of experience

This chapter introduces the challenge of explaining the phenomenal character of experience. Starting from the work of Ryle and Dennett, I shall formulate an explanatory ideal. Subsequently I discuss a difficulty for traditional approaches, which motivates a turn to an alternative framework for thinking about experience.

1. Introduction

Philosophers have always been fascinated by the workings of the mind. Since the twentieth century this fascination has often been naturalistically framed, in terms of the question how we can relate mental phenomena to the natural world. This development has greatly deepened and widened our insight in the mental domain. Believing, hoping, deciding, fear and joy, perception and deliberate action: by relating such phenomena to the rest of the natural world we increase our understanding of ourselves and of the world we live in.

One aspect of our mental lives that seems particularly hard to relate to the natural world is conscious experience. Conscious experience is not the same as experience *sec*. We can say of robots programmed to revise their behavior on the basis of past performances that they learn from experience and in that sense have experiences. However, it would be odd to ascribe *conscious* experiences to these robots. Present-day robots seem to lack something that we, humans and animals, are capable of. Like us, they can be said to react to stimuli, but we do not associate their reactions with conscious experience.

It is useful to distinguish two kinds of issues regarding conscious experience. First, there are issues concerning behavioral capacities associated with consciousness. When we are conscious of something, we can often act in accordance, for example by telling others what we perceive. Whether or not such capacities require or imply consciousness, it is widely acknowledged that at least in our case conscious experience is associated

with such behavioral capacities.[1] Second, there are issues concerning the content and the so-called phenomenal aspects of consciousness. The latter issues have caused the greatest difficulties. In this thesis I shall focus on the *phenomenal* aspects of conscious experience in particular. This means that the focus lies on what the experience is like for the person, independent of the question whether the experience is veridical or not, which would be a question concerning the content of experience.

Imagine that you experience the smell of a rose. Clearly, this experience differs from the smell of coffee, or the visual experience of the rose. But how can we account for these differences? The point here is not that we are able to behave differently towards different things. The point is that we subjectively experience things in a certain way. Experiences, such as the experience of smelling a rose, have a certain *phenomenal character*. Some use the term 'quale', or its plural 'qualia' to refer to this phenomenal or qualitative character of experience. Or, as it is often put, there is 'something it is like' to perceive, and what this is like differs from experience to experience. It seems that this qualitative character cannot be fully captured in terms of cognitive associations or affective appreciations: Whether or not you associate a particular smell with a rose or a cup of coffee, and whether or not you have any preference for one smell over the other, the smell of coffee is markedly different from the smell of a rose. The phenomenal character of experience is widely considered as one of the most puzzling aspects of the mind. How should we account for it?

The point should be carefully understood. It is not that roses actually smell like roses, or that some roses are in fact red. The point is that these truisms do not in themselves provide an explanation for the quality of our experience of the scent or color of roses. In fact, they are not much more than descriptions of the phenomena that a scientific account of experience tries to explain. The smell and the color are a matter of the environment as it is perceived by us; they are at what Gibson called the 'ecological' level of description (Gibson 1979; Thompson 1995). Properties of the outside world may play a role in explaining the character of experience, as does our

[1] Compare our capacities with the robot case. Clearly, there is much we can do that is still lacking even in the most advanced robots known today. It is not just that there is no human-level reflection and communication to be found in robots. It seems that artificial systems have not even matched the behavioral flexibility of insects, let alone that they match the flexibility of animals to which most of us will more confidently ascribe conscious experience. If an artificial system would have the whole range of behavioral capacities that we have, we would certainly be more prone to ascribe conscious experience to it.

physical makeup. However, when confronted with the same physical surroundings, newborn babies do not perceive colors in the same way as we do. Neither do dogs or birds for that matter. In the case of human babies, normal color vision still needs to develop, and in case of the dogs and birds, the different photosensitivity of the eyes will lead to differences in color vision, even if they are raised in a similar environment as humans. What this shows is this: the existence of objects that may appear red to grown-ups with normal color vision does not in itself explain the phenomenal quality of these adults' experience of red. The perceived properties of our environment do not in themselves explain our experiences.

What I aim for is an understanding of how experiencing works, *such that this explains the phenomenal character of experience*. Acquiring knowledge of the mechanisms of experience does not automatically yield an understanding of the phenomenal character of experience. Consider for example the processes involved in catching the smell of freshly made coffee. Hundreds of aromatic compounds have evaporated from the hot coffee, moving through the air until they reach your nose. As you breathe in, volatile substances diffuse at various speeds through a layer of mucus, reaching olfactory receptors which in turn evoke neural activity. This modulates neural activity in various parts of the brain; you may sniff and this may help to catch the smell. And at some point you consciously experience the coffee. Now it should be clear that, in order to account for the phenomenal character of your experience, a list of details concerning the molecular and cellular processes will not do. For all we know, the exact molecular mechanisms involved in the modulation of olfactory receptors may not even matter at the level of the person's experience: if the same aromatic molecules could be picked up by a different receptor with the same response functions, we would not expect a difference in experience. The challenge for an account of phenomenal experience is to explain which processes make a difference for experience and why.

In this chapter I aim to further clarify the challenge posed by the *phenomenal character* of experience, and I shall motivate the turn to the sensorimotor approach, which shall be discussed in detail the next chapter. I shall first explain how the philosophical analysis of Ryle (1949) has helped to relate mental phenomena to the natural world by placing the mental in the context of a person's skills, behaviors, and sensory encounters with the world (Section 2). Next I shall discuss how the work of Dennett (1991) can further contribute to our understanding by relating consciousness to processes at the so-called *subpersonal* level of underlying processes. But I shall also conclude that further work is needed to provide a satisfactory account of the phenomenal character of experience (Section 3).

We can sharpen the formulation of the challenge by drawing on Dennett's distinction between personal and subpersonal levels of description. Ideally an account of phenomenal experience offers characterizations with *dual currency*, as Humphrey (2000) has put it; such characterizations apply both to the personal level phenomenon of experience as well as to subpersonal processes (Section 4). I will discuss a common way to think of phenomenal experience, according to which experience is regarded as an *inner model* of the environment. This approach seems to face serious difficulties in aligning personal and subpersonal levels of description (Section 5). We shall therefore turn to an alternative framework, in which experience is viewed as *skillful engagement* with the environment (Section 6). An advantage of this framework for approaching the dual currency ideal is that its description of subpersonal processes remains closer to the way human beings experience the world. We shall see in Chapter 2 how the sensorimotor account offers a way to flesh out this perspective on the phenomenal character of experience.

2. Ryle on the mind: a critique on the 'ghost in the machine'

"There is a doctrine about the nature and place of minds which is so prevalent among theorists and even among laymen that it deserves to be described as the official theory. (…) It will be argued here that the central principles of the doctrine are unsound and conflict with the whole body of what we know about minds when we are not speculating about them." (Ryle 1949, p. 13.)

Thus starts Gilbert Ryle's influential masterpiece, *The Concept of Mind*, back in 1949. The doctrine Ryle referred to was dualism, according to which human beings consist of two very different things: a body and a mind. The body is supposed to be a spatial thing open to public observation; the mind is thought of as a non-spatial private realm, accessible only to the possessor of the mind. Mind and body are supposed to exist independently of one another, so that when the body dies, the mind may continue to function. Ryle associated this doctrine with the Cartesian picture of the mind, and he famously referred to it as 'the dogma of the Ghost in the Machine'. *The Concept of Mind* made a major effort to debunk this picture, which Ryle thought of as "a myth".[2]

[2] Although the picture under attack here is widely attributed to Descartes, I cannot do full justice to Descartes' views here. The main motivation for Descartes' dualism

The Cartesian dualist supposes that a process must be either a mental process, or a material process, so that a process can never be mental and material at the same time. Given that material processes are thought of as mechanical, the dualist presumes that a certain movement must either be caused by the mind, or else is caused mechanically. According to Ryle this is a logical mistake, or more precisely, a *category mistake*.

The idea of a category mistake can readily be understood by contrasting the following examples. It is perfectly reasonable to say: 'either he arrived there by plane, or he arrived by train (but not both)'. Planes and trains are mutually excluding forms of transport, and as such they are things of the same logical category. However, as Ryle points out, it is absurd to contrast 'she came home in a flood of tears' with 'she came home in a sedan-chair', as if these are necessarily opposing events. The reason is that in these sentences 'a flood of tears' and 'a sedan-chair' are not of the same logical category. To believe that 'coming home in a sedan-chair' must exclude 'coming home in a flood of tears' is to commit a category mistake.

When the Cartesian dualist asserts that mind and matter exclude each other – that something must be either mental or material – this presupposes that mind and matter are of the same logical category, like planes and trains in my example above. As Ryle puts it:

> "The belief that there is a polar opposition between Mind and Matter is the belief that they are terms of the same logical type." (Ryle 1949, p. 23)

Ryle argues that the concepts of mind and matter are not of the same logical type. If he is right, Cartesian dualism makes a category mistake. Dualism sketches a mistaken picture of mental phenomena, and in his analysis of the use of mental concepts Ryle attempts to correct this view.

2.1. Mental and behavioral

Consider someone giving a thoughtful speech. Does it make sense to set up a contrast between on the one hand the material occurrence, the 'mere behavior' of this person, and on the other hand the 'thoughtfulness' of his speech? Does the *thoughtful* nature of speech refer to something else than

---

appears to be intelligence or thought, not phenomenal experience. Descartes believed that intelligence, as found in humans but not in animals, could not be the result of allegedly mechanical bodily processes without interfering mind, so that something else seemed required as well. For ideas related to Ryle's critique, see also Wittgenstein (1953).

the speech itself? Is it something additional that accompanies the speaking, such as a process inside the person's mind? Of course not, says Ryle. The difference between thoughtful speech and babbling does not lie in the fact that the former is accompanied or caused by an inner thought process, whereas the latter is not. As Ryle explains, the thoughtful speaker is not doing two independent things: speaking and thinking, the one behavioral, the other mental. When we say that someone is speaking thoughtfully, we comment on a quality of the speech, not on a hypothetical process in the mind of the speaker.

If we forget this, Ryle argues, logical absurdities result. At first sight it seems natural to call an action 'intelligent' because, and only because, it is caused by intelligent thought. But on closer inspection we see that this approach confronts us with a vicious infinite regress. If we say that a performance is intelligent because it results from an intelligent operation, e.g. an operation in the mind, then we must also hold that the operation itself can only be called intelligent because it resulted from an intelligent procedure. And so on. Since this 'intellectualist legend', as Ryle calls it, leads to the bizarre consequence of an infinite regress, we should reject it:

> "The crucial objection to the intellectualist legend is this. The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle." (Ryle 1949, p. 31)

Intelligence cannot be dependent on intelligence. The question of course is: where to block the regress? If our only concern would be to avoid the regress, we could suppose that an action is intelligent because of certain inner operations, while at the same time denying that these inner operations require further operations in order to be intelligent (a position we may call 'intellectualism without regress'). True intelligence, we may then suppose, lies in the person's mind, while the intelligence of a speech is merely 'derived intelligence'. At first sight, it may seem that such a notion of inner intelligence would offer a way out.

However, a difficulty for this model arises when we ask ourselves how we should distinguish 'intelligent' inner operations from 'stupid' ones. The reason for making this distinction would supposedly lie in a judgment of the person's action. After all, we cannot see the person's inner acts in the way that we can see his ordinary actions. But how then could we say that the inner acts can by themselves be intelligent? It is in the context of interaction with other human beings that we learn to use mental predicates, and hypothetical inner processes cannot form the criteria for

the correct use of these concepts. By focusing on how we learn and use mental vocabulary, Ryle argues, we realize that the meaning of mental predicates such as 'intelligence' lies in their application to human activities and capacities; they are not descriptions of accompanying inner processes in an extra-material 'para-mechanical' mind.

By objecting to the 'para-mechanical' interpretation of mental concepts proposed by the Cartesian dualist, Ryle certainly does not propose to reduce mind to mechanism. Although his arguments are primarily directed towards a dualistic view of the mind, they equally apply to allegedly 'mechanical' physiological processes in the brain. As he writes: "If my arguments have any force, then these concepts have been misallocated in the same general way, though in opposing particular ways, by both the mechanists and the para-mechanists" (Ryle 1949, p. 310). By characterizing someone as intelligent we comment on this person's abilities or certain tendencies in the person's behavior. Whatever differences there are regarding the 'wires and pulleys' inside the head, they cannot distinguish between an intelligent performance and an unintelligent one.

When we comment on someone's chess-playing capacities, we are not referring to any accompanying shadow-operation preceding the act, such as a hypothetical mental act in which the person goes over pieces of knowledge in his or her head. In the words of Ryle, *knowing how* to do something, or having a *skill*, cannot be reduced to the possession of explicit *knowledge that* certain principles can be applied. Not only do we often know how to proceed without being able to formulate explicit rules that can be followed. Also, and more importantly, even when we do know the rules that can be applied, we can be better or worse in applying these rules. Thus the skillful performance of action cannot be fully understood by reference to the agent's explicit knowledge. It is not typically someone's knowledge that we refer to when we appreciate someone's mental capacities. So even if we were to think of knowledge as an inner state, much of our mental vocabulary would still concern the person's skills and performances.

2.2. Sensory experiences and imagery

In the introduction I distinguished between behavioral and phenomenal aspects of the mind (or, more precisely, behavioral and phenomenal aspects of conscious experience). Ryle's approach covers both aspects, and

it covers them in basically the same way.[3] On the behavioral side, as we have seen, Ryle attacks the 'intellectualistic legend', pointing out the trouble of an infinite regress. Similarly, Ryle argues that traditional approaches of for example tactile or visual experiences set off a vicious regress. As illustrated below, such a regress would impair the understanding of phenomenal aspects of the mind.

When we observe something, we may say this implies that we have sensory experiences (or 'sensations') evoked by the observed object. Such sensory experiences are paradigmatic examples of the things that populate what we, for want of a better term, could call the stream of consciousness. How should we conceive of this? What kind of 'thing' is a sensory experience? When we observe something, or when we experience after-images, it is sometimes said that we perceive or experience our sensations, e.g. that we experience a sensation of red. But as Ryle stresses, we should not think of sensory experiences as objects of observation, for that would again lead to an infinite regress:

> "If sensations are proper objects of observation, then observing them must carry with it the having of sensations of those sensations (…). And this is clearly absurd." (Ryle 1949, p. 197)

As Ryle rightly notes, it would be a serious mistake to suppose that consciously experiencing things involves inner acts of observing 'sensations'. For there is no observer in the head to do the observing, and sensations are not even the kind of things that can be objects of observation in the first place. Just as being alive is not the result of being inhabited by vital spirits, an action is not intelligent or free or conscious because it is caused by an inner act of intelligence or free choice or consciousness. Neither is perception conscious because of the conscious experience of 'perceptions'.

A similar point can be made concerning mental imagery. When picturing or visualizing a scene, Ryle writes, it may be tempting to suppose that there

---

[3] Of course the conceptual connections differ for the mental vocabulary concerning phenomenal and non-phenomenal aspects of our lives. For example, as Ryle has noted, there is a sense in which perception is not a process. Just as noon has no temporal extension, and the scoring of a goal is not a process in which at some point a goal is half-scored, there is a sense in which perception verbs declare a terminus: "At any given moment either I have not yet seen it or I have now seen it. The verb 'to see' does not signify an experience, i.e. something that I go through, am engaged in" (Ryle 1954, p. 103). Here I shall focus on a key point of Ryle's analysis which applies also to the phenomenal conception of perceptual experience.

must be something inside the mind that is being perceived, although not with one's eyes, but with the mind's eye. We may then suppose that what we perceive is not a picture or photograph, "but some counterpart to a photograph, only made of a different sort of stuff" (Ryle 1949, p. 234). But when we see something 'with our mind's eye', is there anything like such a picture in the mind? According to Ryle there isn't: "Much as stage murders do not have victims and are not murders, so seeing things in one's mind's eye does not involve either the existence of things seen or the occurrence of acts of seeing them"(Ryle 1949, p. 233).

Ryle contrasts the situation with seeing a real picture of a face or scene. Based on a picture of a friend's face, we may vividly 'see' or visualize the friend's face in detail. We can then speak of a lifelike picture. But this does not imply that the picture accurately replicates the lines and colors of the face, for the lifelike picture may be a cartoon rather dissimilar from an equally lifelike oil painting. The picture's being lifelike is a matter of the degree to which it helps the viewer to 'see' the depicted person, and this visualizing is not at all guaranteed by an accurate replica. We can then make a variation on the argument in the quote above: 'If visualizations are proper objects of sight, then seeing them vividly must carry with it the having of visualizations based on those visualizations. And this is clearly absurd.'

Of course this is not to deny that we are capable of mental imagery. It is to deny that mental imagery must involve (para)mechanical images, conceived as objects inside an entity called the mind. No picture is self-interpreting, and proposing that we have something like pictures in our heads cannot in itself explain what it is to interpret or perceive a picture in a certain way.

2.3. Upshot and remaining questions

Where does this bring us? I think Ryle is right to reject dualism, which contrasts mind and matter as if these belong to two different realms. One way in which the false opposition can be avoided is by appreciating that 'mental skills' cannot always be contrasted with 'behavioral skills'. Nothing about behavioral dispositions or capacities makes them opposed to truly cognitive, mental dispositions or capacities.

Through the notion of skills, or through the notion of disposition, we can see why mind and matter do not have to compete for validity, because behavioral skills or dispositions can be readily related to other parts of the natural domain. For example, if someone is exercising his or her mental capacities, and if these can be taken as behavioral capacities, clearly no violation is required of any of the natural laws governing the material

world (which are of course formulated without specific concern with mental phenomena). The reason is, as we have seen, that remarks about dispositions or skills are not of the same logical category as remarks about natural laws. We may compare this to the difference in chess between, on the one hand, remarks on defensive and offensive strategies and, on the other hand, remarks on the rules: these concern different categories, and clearly the strategies cannot conflict with the rules of the game. There can be no opposition between conforming to the rules on the one hand, and displaying a strategy on the other. Similarly, it makes no sense to say that a bodily movement either conforms to the laws of physics, or it is an intelligent move (but not both).

As far as phenomenal experience is concerned, Ryle's analysis can help to avoid not only a dualistic 'ghost in the machine' conception but also the Cartesian trap of infinite regress. The analysis of mental phenomena should avoid infinate regress, whether the analysis concerns phenomenal or non-phenomenal aspects of the mind. No inner acts of experiencing, and no experienced inner objects, are required for human beings to experience, for it is not in virtue of some inner operation that we can be said to experience. One of the lessons we can draw from Ryle's reflections is that if we are to understand mental concepts, we should take care not to jump to (para)mechanical conclusions about objects in minds. Intelligent action is not 'mere action' plus an additional mechanical or ghostly inner process of 'intelligence'. Nor is conscious experience a matter of 'mere physical responsiveness' to stimuli plus an additional mechanical or ghostly inner process of 'consciousness'.

At the same time, the question remains how we should relate the phenomenal character of experience to processes involved in the experience. Of course coffee and roses have a different scent, but what explains the phenomenal character of the scent of coffee, or the smell of a rose? Ryle's focus on what it means to have sensory experiences leaves this question unaddressed. We may agree with his analysis of mind-talk: "To talk of a person's mind is not to talk of a repository which is permitted to house objects that something called 'the physical world' is forbidden to house; it is to talk of the person's abilities, liabilities, and inclinations to do and undergo certain sorts of things, and of the doing and undergoing of these things in the ordinary world" (Ryle 1949, p. 190). But this does not provide us with a positive account connecting the way in which we are subjectively undergoing various things to the processes underlying these experiences.

It is important to note that Ryle's rejection of the dualist picture of the mind situates our mental capacities, dispositions, or attributes at the level of *persons*. People can be smart, careful, skilled and hot-tempered. People

can be glad or disappointed. And people – not inner processes – can calculate, write, solve puzzles, or tell stories. The persons that Ryle talks about are embodied beings that perceive and act, exercising their behavioral skills and using their sense organs and in a variety of different situations. But although his account of mental phenomena concerns embodied beings, embedded in an environment, not much is said about the processes at play within these persons, as they interact with their environment. In the next section I shall introduce neural processes within a broadly Rylean non-dualist view of the mind. We can then sharpen our understanding of conscious experience by relating the person's experience to the processes underlying the experience.

3. Dennett on consciousness: bringing the brain in view

An influential attempt to relate conscious experience to the natural world has been undertaken by Daniel Dennett. The general idea behind this attempt is that mind and matter are not different kinds of stuff, but that thinking about mind and thinking about matter involve taking different interpretative stances (Dennett 1987). This general idea sits well with a Rylean approach. For Ryle, too, denied that mind and matter are different things of the same type; and he, too, argued that thinking about mind and matter implies the application of different logical categories. Both philosophers reject Cartesian dualism. No more than Ryle does Dennett believe that mental attributes derive their meaning from reference to inner processes, and he joins Ryle in stressing that mental phenomena always occur at the level of whole persons. But there is also an important difference between the two. For Dennett does, whereas Ryle does not, bring the brain into the picture. By relating mental phenomena to neural processes, Dennett attempts to further integrate consciousness in a naturalistic framework.

3.1. Personal and subpersonal

Dennett draws a distinction between *personal* and *subpersonal* levels of description (Dennett 1969). I will argue that he is right in claiming that this distinction is vital when it comes to avoiding category mistakes of the kind outlined by Ryle, and that it can help to build a bridge between consciousness and matter. While I believe that Dennett hasn't yet provided a satisfactory account of the phenomenal character of experience, I hope to show that we can make use of his analysis to sharpen our view of the challenge. Indeed, as I shall argue in Section 4, the way to do justice to

conscious experience from within a scientific worldview is by finding the right way to bring subpersonal processes in view.

When we ask why a particular person performed a particular action we raise a question at the *personal level*. At this level, we explain a person's actions by referring to his or her reasons, e.g. the beliefs and desires that together make the action intelligible. At the personal level, the fact that someone has an experience of red can be explained by referring to the fact that the person is looking at a red object, that the lighting conditions are all right, and that the person has normal color vision. However, when we turn to lower levels of description, to *subpersonal levels*, there are no such things to be found as reasons, beliefs, and desires, and we are no longer speaking of a person looking at objects. At the subpersonal level, there are light waves impinging on the retina, there are neural impulses and there is synaptic activity – no wishes, thoughts, or hopes.

Before we go on, two terminological remarks are in order. First, the term 'personal level' is here used broadly, applying equally to human beings and animals. By saying that a fox sees a rabbit, we make a 'personal-level' remark about the fox. Second, the expression '*the* subpersonal level' should not be taken to mean that there is only one such level. This is clearly not the case. Sub-atomic, atomic, molecular and cellular levels are all subpersonal. Neural processes are often thought of in terms of information-processing or computational processes. But even such relatively high-level interpretations of neural processes must be considered as subpersonal, as they are not meant to imply that the person in which they occur is indeed computing. After all, the person in question might as well be playing Ping-Pong.

Dennett's personal/subpersonal distinction can help to block category mistakes. By clearly realizing at which level a particular phenomenon occurs, we may avoid getting entangled in confused questions that do not have real answers. Sensory experience, for example, is a typical phenomenon that occurs at a personal-level: only persons can have sensory experiences, and it would be a serious confusion to attribute sensory experiences to retinal or neurological processes. Of course, whenever there is a normal sensory experience there is such a thing as sensory stimulation. But the fact that for example visual experience and retinal stimulation usually go hand in hand, does not imply that the two phenomena are of the same level. We experience colors, and we do not experience retinal stimulation. Retinal stimulation may then form a subpersonal precondition for normal visual perception, but it is not itself a phenomenon at the level of the person, and we should not get ourselves involved in questions such as how we can experience our retinal image. Here is another example. Actions are often done for a reason, and whenever there is an action, there must of

course be neural processes involved. Reasons and neural processes thus often enter the scene jointly. However, this does not alter the fact that the two phenomena are at a different level of description. By recognizing this, we avoid category mistakes such as the mistake to suppose that a bodily movement must either have been done for a reason, or it must have a neurophysiological basis (but not both). By situating mental and material phenomena at their own level of description we can thus avoid a false polar opposition between mentalistic explanations and physiological explanations.

The point is that subpersonal processes, considered as such, do not themselves display the personal level phenomena of interest. For example, while we are subjects of experience, Dennett insists that at a subpersonal level no subject of experience is to be found. For if a subpersonal account of conscious experience were to have a subject of experience as an element of the account, explaining that we are subjects of experience would not even have been started – the problem would just have been shifted. Leaving out the subject from a subpersonal account therefore is a precondition for a good subpersonal analysis of the subject (Dennett 1991; 2001).

3.2. Consciousness and the brain

Dennett wants to do more than just give instructions on how to avoid confused questions and category mistakes. He also wants to know how the two levels of description hang together. How can we characterize the subpersonal processes underlying mental phenomena?[4] And more particularly, how can conscious experience be brought in connection with the physical processes inside the brain? A large part of Dennett's *Consciousness Explained* (1991) is devoted to precisely that question. He writes:

> "Somehow the brain must be the mind, but unless we can come to see in some detail how this is possible, our materialism will not explain consciousness, but only promise to explain it, some sweet day." (Dennett 1991, pp. 41-42)

[4] Here I shall focus exclusively on the relation between personal and subpersonal levels of description for the specific case of phenomenal experience. For other mental phenomena very different views on the personal-subpersonal relationship may be more appropriate. For Dennett's views on beliefs and other propositional attitudes, see Dennett (1991b).

Dennett's answer, in a nutshell, is that we should not focus on the intrinsic properties of neurons and neural processes. His way to explain how consciousness relates to subpersonal processes is by accommodating the *effects* or the *impact* of neurons. It is the contribution of neural activity to phenomena at the personal level, such as our ability to see, hear or feel, that makes the neural activity relevant. It is what neurons do for the person, and nothing else, that makes them relevant to consciousness.

At the personal level, we experience the world from a certain point of view and we can judge the approximate moment of experience. To properly see an object, it must be in clear sight; the light from the object must enter the eye, resulting in an appropriate neural response. But what is an appropriate neural response? And what is it about neural processes that makes that opening our eyes results in conscious experience? To address these questions, we will zoom in on the inner workings of the brain, at small spatial and temporal scales. The question then rises: must there, at the subpersonal level, be an exact time and place within the brain 'where consciousness happens'? Of course, if there is such a place, we want to know exactly where it is and what happens there. But Dennett famously argues that there need not be such a place. If he is right, this will throw considerable light on the subpersonal basis of experience.

Let us consider a central argument. This argument is framed in terms of the processing of *information* in the brain. Neural activity is interpreted as carrying information, e.g. about the color or shape of the objects in sight. This notion of information is a subpersonal notion: not all information has to play an active role for the behavior of the person, and the information need not be consciously experienced. The question may then be raised, what makes that some information is consciously experienced? One may suppose that it is the reaching of a certain privileged place in the brain – the place where consciousness happens – that makes the information available to the person: reaching this place is entering consciousness. Dennett's argument is meant to show that this idea has absurd consequences, and should therefore be rejected.

Consider a subject that is confronted with two lights which flash in turns. After a while, just as in a movie composed of different pictures, the subject will see one light moving between the two extreme positions where the two flashing lights are actually located. Dennett argues that consideration of this phenomenon shows the absurd consequences that follow from the hypothesis that there is a 'Cartesian theater', i.e. a place in the brain where neural activity reaches consciousness, such that the sequence in which information enters this place determines the sequence in which the information becomes conscious (Dennett 1991; Dennett & Kinsbourne 1992). Suppose, for the sake of argument, that there is such a

Cartesian theater. Now consider two scenarios for explaining the fact that the subject sees one moving light. According to one scenario, there is a neural delay mechanism such that only after the second light flashes, information concerning the first light enters the Cartesian theater, followed by information about the apparent moving, which in turn is followed by information concerning the second light. In this scenario, all experience is delayed until there is the experience of a moving light. According to the second scenario, there is a revision of experience: first information concerning one of the flashing lights enters the Cartesian theater, then the information concerning the second light, subsequently the information that there has been an in-between light *just before* the second light. Here, experience is revised in such a way that the subject would never know that experience has ever been different.

The crux is that the subject is unable to say which scenario is at work. Whatever differences between the scenarios, the subject cannot identify them, for – as far as the subject can tell – in both cases the same experience results. In this sense, the two scenarios differ in ways that remain forever inscrutable in terms of subjective experience. This raises the question: if subjective experience cannot differentiate between the Cartesian theater models, what can? For lack of other reasons to take the model seriously, Dennett and Kinsbourne conclude that there is no good reason to accept the Cartesian theater model and the differences it postulates (Dennett 1991; Dennett & Kinsbourne 1992). Thus they claim that there is no privileged place in the brain where information is, at the spot, transformed into consciousness.

The idea of differences in subjective experience that make no difference to the subject seems highly problematic indeed. Still one may wonder whether there may be evidence from neuroscience in support of the Cartesian theater model. But Dennett and Kinsbourne argue that in fact the evidence points in another way: there is no neuroscientific reason to believe that there is a subpersonal place where 'it all comes together' in a central arena of consciousness. Rather, there are 'multiple drafts', or parallel streams of information processing in the brain, which may even contain conflicting information. At a timescale of microseconds there need not be a fact of the matter about which of the streams underlies conscious experience. They conclude that "the brain itself is Headquarters, the place

where the ultimate observer is", and that there is no deeper headquarters within the brain (Dennett & Kinsbourne 1992).[5]

What does this mean for the processes underlying conscious experience? How should we conceive of these processes? Suppose that parallel streams of neural activity contain conflicting information, and let us assume that neuroscience can tell which stream contains what information (e.g. based on reliable correlations between neural activity and environmental features). It should then in principle be possible to identify the processes that are relevant to conscious experience, on condition of course that subjects can provide us with trustworthy reports of their subjective experiences. We would then be able to distinguish neural streams that *are* reflected in subjective experience from other neural streams that are *not*. This would raise the question: what is it about a neural process which makes that it contributes positively to conscious experience? In other words, what is special about the processes underlying conscious experience?

In *Consciousness Explained* Dennett argues it is not a particular intrinsic property that makes some neural streams special. Since there is no Cartesian theater, there is no need for peculiar intrinsic properties of the activity in this place. In fact, Dennett claims, intrinsic properties of neural activity are irrelevant to the person's conscious experience. Rather, it is the *effect* or *aftermath* of neural processes that plays a crucial role here. Dennett has put forward his original and illuminating metaphor of 'fame in the brain' to explain the point (e.g. Dennett 2001).[6]

Consider a famous speech. What is special about the speech, why can it be considered famous? Clearly, a speech can only be famous in virtue of what happens after the speech: people talk about it, refer to it, and perhaps

---

[5] This is not to say that there may be no specialized subsystem that plays a particularly important role in promoting consciousness, such as the subsystem proposed in Baars' 'global workspace model' (Baars 1988; discussed in Chapter 3 of this thesis). Such a subsystem may very well exist and facilitate consciousness (Dennett 2001). The point here is that the activity of such a subsystem takes time, and that it is not in virtue of 'entering' such a subsystem that information becomes conscious. Rather, it would be the role fulfilled by this subsystem – and the effect it has on further processes – that makes for its specific contribution to conscious experience.

[6] Another part of Dennett's answer has to do with language: he believes that language implements 'a serial virtual machine' on the parallel architecture of the brain. He argues that through language the kind of informational organization is achieved that is necessary for (our kind of) consciousness (Dennett 1991; 1995). It is beyond the scope of the present discussion to address these ideas.

even quote from it. A famous speech need not be particularly good, and an excellent speech need not become famous. And being famous is not an intrinsic property of the speech. It is a matter of what happens next. Just as a speech counts as famous only in virtue of what happens afterwards, neural activity counts as conscious activity (or more carefully, as activity underlying consciousness) only in virtue of its *aftermath*. What makes neural activity relevant to consciousness are not the intrinsic properties of the neural activity, it is its 'fame' that does it. Just as you cannot read off fame from a speech without considering its impact, one cannot read off consciousness from neural processes as if they mysteriously 'glow in the dark'. (And if the neural activity underlying consciousness were to have some curious intrinsic feature, this feature would matter only in as far as it contributes to further activity.) To use a different metaphor, consciousness is like political influence, or *clout*: "When processes compete for ongoing control of the body, the one with the greatest clout dominates the scene until a process with even greater clout displaces it" (Dennett 2001, p. 225). On this view, there is nothing special about the processes underlying conscious experience, except that they happen to have different effects on subsequent processes.

The 'fame in the brain' account is about potential control of the body, about personal skills and capacities, such as memorizing, planning, producing verbal reports, using information for carrying out actions, and so on. It is therefore about potentially extra-neural 'fame'. For on Dennett's view, the personal level is essential to theorizing about the mind. It is only because of its facilitation of personal level skills and behaviors that neural activity can count as the neural basis of consciousness.

To summarize, Dennett argues that not only there is no immaterial ghost in a material machine, but that it is also the case that no part of the organism is by virtue of its own intrinsic properties relevant for conscious experience.

3.3. Upshot and remaining questions

The upshot of Dennett's arguments is that there is no reason to suspect that anything out of the ordinary occurs in the subpersonal processes underlying consciousness. Still, even if the 'fame in the brain' model is right, it does not offer a full-blown account of conscious experience. For it is one thing to accept that no mysterious processes take place inside the brain. But it is quite a different thing to actually understand why we have the particular experiences we have, and why these differ from each other in the way they do. What explains the specific character of our experiences? Until

we can at least discern the outlines of an answer, we remain in the dark regarding the natural basis of conscious experience.

That Dennett's 'fame in the brain' model provides no account of the phenomenal character of experience can also be phrased in the following way. Why would the 'fame' of one neural process come with the experience of red, while the impact of another neural process is associated with the experience of blue, the experience of touch, or of the smell of a rose? It is hard to see how explanations in terms of *aftermath* can get hold here. We could distinguish the aftermath of a neural process from the aftermath of experience. At the subpersonal level, experience involves a neural process together with its aftermath. At the personal level, the aftermath of experience consists for example in reports of experience. As the personal level is the starting-point for his theorizing about the mind, Dennett's characterization of the subpersonal aftermath is grounded in the characterization of the aftermath of experience. Now it is true that the experience of something as being red may be connected to the possibility to say that something is red, or to otherwise behave in red-appropriate ways (which may be, say, the selection of a ripe apple based on its color). And it is also true that the particular saliency of a bright red object against a pale background may be partly understood at the subpersonal level in terms of the role played by the processes that are involved in the experience of the color. But something important is left out here. It is not the role played by consciousness or by the processes underlying consciousness that puzzles us. It is the specific 'feel' that comes with the fulfilling of this role. The challenge remains to explain why certain subpersonal processes come with a specific phenomenal character (such as the experience of red), while other processes come with a very different phenomenal character (such as the experience of green or of the smell of a rose).[7]

The possibility to provide an informational interpretation of neural activity does not fundamentally alter the situation. For the question remains why this proposed information should come with one experience rather than another. Differences between the proposed neural carriers of information will be discussed in Section 5 below (under the heading of 'inner models'). For now, note that, again, it seems that pointing at differences in aftermath does not do the trick. Consider for example the different names we use for colors or roses. The experienced smell of a rose

[7] This limitation of Dennett's functionalist account applies to functionalism more generally. Functionalism may be true and we may have excellent reasons for embracing it, but this does not in itself explain the specific character of experience. It just forms a theoretical background for the development of such an explanation.

would still need to be accounted for after we have accounted for the verbal identification of the smell. Of course there is much more to the aftermath than the subsequent use of words; there are many ways in which we may behave differently, and many ways in which subsequent thought may take a different course. But affirming the necessary existence of differences in aftermath does not in itself explain the differences in the phenomenal character of experience, as long as the link between the specific aftermath and the specific phenomenal character of experience remains obscure.

Dennett's criticism of the Cartesian theater model of consciousness is a valuable addition to Ryle's approach, for it brings subpersonal processes in view while avoiding category mistakes. His 'fame in the brain' model thereby helps to get a clear mind on the neural basis of conscious experience. But even if it does, we are still saddled with the question how we can relate the phenomenal character of experience to subpersonal processes.

4. The dual currency ideal

How should we account for the phenomenal character of experience? We are now in a position to formulate an explanatory ideal, drawing on the personal/subpersonal distinction discussed above.

The fact *that* subpersonal processes have a relation to conscious experience is obvious to any naturalistic philosopher of mind. The question is only *how* they are related. Clearly, the relation is not of a dualistic kind. In a dualistic model, neural or other subpersonal processes 'give rise to' or 'are linked with' *something else*, namely phenomenal experience. Non-dualistic approaches reject this view. They try to make intelligible how the personal level phenomenon of conscious experience is in agreement with subpersonal characteristics of the world, without assuming that the one causes or brings about the other. In what way should we then relate the description of the phenomenal character of experience to our understanding of the natural world?

The difficulty of our task is increased by the mundane fact that it is so hard to describe the phenomenal character of our experiences. How, for example, can we articulate what it is like to smell a rose? We could try to describe it by referring to concepts relating to other senses, speaking of a sweet, round smell, but I would be at a loss if I had to characterize what precisely the sweetness of the smell consists in. Training may of course improve my capacities to describe my experiences, and some people are better at it than others, but a certain ineffability seems to remain. Consider another example. Although we can see and touch various things, most of us will have great difficulties to say what precisely is characteristic of the

experience of sight, and how this differs from the experience of touch. Again, we cannot avoid a certain ineffability.

Even if we were able to produce a satisfying personal-level description of the phenomenal character of experience, our problems would not be over. True, such a description would be helpful in that it sharpens our view of the explanandum, and thus facilitates an explanation of it. But in itself such a description doesn't even rule out dualism, and we would still encounter the problem of relating this description to a description of phenomena at a subpersonal level. As long as this problem has not been solved, a naturalistic account of conscious experience has not been given. At best we would then have a narrative of our experiences, but as long as this narrative remains unconnected to non-experiential descriptions, a positive naturalistic understanding of experience is lacking.

If we wish to relate phenomenal experience to subpersonal descriptions of the natural world, perhaps the best way to proceed is to work on both sides of the equation (e.g. Humphrey 2000; Thompson 2007). We may then use our knowledge of subpersonal processes to carefully reconsider our perceptual phenomenology. We may thereby find a description of experience that can more readily be aligned with our ideas on how perceiving works. At the same time we may use our reflections on the phenomenal character of experience to inform our analysis of subpersonal processes. Careful consideration of perceptual experience may not only contribute by clarifying the explanandum; it may also provide a hint about the explanans.

Ideally this would yield a way to describe subpersonal processes, such that the description applies also at the personal level to the phenomenal character of experience. In other words, we may find characterizations that have what Humphrey has called *dual currency* (Humphrey 2000), in that they can be interpreted in terms of subpersonal processes, while they also apply to the level of the phenomenal character of the experience.[8] With such characterizations, we would be able to link descriptive features of experience to our understanding of the natural world. Dual currency

[8] Humphrey frames the challenge of finding 'dual currency concepts' more narrowly, as the challenge to understand how it can be, as he supposes we must all assume, that a certain sensory quality is the very same thing as a certain brain state. His particular proposal, which I shall not further discuss, attempts to understand sensations as 'a representation of what's happening to the organism', and perceptual experience as an inner response to sensory stimulation, in which the stimulation is 'interpreted as a sign of what is happening out there' (Humphrey 2000).

characterizations would be explanatory characterizations, in that they make intelligible how the characteristics of a high-level phenomenon are in agreement with lower-level characteristics of the world.[9]

Today, the dual currency ideal is widely accepted, although it may be expressed variously. A way to formulate the explanatory ideal is in terms of an 'isomorphism' between personal and subpersonal levels. Of course different theorists have different ideas about how to such an isomorpism can be found. For example Noë and Thompson, who criticize the idea that there must be a structural match or isomorphism between neural processes and experience, do accept a personal/subpersonal isomorphism as a constraint on the explanation of experience (Noë & Thompson 2004, p. 26). The challenge to relate experience to subpersonal processes need not be narrowly conceived as the challenge to relate it to neural processes.

A dual currency account would avoid the pitfall of providing 'nothing but' a list of the subpersonal processes that correlate with experience. Indeed, the very point of a dual currency explanation is to provide a coordinating account, in which one can see how the relevant subpersonal processes contribute to the high-level phenomenon of interest; this would answer the question why particular subpersonal phenomena should come with particular experiences and not with altogether different experiences. In such an account, it is in the organization of subpersonal phenomena that the personal level phenomena reveal themselves. Importantly, the personal level phenomenon is thereby captured without succumbing to the dualistic 'something else as well', like a ghostly entity posited to complement the basic biological processes involved in conscious experience.

5. Inner models and the explanatory gap

How can we characterize the relevant processes of perception, such that these processes explain the phenomenal character of experience? Evidently, much has happened since Ryle, and the most prevalent doctrine of today differs in important respects from the dualistic doctrine discussed

[9] Given that such an account would bridge the gap between personal and subpersonal levels of description, there is a sense in which the personal/subpersonal distinction breaks down for the particular case of phenomenal experience. At the same time, the distinction remains crucially important, as it allows one to make claims about one level without commitment to claims concerning the other level. Note that dual currency claims are open to critique based on personal-level considerations as well as from subpersonal considerations.

in Section 2 above. Presently the most popular framework for approaching perceptual experience hypothesizes the existence of *inner models*. The phenomenal character of experience is then viewed as a matter of the characteristics of the proposed models inside the brain. While Rylean considerations should make us conclude that having an inner model is not logically required for vision, it may be proposed that something like an inner model is in fact involved in perceptual experience. Although this view appears to have a strong intuitive appeal, I shall point out crucial difficulties of the inner model approach for reaching the dual currency ideal.

The natural starting-point for the inner models approach is a focus on the brain. Convinced that there is no mysterious mind-stuff, philosophers have concluded that conscious experience must be a brain process. As it is expressed in a major textbook of neuroscience:

> "Philosophically disposed against dualism, we are obliged to find a solution to the problem in terms of nerve cells and neural circuits." (Schwarz 2000, p. 1318.)

It must be noted that this conclusion does not logically follow from the rejection of dualism. As Susan Hurley has pointed out, there is no 'magical boundary' around the brain which ensures that the relevant subpersonal processes must be the processes inside the brain alone (Hurley 2010). Still, there is a widespread tendency to focus on internal processes. Importantly, these processes are then typically interpreted as models or representations of the world (e.g. Marr 1982; Metzinger 2003; Revonsuo 2006; Damasio 2010). Applied to visual perception, seeing an object then involves having something like a 'picture' in the brain. This 'picture' should of course not be imagined as an ordinary picture, but it is proposed that the brain somehow contains a model or map of the surrounding world, and that the person's experience is contained in this model.

A commitment to inner models need not be a commitment to a Cartesian theater, a place 'where consciousness happens' due to the intrinsic properties of this place. Indeed, inner model accounts may accept Dennett's view that there is no place in the brain where neural activity gets transformed into phenomenal experience. But while no such miraculous transformations occur, these accounts do assert that the inner states themselves embody the conscious experiences. It may then be based on their *aftermath* that we can single out neural processes as relevant to conscious experience, as Dennett stresses, but this does not imply that the conscious experiences and the inner processes cannot be the very same thing. While on Dennett's view consciousness is a matter of what the brain does, one might hypothesize that this involves processes embodying a model of the world.

A potential advantage of this position is that it provides a natural way to interpret the distinction between the person's conscious experience and the behavioral expression of the experience. By focusing on the proposed inner models, it may be thought, one can then focus on what the experience is like for the person, irrespective of the role the experience plays in the person's behavior. While Dennett's emphasis on the aftermath of neural activity is often taken to imply a bias towards the behavioral at the expense of a focus on phenomenal experience itself, the inner model theorist can avoid such a bias. The proposed inner model must play a role for the perceiver, but it need not always play a role in bringing about any behavior.

So suppose that phenomenal experience depends on the local properties of models in the brain (be it in virtue of non-intrinsic properties of the local activity). How then can we conceive of the subpersonal processes underlying our experiences? What is it about the relevant neural activity that makes that it comes with a particular phenomenal character? Can we relate the personal level of phenomenal experience to its underlying neural processes, and if so, how?

Perhaps we cannot make the link in practice. When we view inner processes in terms of the processing of information, we take a high-level perspective on the inner goings on, and it may not be practically possible to link this level to the lower-level properties of the alleged inner model. As David Marr has put it, it might be that the 'complexity barrier' is too great (Marr 1982, p. 349). The idea here is that an informational interpretation relates to the specifics of neural processes as a running computer program relates to its hardware implementation. Even in case of a simple computer program we could not readily understand the high-level computations in terms of its precise physical realization. So on this particular interpretation we should not expect that we could intuitively link the levels of description in the immensely complex case of neural information-processing.

In a more optimistic spirit, however, attempts are made to find a neural interpretation of the proposed inner models of experience. For this, inner model theorists have searched for correlations between neural processes and personal-level indicators of experience (e.g. Metzinger 2000). Some neural activity may correlate with the experience of seeing red, other activity with the experience of seeing green, of hearing a high tone, smelling a rose, or seeing a straight line. No doubt there will be differences in the neural activity underlying these experiences. Of course some features of the correlating activity may simply be irrelevant at the level of phenomenal experience. For example, the precise location in the head may not matter for the phenomenal character of an experience – the same activity could presumably have been taking place elsewhere, had the brain been wired up differently. The question then is which of the correlating neural properties

can help to provide a subpersonal interpretation of the character of experience. The challenge is to find neural correlates of consciousness and to find out how these would embody inner models.

There are two potential problems associated with this research program, which I shall address in turn. First, there may not be anything inside the brain that functions as an inner model, in which case one's search would be in vain. Second, even if there are neural processes that function as inner models, there are reasons why this research program may still fail to find dual currency explanations of the phenomenal character of experience.

5.1. Are there inner models?

The reliance on internal models or representations is certainly not to be taken for granted. A famous example that has led to some reservations about the appeal to inner models is provided by the study of Tetris (also discussed in Clark 1997). In this computer game configurations of blocks enter the screen from above, falling down and accumulating below. The goal is to rotate and move them in such a way that a gapless horizontal line below results (which will then disappear). As the game continues, the blocks fall faster and faster, so that the player has to be quicker and quicker in producing a good orientation and horizontal position of the blocks.

Traditionally, the cognitive challenge of a game of Tetris would have been generally analyzed as follows. First, the agent perceives the present state of affairs, presumably leading to the build-up of an inner model. Second, the agent considers the possibilities and computes the best possible action on the basis of this inner model. Third, the chosen action is executed. A perhaps surprising consequence of this traditional model is that investigation of the cognitive challenges of Tetris need not even involve actual Tetris playing. After all, what matters is what output a Tetris player generates. This output could be described for example as "shifting the block three to the right, turning it two quarters counter clockwise", and the actual performance of these actions would appear irrelevant for the problem-solving.

Kirsh and Maglio (1994) investigated how skilled Tetris players actually play. It turns out that players often start to move or rotate the blocks before they have even had time to think about the best orientation. Also blocks are sometimes moved all the way to the side, only to move them a few steps back again, so that the player can easier judge the precise horizontal position of the blocks. Such findings justify a distinction between pragmatic action and epistemic action (Kirsh & Maglio 1994). Pragmatic actions bring a goal physically closer, as when you move a configuration of blocks closer to the desired endpoint. Epistemic actions contribute to solving a cognitive

or perceptual problem, as when blocks are moved aside or rotated to judge the best place and orientation of their destination. Epistemic actions may help to judge the location of the blocks, as when blocks are moved all the way to the side, just to make it easier to judge how to get them at their desired end point. Kirsh and Maglio conclude that skilled Tetris players often physically rotate the blocks to reduce the internal computational requirements of the game.

The point is that the changing of the perceptual situation can contribute to the skillful performance of the game. On a more traditional view, what is in fact the computational upshot of a temporally extended process would be seen as the cognitive problem faced by the agent at a single moment (e.g. the problem to compute the best pragmatic move, before performing any action). In other words, it would seem that the agent has to do a lot of 'mental gymnastics', as Chemero (2009) calls it, to find out what he or she should do. While this traditional view has a significant role for inner models, the work of Kirsh and Maglio showed that this role may have been seriously overestimated.

In fact it has been argued that the best framework for explaining behavior may not at all be based on internal models of the world, but that it may better be cast in terms of dynamically unfolding patterns of organism-environment interactions (Keijzer 2001). Within such a dynamical view, no inner representations are required, leading to the idea that representational interpretations of subpersonal processes are superfluous, and can be discarded (Chemero 2009). More dramatically, an analysis of cognitive science by William Ramsey suggests that the very reliance on inner models in present-day cognitive science is limited, even where inner representations are not explicitly rejected (Ramsey 2007). The term 'representation' is often used all too casually, leading Ramsey to ask whether there is actually anything functioning as a representation or inner model in the accounts found in cognitive science. In other words, it can be questioned whether an informative notion of representation figures in scientific models, such that not just any process will count as a representation. To answer this, Ramsey has scrutinized theories, asking whether a notion of a representation or inner model does any explanatory work there. While the term 'representation' is ubiquitous in present-day cognitive science, his perhaps surprising conclusion is that cognitive science has actually shifted away from a representationalist framework. If this is right, and if we suppose that cognitive science is right to refrain from reliance on inner models, we may question whether we should rely on inner models for explaining the phenomenal character of experience.

5.2. Explanatory difficulties with inner models

As said, there is a further difficulty with the inner model-approach, even if there are neural processes that function as inner models. The trouble is that the search for the subpersonal basis of inner models may fail to get beyond mere *correlates*. Imagine for example that the visual experience of white is correlated with the activation of a neural group in a particular part of the brain. What is it about this particular neural group that makes it result in the experience of white rather than in the experience of black? And why does it correlate with visual experience rather than with experience of a very different kind? Let us assume that there are different types of neurons involved, or different spiking frequencies, or that there are different neurotransmitters at play. What then would it be about these neurons, spiking frequencies or neurotransmitters that would ensure that the experience of white results? Suppose that the experience correlates with spiking frequency of neurons, why then would it be, say, higher spiking frequency rather than lower spiking frequency that correlates with whiteness? It seems that whatever correlating neural activity is found, always the question can be asked: what is it about this neural process that makes it result in this experience rather than in a very different one? Seen from this perspective, the prospects for a neural explanation of the phenomenal character of experience are rather bleak (e.g. Taylor 1962; Levine 1983; Chalmers 1996; Hurley & Noë 2003; O'Regan 2011). As Hurley and Noë, referring to color experience, formulate it:

> "Neural properties are qualitatively inscrutable. If you were to land in the visual system as a microscopic alien, you couldn't tell, by looking around at the local fireworks, whether experience was happening, or whether, if it was, it was visual experience, or whether, given that it was visual, it was visual experience as of something red." (Hurley & Noë 2003, p. 132.)

In short, there appears to be an 'explanatory gap' (Levine 1983) between neural activity and conscious experience. More precisely, as Hurley and Noë (2003) and Chalmers (1996) have pointed out, we should distinguish the *comparative gaps* from the *absolute gap* (the terminology is Hurley and Noë's; Chalmers speaks of questions regarding *character* and *existence* respectively). The comparative gaps concern the phenomenal character of experience; the main challenge here is to explain the differences between experiences, e.g. to explain why certain subpersonal processes come with the experience of red rather than the experience of green, or why they come with a visual experience rather than an auditory experience. The absolute gap concerns the very existence of consciousness, i.e. the question

why a process comes with conscious experience at all; the challenge here concerns the difference between conscious and not-conscious. Setting aside the absolute gap for now, it seems that the trouble with neurophysiological correlates of experience is that we see no way to overcome the comparative gaps. No matter how much neuroscientific knowledge we acquire, it seems, it will remain obscure why the processes studied by neuroscience should be accompanied by a particular phenomenal 'feel'. Inner model-oriented accounts saddle us with the difficulty that neural properties – conceived as embodying inner models – seem too dissimilar from the properties of phenomenal experience. This, and the growing skepticism concerning the existence of inner models, motivates the turn to an alternative approach, in hope for better prospects.

6. Perceptual engagement: towards a skill-oriented approach

Given the difficulties associated with inner models, it seems that there is every reason to explore an alternative. Of course one might look for such an alternative within the brain, in terms of an account of neural processes that does no involve models. But the nature of the difficulties for the inner model approach gives us an indication of what a more promising alternative may look like.[10]

There is no doubt that neural processes are crucially important, but it looks as though the inner model characterizations of these processes fail to make contact with phenomena at the personal level, in particular with the phenomenal character of experience. It seems that local properties of inner processes are too far removed from the scale at which we live our lives as conscious persons. What we need are characterizations of the relevant subpersonal processes that remain closer to the way human beings as a whole experience the world. For this we may need a different way to bring subpersonal processes into the picture.

Traditional approaches have searched for the answers at the neural level itself, rather than at the level of the skills in which the neural processes participate. This way, the inner model oriented approaches could respect the intuitive distinction between phenomenal aspects and

[10] My main aim is to develop a positive account of the phenomenal character of experience and I do not think we should rule out the possibility that different approaches can yield good alternative ways of approaching the dual currency ideal. While I shall argue that we have reasons to reject inner models, I do not aim to argue against the possibility of internal, not model-based approaches.

behavioral aspects of experience. But the notion of skill need not be taken as a behavioral notion; perceptual skills need not reduce to behavioral skills. Alternatively, we may therefore focus on distinctively perceptual capacities, without introducing a behavior-oriented bias. Perhaps a skill-oriented view, without an additional commitment to inner models, can then help to understand the phenomenal character of experience after all.[11]

Such a skill-oriented view requires a major change of perspective on the subpersonal processes, compared to inner model oriented accounts. Susan Hurley has drawn the contrast in terms of the difference between 'horizontal modularity' and 'vertical modularity' (Hurley 1998; 2001). Traditional approaches analyze subpersonal processes as subsequent stages of input processing, leading to a central stage, which in turn results in neural precursors of motor output. On this view, a sequence of 'vertical modules' lies between input and output; inner models are a stage in this sequence. The alternative analysis considers whole dynamic loops of interaction involved in personal level capacities: 'vertical modules' involving input as well as output.

A prominent example of a vertical modular approach is provided by Rodney Brooks' robots, which are organized by adding different 'layers' of input-output routines (Brooks 1991). For example, one layer could be devoted to object-avoiding locomotion, while a second layer was devoted to the reaching of distant places; the second layer could then influence the robot's locomotion without itself being concerned with object-avoidance. Brooks argued that within such architecture, no internal representations had to be present to yield intelligent behavior. Rather than focusing on sequential 'stages' of processing, this invites a 'horizontal' analysis in terms of dynamic patterns of interaction between the agent and its environment. As Hurley and others have argued, we may get a better view of perceptual experience by adopting such a 'horizontal modularity' perspective.

At the personal level, the starting point for such an alternative approach can be provided by the notion of experience as a skillful mode of engagement with the environment. Human beings experience the world,

[11] Some theories are skill-oriented in a different sense, namely in that they connect perceptual content strongly to the perceiver's possibilities for action, proposing that the content of perceptual experience is partly constituted by the behavior allowed by a situation (e.g. Grush 1998; Ward, Roberts & Clark 2011). Such theories allow for a representational interpretation of experience, in terms of action-oriented representations (Ward, Roberts & Clark 2011). In this thesis I shall not discuss this particular proposal, although I shall elaborately discuss the relevance of action to perception.

and they do so from a certain bodily point of view, as they engage with their environment. While it may have seemed natural to think of visual experiences in terms of something like 'images', experience need not be conceived in terms of possessing inner models. What matters, on this view, is rather the way in which we are (or seem to be) engaged with the environment, and neural processes are only relevant in as far as they contribute to this (real or apparent) engagement. To address the phenomenal character of perceptual experience, we may then have to focus on the characteristic patterns of engagement – of *perceptual* engagement – rather than on the local neural properties.

6.1. Vision as dynamic engagement: the case of change blindness

This shift of perspective can be illustrated with a brief excursion into the richness of visual experience. Most of us will agree that our visual experience is very rich indeed. One only has to look up and appreciate the visual world in all its colorful detail to convince oneself of this. At the same time, there is a blind spot of considerable size quite central in the visual field of each eye and the resolution of our eyes is much worse in peripheral vision compared to focal vision. Also the frequent and typically unnoticed eye movements cause considerable retinal smearing, resulting in the temporal inability of the eye to effectively process information. This has led theorists to wonder how vision can seem so good despite the apparent flaws in the underlying mechanisms (e.g. O'Regan 1992).

Suppose you walk down the street as someone asks you directions. As you start explaining to the enquirer, you frequently look him or her straight in the eye. The broad daylight ensures that you should get a rich impression of the enquirer. Then two passers-by carrying a large object rudely interrupt your explanation by moving between you and the enquirer, temporally hiding the enquirer from your view. As they have passed, you continue giving directions as if nothing happened. If you were the subject of a surprising experiment by Simons and Levin (1998), what would in fact have happened is that during the short interruption, unbeknownst to you the enquirer has been swapped by someone else. As it turns out, many subjects fail to notice that the person standing in plain sight in front of them is replaced by someone else.

This phenomenon, in which subjects are remarkably 'blind' to visual differences or changes, is called *change blindness*. Very slow changes in a scene happening in plain sight may go unnoticed. Also when two pictures are viewed subsequently, under some conditions, e.g. when the changing of the pictures is accompanied by a short flash or when it coincides with a

blink, very drastic changes in the picture can go unnoticed by the viewer (e.g. Rensink, O'Regan & Clark 1997). Change blindness is generally considered as a surprising phenomenon, and people tend to overestimate their ability to detect changes (Levin *et al*. 2000). For how can we miss changes that occur in clear view? The interpretations that can be given of this phenomenon offer a good illustration of the new views on perceptual experience. For purpose of contrast, let us first consider a traditional interpretation of this phenomenon.

On a traditional interpretation, perceptual experience is primarily a matter of what happens inside the perceiver. Given the apparent richness of our visual world, the inner activity is typically conceived as building up an elaborate model or 'picture' of the world. One could then interpret change blindness as an indication that the subject has one rich inner model of the world replaced by another. From this perspective, change blindness is regarded as a symptom of the limited capacity of monitoring the differences between the elaborate models (e.g. Simons & Levin 1997). For example, one inner representation may get 'overwritten' by the other, or there just is no comparing mechanism for the different representations (for a brief overview of several accounts of change blindness, see Simons 2000).

Classical interpretations in terms of rich inner models cannot be logically ruled out by current change blindness evidence. But proponents of a more dynamical perspective of perception, which stresses the perceiver's interaction with the environment, offer an alternative interpretation which places less demand on perceptual systems. On this skillful engagement oriented interpretation of perception, perception is seen as first and foremost an active and exploratory encounter with the world. From this general perspective, two somewhat different interpretations can be given to change blindness. On one interpretation, our experience is less detailed then we think it is: the richness of experience is an illusion. On the other interpretation, experience is rich, even though no rich inner models are involved.

According to the first of these engagement-oriented interpretations, our visual world – as perceived at a certain moment in time – is not as rich as we normally take it to be: the richness of experience is *illusory* (Dennett 1991; O'Regan 1992; Rensink, O'Regan & Clark 1997). Rather than consisting in a rich inner representation that ensures the subjective presence of the visual world, seeing is "an active process of probing the external environment", and the richness and presence of the visual world "are actually an illusion, created by the fact that if we so much as faintly ask ourselves some question about the environment, an answer is immediately provided by the sensory information on the retina, possibly rendered available by an eye movement" (O'Regan 1992, p. 484). Perhaps there is no

rich inner model of the scene, although we may mistakenly take ourselves to possess such a model. The fact that people are surprised by change blindness could then be explained by the fact that our intuitions concerning the richness of experience are largely based on active encounters with the ever-present world. On a timescale of seconds, we may see quite a lot. But at one particular moment our perceptual experience may be very limited. We fail to notice how little we see.

Given the conception of perception as a skill, a lack of a detailed inner model does not imply a lack of detail on the level of perceptual experience. More recently, a somewhat different interpretation has therefore become available, which claims that "it is not necessary to represent all the detail internally to see the elaborate detail of the world" (Noë, Pessoa & Thompson 2000, p. 104). According to this interpretation, the richness of visual experience is no illusion; it is a genuine feature of the process of experiencing (Noë, Pessoa & Thompson 2000; Noë 2001). Our surprise upon finding out our susceptibility to change blindness may then have a different origin than an intuitive commitment to a traditional 'rich inner model'-account of the basis of perceptual experience. It may be simply based in the fact that we often do notice environmental changes. Our confidence in the ability to detect *changes* does not logically require a confidence in the possession of a rich inner model of the world at any particular point in time.

We should not presuppose that having a rich impression of the world requires the possession of a detailed inner model. What is present may just be the rich environment itself, together with a whole set of perceptual skills. The perceiver then knows how to access the available environmental detail: "although the brain may not construct a detailed model of the scene, the environment is detailed, and the mobile and exploring animal is able to discover that detail by active exploration" (Noë, Pessoa & Thompson 2000, p. 102). Just as the agent's ability of skillful Tetris playing requires less inner computations than one might expect, the perceiver's ability of skillful perceptual engagement with the environment may involve less elaborate inner models than one may have thought – and perhaps we may even discard inner models altogether.

6.2. Expanding our view: the sensorimotor approach

This skill-oriented perspective opens the possibility to appeal to a wider range of processes than the internal model focused approach concentrated on. Perhaps the exclusive focus on the local properties of neural processes has prematurely limited our view, by blocking different ways to bring the brain in the picture. We may then overcome the limitations of this focus by

characterizing experience in terms of the relevant patterns of perceiver-world interaction. Such an approach has recently been proposed by Kevin O'Regan, Alva Noë, Susan Hurley, Erik Myin and others (e.g. O'Regan & Noë 2001; Hurley & Noë 2003; O'Regan, Myin & Noë 2005). This approach is known as the *sensorimotor* approach, because it stresses the patterns of dependency of sensory stimulation on bodily action. From this perspective, visual experience has been explored in particular detail (O'Regan & Noë 2001), but the approach has also been applied to other experiences, such as olfactory experiences (Cooke & Myin 2011), as well as to more general aspects of perceptual experience (O'Regan, Myin & Noë 2005). A key idea behind this approach is, as Hurley and Noë have put it:

> "To find explanations of the qualitative character of experience, our gaze should be extended outward, to the dynamic relations between brain, body, and world." (Hurley & Noë 2003, p 132)

By referring to characteristics of this dynamic engagement, the sensorimotor approach aims to articulate what experiences are like, while it can be firmly committed to the challenge to relate the phenomenal character of experience to subpersonal processes. By focusing on the sensorimotor processes that characterize our engagement with the environment, this approach aims to overcome the difficulties of inner model oriented views. Because of its broadened view, this approach holds significant promise for keeping the description of the relevant subpersonal processes in touch with the phenomenal character of experience. In the next chapter I shall therefore turn to the sensorimotor approach in an attempt to acquire a better naturalistic understanding of the phenomenal character of experiences.

Chapter 2
The sensorimotor approach to phenomenal experience

How can we understand perceptual experience in relation to the natural world? The first chapter discussed the problem of the phenomenal character of experience, which presents a particularly intriguing challenge for a naturalistic understanding of mental phenomena. Here I explicate the approach I shall adopt in taking up this challenge, namely the sensorimotor approach.

1. Introduction

You open your eyes and behold the world. A split second before, light composed of a whole spectrum of different frequencies has fallen on objects, which reflected part of the light. When some of this light reached your eyes, it was refracted, and at the back of your eyes it hit photosensitive pigments, influencing the biochemical processes in the retina. Electrical signals were the result, modulating the activity of your brain, in some places more than in others. Perhaps this will express itself in a comment you make on what you see, or in some other behavioral response. But why, we may wonder, should all these processes come with this particular experience?[1] Although we believe that somehow our experience must be physically realized, the question is how this can be the case.

More particularly, as discussed in the previous chapter, the challenge is to relate the *phenomenal character* of perceptual experience – the specific way in which a person subjectively experiences his or her environment, considered independently of questions of veridicality – with the processes at the *subpersonal* level of description, such as the stimulation of the sense

[1] The issue is often framed in terms of 'brain states' and 'conscious states'. I shall speak more broadly of 'physical' processes rather than neural processes to avoid exclusive focus on the brain. To stress the dynamic nature of perception, I shall speak of 'processes' rather than 'states'. Imagine a frozen world, in which absolutely *nothing* moves – not even the molecules inside the heads of the frozen people. In this chilling scenario, would there be experience? Certainly not!

organs and the neural activity, which can be considered independently of their roles for the person as a whole. Ideally we would find a characterization of experience that fits both the description of what the experience is like for the person, as well as the description of the underlying processes that enable the experience. Such a characterization would have *dual currency*, as Nicholas Humphrey puts it, in that it can be interpreted in personal-level terms as well as in subpersonal terms (Humphrey 2000; see Chapter 1). Such a characterization would provide a naturalistic understanding of phenomenal experience by giving it its rightful place within our conception of the natural word.

On a traditional view, our perceptual experience is thought of in terms of the possession of an *inner model* of the world. But as pointed out in the previous chapter, dual currency characterizations relying on this 'inner model'-conception are notoriously difficult to find. Indeed, there appears to be an explanatory gap between phenomenal characteristics of experience and neural process. The difficulty is not just to find neural processes that correlate with experience. It is to make sense of these correlations. It is to explain why some neural activity should come with, say, the experience of something round rather than the experience of something angular, or visual experience rather than auditory experience. A difficulty for inner model approaches is that it appears impossible to read off the nature of an experience from neural properties.

The approach I shall discuss in the present chapter conceives of perceptual experience as *skillful perceptual engagement* with the environment, rather than as the possession of an inner model. By approaching perception as a skillful engagement with the environment, we may find a description of subpersonal processes that remains closer to the way human beings experience the world. The focus of my discussion will lie on the character of perceptual experience. But in as far as non-perceptual experiences such as mental imagery are derivative of perceptual experiences we may expect that consideration of perceptual cases could help to account for such non-perceptual experiences as well.

Below I shall explain how the skill-oriented perspective is worked out in the *sensorimotor account* of perceptual experience, and how this account approaches the dual currency ideal for the specific character of experiences. My main aim here is to explain the sensorimotor account, as proposed in a cluster of papers and books starting with a paper by Kevin O'Regan and Alva Noë (2001a). In explicating this account, my main focus shall be on visual experience. No doubt different experiences may vary in the extent to which they depend on sensory and motor factors respectively; by acknowledging the contribution of both factors in experience, the sensorimotor account develops a conceptual framework that can apply to

perceptual experience quite generally. In explicating the sensorimotor account, I shall particularly emphasize the way in which it relates personal and subpersonal descriptions of phenomenal experience.

This emphasis is not obvious. Indeed, within the sensorimotor literature, part of the reference to subpersonal processes has been negative: it is stressed that experience is a personal-level phenomenon and it is denied that there have to be subpersonal inner mechanisms for producing inner models of the environment (O'Regan & Noë 2001a). Still, as Evan Thompson has observed, it is precisely by linking the mental and the material through a dual currency characterization that the sensorimotor account may explain aspects of the phenomenal character of experience (Thompson 2007, p. 256). By discussing the sensorimotor account in light of the distinction between personal and subpersonal levels of description, I aim to explicate its distinctive explanatory contribution, highlighting crucial advantages of a skill-oriented approach to phenomenal experience.

The sensorimotor account of perceptual experience appeals to the action-dependence of perception. I shall therefore introduce the sensorimotor approach by explaining the relevance of action to perception, drawing on the work of Hermann von Helmholtz, Jacob von Uexküll, James Gibson, Susan Hurley, and others (Section 2). Next, I shall explain the core commitments of the sensorimotor account (Section 3). I conclude by highlighting potential explanatory advantages of the sensorimotor account and by discussing the way in which the account approaches the dual currency ideal (Section 4).

2. The action-dependence of perception

Perception, or the sensory experience of environment, may come so effortlessly that we tend to think of it as something simple and passive. Often there isn't much we have to do in order to perceive: having our eyes open suffices to see, and noises may be heard whether we are actively listening or not. But the ease with which we perceive defies the complexities of the mechanisms involved, and as soon as we do act this influences perception in various ways. In this section I shall discuss the ways in which action influences perception and I shall briefly attend to the inner mechanisms involved. In the next section we shall then see how the sensorimotor account draws on the role of action in perception, as well as on the relevant inner mechanisms, to explain the phenomenal character of perceptual experience. There it shall be suggested that also the more passive cases of perceptual experience, where the perceiver does not move

around, can be seen as special cases which can be better understood in light of the action-dependence of perception.

Before I turn to the ways in which action contributes to perception, let me set the stage by pointing out two well-known facts about the way in which perception is dependent on the stimulation of our sense organs.

First, experience is context-sensitive. For example, one spot of grey paint can be made appear light or dark, and even yellow or blue, depending on what you paint around it. What this shows is that color experience is not just a matter of the local sensory input coming from a certain spot. The explanation for the experience of the color at a spot must therefore appeal to a larger pattern of sensory input, which includes, as an additional factor, the sensory input of the surroundings of the spot.

A second example shows that explanation of sensory experiences may have to appeal to other sensory modalities as well. The well-known McGurk-effect in speech perception is a clear demonstration here: when you see a video in which a mouth is shaping one phoneme, while the sound of a different phoneme is displayed, the experience of the sound may be altered by the sight of the apparent speaker (McGurk & MacDonald 1976). Apparently, vision may affect auditory experience. The explanation of sensory experience cannot always be restricted to input from one sense modality only: additional factors may have to be taken into account.

Below I shall first discuss how action contributes to perception by changing the patterns of sensory stimulation (Section 2.1). Then I shall explain how action contributes to perception also by providing an additional factor besides sensory stimulation. Just as contextual factors from within the sensory domain contribute to perception, as in the first two examples above, factors related to action contribute to perception as well (Section 2.2). Subsequently I shall briefly attend to the inner mechanisms involved in the latter type of action-dependence of perception (Section 2.3).

2.1. Perception as active exploration

Suppose you are trying to find a flashlight in the dark. As you move your hands across the table where you expect to find it, you encounter various objects. Some of these you ignore, others you recognize only after a brief manual exploration. When you touch the flashlight, after a brief exploration you notice that you have found it. In this case, it is evident that the tactile exploration of objects makes their recognition easier than it would be at the basis of a one-shot sensory contact. By exploration we get a richer impression of the environment.

As several authors have pointed out, visual experience is not so different (e.g. MacKay 1967; O'Regan & Noë 2001a; O'Regan 2011). Suppose you

switch on the flashlight: it is then by shining around that you get a rich impression of the room. But even under bright daylight, your visual sampling of the scene greatly enhances your experience. One reason for this is that peripheral vision has a limited resolution, resulting in limitations that can be overcome by eye movements. In addition, already a subtle movement of the head can make it easy to see which parts of the stimulus belong to the same object, and how objects are spatially orientated with respect to each other. Given the fact that adjoining parts of the stimulus are likely to belong to the same object if they stay adjoined when the perceiver moves, movement disambiguates the stimulus. Just as in the case of feeling an object picked up in the dark, the sensory patterns reveal much of the nature of the object. Rather than contemplating the most probable layout of objects on the basis of a single 'snapshot' of the scene, we can move around to facilitate visual perception. Much of vision is a matter of active exploration (Findlay & Gilchrist 2003).

Due to active motion through the environment, we have an abundance of sensory information at our disposal, and it is well-known that action can thereby help to disambiguate stimuli (e.g. Gibson 1966). Moreover, also the change of perspective itself can provide the occasion for perception, e.g. when you experience your own movement. Consider for example the optical flow that results from driving towards the horizon, or landing with an airplane (Gibson 1950; 1979). When you look to the front, the optical flow originates in a point of expansion (see *Figure 1*). This is where you are going. If the point of expansion moves, this indicates that you are turning, and the velocity of the optical flow indicates your speed. Clearly, there is important information available in such dynamic patterns of optical change. Such temporally extended patterns may therefore be at the basis of perceptual experience.

Consider the patterns of optical change that result as you walk through a room. As you move around, parts of the room that were occluded by tables or chairs come into view, and the relative position of objects from your point of view is changing. The ways in which these changes take place are characteristic for the spatial relations between the objects. When you move, the visual system does not have to infer these spatial relations on the basis of independent impressions. Instead, James Gibson argues, these relations are themselves specified in the dynamic patterns of optical change (e.g. Gibson 1966; 1979). It may therefore be your sensitivity to these patterns, rather than your capacity to interpret short-lived sense impressions, that allow you to perceive the spatial layout of the world.

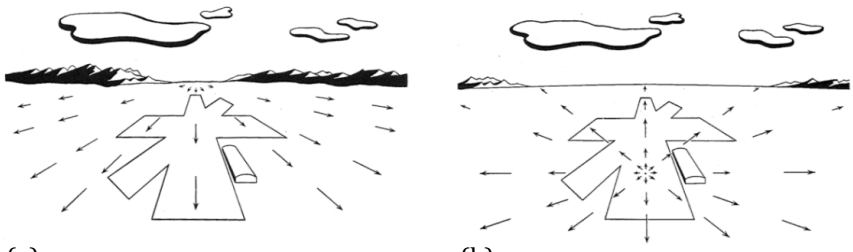

Figure 1. The points of optical expansion in horizontal movement (a) and movement towards the ground (b) show where the perceiver is heading. The velocity of the optical flow indicates the perceiver's speed. Drawings from Gibson (1950).

Psychologists in the wake of Gibson aim to identify the information that we are using. By focusing on the environmental preconditions of perception and on the abundance of available environmental information, they aim to avoid premature conclusions concerning the inner processes involved in perceptual experience. For example, contrary to a static 'snapshot' conception of vision, we may be perceptually sensitive to dynamic patterns of stimulation, such as the above-mentioned patterns of optical flow. But it remains to be seen how this sensitivity is realized.

Until now we considered how bodily movement contributes to perception by changing the sensory stimulation. Active engagement with the environment can make environmental detail available for scrutinizing, and action can introduce dynamical patterns of sensory stimulation that may themselves form the basis of perception. In such cases, action functions as a means for changing the sensory input; the contribution of action to perception is then said to be *instrumental* (Hurley 1998; 2001). For instrumental action-dependence, passive movement, such as illustrated in Figure 1 by the view from an airplane, serves as well as active bodily engagement (Taylor 1962, p. 319). But there is an additional way in which action may contribute to perception, which does not hinge on the changing of sensory stimulation. This is the topic of the next subsection.

2.2. The noninstrumental action-dependence of perception

Perceptual experience may differ as a result of motor activity, even if the sensory input – the affecting of the sense organs by stimuli from outside the perceiver – remains unaltered. If this is the case, there is a *noninstrumental*

contribution of action to perception (Hurley 1998; 2001).[2] In cases of noninstrumental action-dependence of perception, perceptual experience depends on action as a result of mechanisms other than the modulation of stimulation of the sense organs.

For example, suppose you are playing piano with your eyes closed. You could then perceive the spatial position of the keys you touch, were you to attend to this. Now suppose that the same pattern of stimulation on your fingertips would be produced on an inactive open hand. No doubt your tactile experience would be different: you would not have the experience of feeling a piano's keyboard oriented in space. The tactile experience of space depends not on the sensory input alone, but also on the sense of one's own movement or the position of the hand. The contribution of bodily activity to perception therefore cannot be reduced to its bringing about of patterns of stimulation of your sense organs. Also visual perception depends on action in a noninstrumental way, as the following cases show.

A classic case is provided by a study on the development of kittens (Held & Hein 1963). In this study, pairs of kittens were subjected to similar sensory stimulation for a few hours every day: in each pair, for one of the kittens the stimulation depended on the kitten's own movement, while the other received the stimulation irrespective of its movements. When the kittens had been subjected to this treatment for three or more days, their visually guided behavior was tested. While the kittens for which the sensory stimulation had been depending on their own movement responded normally, the other kittens responded strikingly different. For example, they were willing to jump off a 'visual cliff' where normal kittens would find another, presumably safer, route to descend. This change in visually guided behavior is remarkable, because the motor capacities of the kittens in the two groups seemed the same. This indicates that the kittens in the group in which the sensory stimulation had been independent of their own movement, had a specific impairment of vision rather than of basic motor capacities. It appears that development of a normal visual system depends on active engagement with the environment, rather than on the passive exposure to sensory stimulation alone.

Another, more familiar example of the noninstrumental action-dependence of perception concerns the activity of the eyes in normally developed human perceivers. It is clear that the presence or absence of eye

[2] Taylor expresses this by contrasting a 'primary' role of movement with a 'secondary' role: while the secondary role is instrumental, the primary role of movement is a noninstrumental role, requiring that the movement "is effected by the subject's own efforts" (Taylor 1962, p. 319).

movements matters for the visual experience that results from a given retinal stimulation. Suppose you move your eyes across a scene. The patterns of retinal stimulation are similar to the patterns that would result when all the objects in the visual field would move while you keep your eyes stationary. But in the former case, you experience a stable visual world, while in the latter case you would experience the visual movement of the world. The difference between these two cases is not based on differences in sensory input. If you make sure that there are no such differences, so that the patterns of retinal stimulation are the same, you will still be able to experience the difference between self-produced changes and environment-based changes. This indicates a noninstrumental role of action in perception, for it shows that experience varies with action in a way that does not depend on bringing about a difference in sensory input.

You can test this yourself when you close one eye and direct your view to a certain point. Now if you press gently at the lower eyelid of the open eye, you will observe that the visual world appears to move. While doing this, you will notice that you can actually keep your eye focused on the very same target in your visual field. (This is especially clear when you look at a piece of text while performing the experiment, for in this case you can verify that you are able to read the same limited range of words, and that your gaze therefore does not sweep across the text.) What happens, then, is that your visual world appears to move, despite the fact that your eye stays directed towards the same place and the retina is unmoved with respect to the stimulus (Bridgeman 2007). There is a change in visual experience that is not the result of a change in retinal stimulation. The explanation for this is that your eye muscles are actively compensating for the pressure to the side of your eye, and that there is something about this activity that has a noninstrumental impact on your experience.[3]

Further confirmation of this noninstrumental action-dependence of perception derives from experiments with inverting glasses, which, in the words of Taylor, "shattered once and for all the doctrine … that visual perception is a function of retinal stimulation alone" (Taylor 1962, p. 167).

[3] At least since Descartes discussed the example of experiences while pressing the eye, it has commonly but erroneously been presumed that the pressure to the eye always causes a passive movement of the eye (Bridgeman 2007). (This is erroneous, since movement is in fact often counteracted by your eye muscles.) But note that also on that false assumption, the same conclusion can be reached, as Helmholtz does when he attributes the lack of apparent motion in a normal eye movement to the 'the effort of will' (Helmholtz [1867]/1924, volume 3, edn. 1962, pp. 243 forward).

Inverting glasses alter the light entering the eyes, so that sensory stimulation is inverted in the left-right and/or the above-below direction. On first wearing such glasses, the world appears to move when one moves one's head, namely in the direction of the movement of the head, but at twice the speed. The apparent motion depends on head movement, for the retinal stimulation alone would not result in apparent motion of the same magnitude. Moreover, during prolonged wearing of inverting glasses, the apparent motion of the world disappears, while the consequences of movement for the stimulation of the eyes remain unaltered. As Taylor pointed out, both the original apparent movement and its disappearance clearly demonstrate the noninstrumental action-dependence of visual experience. The perceptual consequences of wearing inverting glasses are discussed in further detail in Chapter 6 below.

2.3. Inner mechanisms of noninstrumental action-dependence

The general idea that there is an active contribution of the perceiver to perceptual experience has a long history, dating back to the pre-Socratics (Grüsser 1995). In some early theories of vision, it was proposed that the eyes emitted 'fire' or 'visual spirits', an active contribution of the perceiver which was believed to integrate with the environmental light at the surface of objects. Also today, an active contribution of the perceiver is thought to integrate with the contribution from the environment, but now the integration of these two factors is thought to take place within the brain, where action-related signals and the signals based in the input from the environment interact (see Grüsser 1995 for a historical overview). Let us briefly consider these inner mechanisms, which provide possible subpersonal interpretations of the noninstrumental action-dependence of perceptual experience.

At a neurophysiological level, there are various ways in which action may influence perception through causal routes different from the bringing about of changes in sensory input. We can distinguish two familiar types of mechanisms that may be involved, the mechanisms proposed by the 'inflow' hypothesis, and those proposed by the 'outflow' hypothesis. The schemas depicted in Figure 2 illustrate these distinct neural mechanisms, as they appeared in von Uexküll (1920).

The 'inflow' hypothesis relies on signals deriving from outside the brain. These are signals originating in the activity of the muscles and signals related to the position of the joints, which we may call *proprioceptive feedback*. These signals are obviously relevant for proprioceptive experiences, but they may also alter our perception of the environment (e.g. James 1890; Sherrington 1918). The 'outflow' hypothesis holds that the noninstrumental action-dependence of perception is due to the neural

underpinnings of motor activity, rather than to the muscle activity or resulting body position itself. The relevant signals here are known as *corollary discharge* (Sperry 1950).[4] The idea has already been proposed in the early 17th century by Franciscus Aguilonius (Grüsser 1995), and it can be found for example in the work of Hermann von Helmholtz (1876/1924), who appealed to an 'effort of will' to explain the stability of vision despite eye movements, and in the work of Jacob von Uexküll (1920), who strongly emphasized the interdependence of perception and action. Today, it is believed that both proprioceptive feedback and corollary discharge play a role in perception.[5]

Note that in both models, the receptors (e.g. the photoreceptors in the eyes) modulate the activity in a part of the brain (e.g. in the visual areas), where the activity is also under influence of either proprioceptive feedback (*Figure 2a*) or corollary discharge (*Figure 2b*). Von Uexküll spoke of the underlying neural areas in terms of the 'central receptors', and 'central effectors' in the brain (von Uexküll 1934). Today we speak instead of 'sensory areas' and 'motor areas', but the general idea is the same.

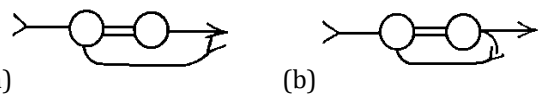

Figure 2. Schematic representations of two mechanisms of noninstrumental action-dependence of perception. (a) Signals originating in the muscles feed back to sensory areas in the brain. (b) Signals originating in a motor area in the brain feed back to sensory areas in the brain. Schemas from von Uexküll (1920/1928, p. 209).

Although the relevance of corollary discharge and proprioceptive feedback to perception is widely acknowledged, there is much that is unknown about the extent to which these processes contribute and the way in the work. Further interpretation of these processes are postponed to later sections, where I discuss the way in which these processes may be connected to phenomenal experience. For now, note that the schemas in Figure 2 capture routes of causal influence among different parts of the perceiver, and as

[4] Another term for (a type of) corollary discharge is *efference copy*, a hypothetical 'duplicate' of a signal initiating a movement which has been proposed to modulate sensory processing (e.g. von Holtz & Mittelstaedt 1950). I choose to use the term 'corollary discharge' because it is more neutral since it does not interpret the discharge as a 'copy'.

[5] Much work has focused on visual stability in relation to eye movements. See for example Bridgeman, van der Heijden and Velichkovsky (1994).

such should be interpreted strictly subpersonally. No conscious sensations of bodily movement or intentional action are implied; no commitment is made regarding the status of the proprioceptive feedback or corollary discharge at the level of conscious experience. For example, you may not be explicitly aware of the movement of your eyes in order for the signals relating to the movement to be effective for your visual experience. No conscious sensation of intentional action or 'effort of the will' need to accompany corollary discharge, and no conscious proprioceptive sensations need to accompany effective proprioceptive feedback.

To conclude, perception is intimately connected to active bodily movement. One way in which action contributes to perception is through the change of sensory stimulation. But we also saw that perceptual experience is not simply a function of sensory stimulation alone: action – understood as bodily movement – may influence perception independent of any change in sensory input. Mechanisms of proprioceptive feedback or corollary discharge can help to provide a subpersonal interpretation of this noninstrumental action-dependence of perception. In the next section we shall turn to the sensorimotor account, which draws on the action-dependence of perception and its underlying mechanisms to explain the phenomenal character of perceptual experience.

3. The sensorimotor account of phenomenal experience

How can we account for the specific *phenomenal character* of perceptual experiences? What, for example, gives visual experience its characteristic visual phenomenology and how can we explain its difference from tactile experience? How can we describe and explain the experience of the softness of a sponge or the straightness of a line? In this section we will turn to the *sensorimotor account* as first proposed by Kevin O'Regan and Alva Noë (2001a,b,c). On this account, our capacity to perceive is based, not on sensory stimulation alone, but on regularities in the way in which sensory stimulation depends on motor action. It is because of these regularities that exploratory capacities can develop (i.e. we may learn to act in a way that facilitates perception), and that perception becomes noninstrumentally action-dependent in a systematic way (e.g. eye movements without change of retinal stimulation will give the appearance of movement of the environment). I shall focus on the way in which this account explains the phenomenal character of experience by relating experience to subpersonal processes.

Sensorimotor approaches, broadly conceived, consider perceptual experience as a function of the relation between sensory stimulation and

motor action. Sensorimotor approaches go beyond the action-dependence of perception discussed in the section above, by claiming that *even if we do not move* at a certain moment, our experience is the implicit acknowledgement of the law connecting sensory stimulation and bodily movements. For example, in spatial vision we implicitly grasp the sensory consequences that would result from movements of our head. As Helmholtz puts it:

> "when we perceive before us the objects distributed in space, this perception is the acknowledgement of a lawlike connexion between our movements and the therewith occurring sensations." (Helmholtz 1878/1977, p. 138)

Sensorimotor approaches hold that perception cannot be understood by considering sensory stimulation alone. But accepting the relevance of sensorimotor patterns for perception does not in itself yield an account of the phenomenal character of experience. What is at stake for Helmholtz, for example, is our knowledge of the world rather than the quality of the experience. While he pointed out factors influencing our perceptual judgments, he did not thereby claim to address the phenomenal character of experience.[6]

In an influential paper, Kevin O'Regan and Alva Noë have proposed a sensorimotor account of perceptual experience that does address its phenomenal character (O'Regan & Noë 2001a). This account is closely related to the work of Susan Hurley (1998; 2001)[7], but O'Regan and Noë

---

[6] See Chapter 4 of this thesis for a discussion of the distinction between the epistemic conception of perception and the phenomenal conception. The limitation of Helmholtz' account is also apparent in his study of vision, where he addressed "our conceptions as to the existence, form and position of external objects", putting aside as much as possible questions of 'psychic energy' or 'the nature of the processes of the mind' (Helmholtz [1876]/1924/1962, pp. 1-2). Helmholtz' claim is limited in a second sense as well, for he applies sensorimotor ideas to spatial perception, but for example not to color experience. On his view, the *quality* of color experience is a function of sensory stimulation, without any motor component (Helmholtz 1878/1977, pp. 118-119). Thus crucial aspects of experience are unaddressed by his assertion of the relevance of sensorimotor patterns. (We shall see below that the present-day sensorimotor account addresses a much broader range of experiences.)

[7] Hurley proposed that perception and action are interdependent and that at the subpersonal level both rely on dynamic sensorimotor patterns. She discussed possible consequences of this view for our ideas about mental phenomena. More in particular she argued that consciousness, self and cognition are not inner processes

(2001a,b,c) more strongly emphasize the phenomenal character of experience. Their account has been further developed in a series of papers and books, including Hurley and Noë (2003), Noë (2004), O'Regan, Myin and Noë (2005), and O'Regan (2010; 2011). It claims that the phenomenal character of experience can be understood at the personal level in terms of patterns of skillful perceptual engagement with the environment. To develop a naturalistic account of these patterns of engagement, the approach appeals to the specific way in which, at the subpersonal level, sensory input depends on motor output, the so-called *sensorimotor contingencies* or *sensorimotor dependencies*.

Below I shall discuss the sensorimotor account of perceptual experience as proposed in O'Regan & Noë (2001a,b,c), focusing on key notions that remain influential in more recent explorations of the sensorimotor view. I shall first explain the key notion of *sensorimotor dependencies* as a way to characterize perceptual experiences at personal and subpersonal levels of description (Section 3.1). Next I shall explain the idea of experiencing as a skill, as expressed in the claim that experiencing consists in the *exercising of the implicit grasp* of sensorimotor dependencies (Section 3.2). Subsequently I shall discuss further constraints that must be satisfied for the presence of full-blown conscious experience (Section 3.3). In all subsections I shall draw connections between descriptions at the personal level and descriptions at subpersonal levels; in Section 3.2 and 3.3 I shall briefly touch on inner mechanisms in particular.

3.1. Grounding perceptual experience in sensorimotor dependencies

The sensorimotor account views perception as a skillful engagement with the environment (O'Regan & Noë 2001a,b,c; Myin & O'Regan 2002; Torrance 2002). The world is regarded as 'outside memory', available to the perceiver, who therefore need not possess a detailed inner model of the world (O'Regan 1992). O'Regan and Noë (2001a) lay great emphasis on their rejection of the view that experience consists in having an inner model or representation of the world, and they discuss a wide range of empirical findings to contrast their general framework with the 'inner model'-based framework. Here I shall focus on the way in which their skill-

sandwiched in between sensory input and motor output: "At the personal level, the self does not lurk hidden somewhere between perceptual input and behavioral output, but reappears out in the open, embodied and embedded in the world" (Hurley 1998, p. 3).

oriented perspective is fleshed out in the sensorimotor account of perceptual experience.

As said, the sensorimotor approach conceives of perceptual experience as a person's mode of engagement with the environment.[8] It is claimed that at the subpersonal level, the agent's perceptual engagement can be characterized in terms of relations of dependency of sensory input on motor action. *Sensory input* and *motor action* will be treated as subpersonal concepts, not to be conflated with the personal level categories of perception and action (Hurley 1998). Sensory input may or may not lead to a person's conscious experience, and to say that there is motor activity is not the same as to say that the person is engaged in an activity. Motor action contrasts with passive movement in that it must be produced by the system itself, but the concept does not imply consciousness or deliberate agency, since the activity may be unconscious and automatic. More generally, motor action differs from bodily movement, in that pushing against a concrete wall involves motor action but no bodily movement to speak of.

To get the flavor of the sensorimotor approach, consider the experience of driving a fast car, say, a Porsche (O'Regan & Noë 2001a,c). What is this experience like? Clearly, the experience does not consist in a simple bodily sensation. Instead, if one were to describe the way it feels to drive a Porsche, one would say things about the characteristic way in which the car responds to curve-taking, the way in which it accelerates when one pushes the gas pedal, the ease and impact of changing of the gears. Skillfully driving a car like this is a characteristic mode of engagement with the road, which clearly differs from the engagement with the road that results from driving a massive truck. Perhaps the best characterization of this mode of engagement is in terms of the way in which what you do affects your sensory situation.

In the case of the Porsche-driving example, however, one might still think of the patterns of change in the sensory situation in terms of a sequence of conscious sensations. On this interpretation, the example nicely exemplifies the idea that a Porsche-driving experience is a temporally extended mode of engagement, but until an explanation has

[8] O'Regan and Noë (2001a) also speak of vision as "a mode of exploration of the world". I use the term 'engagement' here to better capture their view that we need not act in order to perceive. (O'Regan and Noë's view of perception as "exploratory activity" is not meant to reduce perception to action, but to stress the action-dependence of perception and to contrast their view with the idea of perception as an inner state.)

been offered for the conscious sensations, there would be an important sense in which accounting for phenomenal experience has not even started yet. It would be as if one provides a personal-level narrative of experiences ('first I experienced *this*, then I put my foot down and I experienced *that*') rather than an explanatory account linking experience to subpersonal processes. The sensorimotor approach aims to do more than that. The idea is this: if there is such a thing as a Porsche-driving experience, then whether or not you are conscious of a variation of sensory experiences over time, there must be characteristic patterns of sensorimotor dependencies underlying the experience. The challenge for a sensorimotor account of phenomenal experience is to find such patterns.

Consider for example the experience of colored surfaces. You may never have noticed how, depending on the color of a surface, sensory input changes as you move your eyes. Still, there is evidence that action does play a role in color experience (e.g. Bompas & O'Regan 2006a,b). Also, you may never have noticed that your sensory stimulation changes as you turn a colored surface around, altering its orientation with respect to different light sources. Still, the sensorimotor account claims, such patterns of sensorimotor dependencies do determine the way the surface appears (O'Regan & Noë 2001a; Noë 2004; O'Regan 2010; 2011). These patterns form the subpersonal preconditions of normal color vision. An example that will be discussed in a later chapter concerns colors that seem to occupy a somewhat special place in color vision, namely 'pure' forms of red, green, yellow and blue (Chapter 5). As Philipona and O'Regan (2006) have found, at the subpersonal level the experience of these 'pure' colors has particular sensorimotor characteristics, which differentiate the experiences from other color experiences. If this is right, personal-level characteristics of color experience may be given a subpersonal interpretation in terms of differences in sensorimotor dependencies.

On this view, also different perceptual modalities, such as vision and audition, are characterized by different patterns of sensorimotor dependencies (O'Regan & Noë 2001a,b,c; Hurley & Noë 2003). For example, in vision, but not in audition, objects sharply occlude what is behind them. One may look behind an object by moving one's head sideways, thus making available new information that was previously concealed, in a way that depends on the distances between perceiver, object, and background. When approaching an object it expands in one's visual field, thus occluding a larger part of the background. Similarly, auditory stimuli appear louder on approach and they may even obscure other auditory stimuli. This however is no true case of occlusion, for the obscuring effect here is much less dependent on the precise spatial position of the obscured stimuli. Although we are not usually conscious of these dependencies – as said, one

may never have noticed the subtle interplay between actions and their sensory consequences – it is claimed that these dependencies are very much descriptive of visual phenomenology.

The sensorimotor account claims that perceptual phenomenology can be described by such patterns – to change the patterns is to change the experience. But while I expressed these patterns of sensorimotor engagement in terms of occlusion, optical expansion, and loudness, it should be clear that the descriptions above are not just personal-level descriptions of visual and auditory phenomenology. What makes the descriptions particularly revealing is that they can be given a clear subpersonal interpretation. When one moves to look behind an object there is a change in the *sensory stimulation* depending in a systematic way on *motor action*. At the subpersonal level of description, approaching a source of sound clearly brings about a change in sensory input. The more the sense organs are occupied by a loud stimulus, the less they will be mobilized by subtle distant sounds. And while an opaque object blocks light waves from behind the object, the acoustic patterns from behind an object may still reach the ears fairly undistorted.

The affinity between the personal-level description and the subpersonal description of experience can be appreciated even better in the example of experiencing the softness of a sponge. It will be evident that this experience cannot be understood in terms of a specific type of sensory input. Instead, as Myin (2003) points out, the feeling is described by the specifics of the sensorimotor exploration of the sponge, such as the specific way in which the sponge yields under pressure. A description of the experience of softness refers then to the physical activity of sponge-squeezing, not to a passive encounter with sensory stimulation. One may then characterize the sensorimotor exploration in personal-level terms by explicating what the experience of softness consists in, but clearly the exploration can also be given a subpersonal interpretation by describing the dependencies of sensory input and motor action. Thus, according to the sensorimotor account, the relevant sensorimotor dependencies describe the mode of engagement that constitutes the experience.

The 'corporality' and 'alerting capacity' of perceptual experience

The sensorimotor account describes specific perceptual experiences, such as the experience of visual occlusion, in terms of the properties of our engagement with the environment, and we have seen how these personal-level descriptions of perceptual phenomenology square with descriptions of sensorimotor interactions at the subpersonal level of description. To

further illustrate the sensorimotor account of perceptual phenomenology I shall now turn to more general aspects of perceptual experience.

O'Regan, Myin and Noë (2004; 2005) have argued that sensory experience differs systematically from the experience of thought or memories, in that the former but not the latter have a characteristic sensory 'feel'. In explaining their point, they provide an example of what it is like to vividly remember your grandmother. When remembering your grandmother, you may recall her voice and her way of speaking, the smell of her perfume, a characteristic facial expression. The memory may be somewhat similar to the experience you would have when you were actually facing her. Still, the memory somehow lacks perceptual reality. There is something typical for the experience of the actual sensory encounter that is usually not present in memory.

At least part of this phenomenal 'feel' of sensory experiences, O'Regan, Myin and Noë (2005) claim, corresponds to what they call the *corporality* and *alerting capacity* peculiar to sensory experience (O'Regan and Noë (2001b) speak of 'bodiliness' and 'grabbiness'). The corporality refers to the fact that in actual sensory experiences, movement of your body will result in immediate changes in sensory experience, as when you turn your eyes or blink. Thought or memory lacks such corporality. The alerting capacity refers to the fact that sudden environmental movements or unexpected sounds tend to grab your attention. A sudden movement in your visual field will often result in an immediate and automatic saccade towards the movement and in the orientation of your attention towards the place of movement. An unexpected pat on your shoulder has an alerting capacity that is not typically matched by memory, or by the thought of a pat on your shoulder.

At the personal level, corporality and alerting capacity can be considered descriptive of our perceptual engagement with the environment. At the subpersonal level, the alerting capacity of perception will be reflected in low-level features of perceptual systems, such as motion sensitivity of the visual system, which modulate our sensorimotor engagement with the environment. Corporality is straightforwardly implied by the explorative activity at the basis of typical perceptual experiences of the environment. The personal-level description of experience does not remain free-floating, but it can be connected to descriptions at the subpersonal level. Perceptual experiences come with a typical phenomenal 'feel' that may then be given a naturalistic interpretation in terms of the relevant patterns of bodily engagement with the world.

3.2. Experience as exercising perceptual skills

Although the sensorimotor account claims that perceptual experiences are characterized by sensorimotor dependencies, the account also recognizes that the mere presence of systematic sensorimotor dependencies does not guarantee perception. To have the capacity to perceive is to have the right perceptual competences; perceiving is the performance of these skills. As I shall explain below, the sensorimotor account of perceptual experience therefore holds that the perceiver must have *implicit grasp* of the sensorimotor dependencies (this is having a perceptual skill), and that this implicit grasp must be *exercised* in order to perceive (this is the performance of the skill) (e.g. O'Regan & Noë 2001a).[9] At the end of this section I relate these notions to the inner mechanisms involved in perceptual experience.

Some patterns of sensorimotor dependencies play no role in the perceiver's experience. In fact, there are many possibilities to define sensorimotor dependencies that do not lead to perception. For example, before a child has learned to see depth, the sensorimotor dependencies characteristic of depth are clearly in place: when the child moves, the changes in sensory input depend in a systematic way on the spatial characteristics of the environment. Until the child can see depth, however, the systematic patterns in the relation between head movement and changes in the sensory stimulation of the eyes have no perceptual significance. The child still lacks the appropriate neural circuitry for the sensitivity to these sensorimotor dependencies. The question of course is what this 'appropriate neural circuitry' amounts to. I will come back to this below, but for now the point is that the mere presence of sensorimotor dependencies does not suffice for perceptual experience.

The sensorimotor account of O'Regan & Noë (2001a) acknowledges this by claiming that perceiving is a skillful activity of the perceiver, it is something we do. For this, the perceiver must have *implicit grasp*, or

[9] To gain a better understanding of these remarks, consider an analogous case of a behavioral skill. For example, the capacity to cycle implies that the person implicitly grasps the sensorimotor dependencies pertaining to cycling (dependencies related to keeping one's balance, steering, etc.). On the assumption that the relevant skills consist in a characteristic mode of sensorimotor engagement, the assertion that, in order to perceive (or to cycle) one must *exercise* one's *implicit grasp* of the relevant sensorimotor dependencies is no empirical speculation, but a properly *descriptive* assertion (by contrast, the precise role of the brain in the relevant skills is an empirical matter).

'implicit knowledge', of the laws of sensorimotor dependencies. Or as it is sometimes said, the perceiver must be *tuned to* or have 'mastery over' the laws of sensorimotor dependencies.

These metaphors should be treated with caution. The implicit grasp of sensorimotor dependencies should certainly not be understood as the possession of knowledge-that, but rather as the possession of skills, or know-how (e.g. O'Regan & Noë 2001a,c; Myin & O'Regan 2002).[10] As Myin and O'Regan emphasize, "The knowledge of the sensorimotor contingencies is not an independent or separately stored item which is available to the perceiver, but it is *implicit*, present only in the particular ways the ongoing exploration unfolds" (Myin & O'Regan 2002, p. 34). The introduction of implicit grasp should not be considered as the introduction of an extra layer in the sensorimotor account of experience: 'grasping' sensorimotor patterns does not consist in something over and above an embodied sensorimotor engagement with the environment. Instead, the introduction of the concept of 'implicit grasp' is meant to distinguish genuine perceptual engagement from the mere presence of sensorimotor dependencies. To put it as plastically as possible, systematic sensorimotor dependencies could still obtain when you remove the brain – spasms would still have systematic sensory consequences depending on the environment – but without appropriate inner structures, there is no perceptual engagement, no perceiver that is tuned to the sensorimotor regularities.

There clearly are many sensorimotor dependencies which we fail to grasp. By saying that implicit grasp is implied in perceptual experience, as

[10] The distinction between knowing-how and knowing-that is taken from Ryle (1949) (see also Chapter 1 above). While O'Regan and Noë emphasize that "visual experience rests on know-how, the possession of skills" (O'Regan and Noë 2001a, p. 946), they have also asserted that experience "consists in the knowledge that" certain laws obtain (O'Regan and Noë 2001a, p. 949). In particular such a 'knowledge that' formulation has been criticized by Hutto (2005), but his criticism applies to implicit know-how as well. Hutto writes: "While I endorse the general spirit of and the core message of the [sensorimotor contingency] approach – i.e., that the character of experiences is determined by sensorimotor contingencies specific to the various sense modalities – I find the invocation of knowledge … to be treacherous (and ultimately unnecessary)" (Hutto 2005, p. 391). If Hutto is right, unnecessary talk of 'knowledge' can be avoided: talk such as that persons 'have implicit knowledge of sensorimotor contingencies' can be replaced by saying that they are 'adapted to the occurrence of sensorimotor contingencies'. I will not engage in such exercises of translation. Instead I shall explain the sensorimotor account in the terms introduced by O'Regan and Noë (2001a), aiming to explicate how this account may approach the dual currency ideal.

said, the sensorimotor account captures the fact that one may have to learn to perceive. To further illustrate this with an example we already saw in Section 2.2, when you press with your finger against the side of your eye, the world appears to move, despite the fact that you press on your eye yourself. Clearly, then, you lack the implicit grasp of the consequences of pressing with your finger against your eye required for visual stability. While you implicitly grasp the sensorimotor regularities associated with normal eye-movements, you are not tuned to the sensorimotor laws associated with pressing against the side of your eye. In a similar way, infants still lack the implicit grasp of many of the sensorimotor dependencies that adults have learned to rely on.

Having perceptual skills may in part be a matter of mastering the instrumental action-dependence of perception. Consider for example the know-how involved in having a detailed view of a moving object or a person walking by, given the limited sensitivity of our eyes in peripheral vision. For perceivers with eyes like ours, the relevant skills depend on the ability to track objects with our eyes, or to keep the object in central vision by moving our heads. Someone who lacks these capacities may of course accidentally move his or her eyes in a way that keeps something in central vision. But lacking the object-tracking sensorimotor know-how, he or she will lack a common perceptual capacity, namely the capacity to visually explore moving things of interest and to get a good view of them. Visual skills may thus rely on behavioral capacities.

Sensorimotor theorists sometimes speak of 'acting out our experience' (e.g. O'Regan & Noë 2001a), and it is proposed that perception is a kind of 'skillful bodily activity' (e.g. Noë 2004, p. 2). This is a way to say that experiencing involves the exercise of one's grasp of sensorimotor regularities, and to stress that experiences are often depending on temporally extended patterns of active exploration. To get a rich visual impression of the world one has to look around, and in that sense experiencing can be seen as an exploratory activity.

Still, our implicit knowledge of sensorimotor dependencies goes beyond such an instrumental contribution of action to perception. Also the noninstrumental action-dependence of perception is a matter of being attuned to the laws of sensorimotor dependencies. For example, when an object visually appears stationary during a normal eye movement, this reflects your grasp of the sensory consequences of your action. While the eye movement does not contribute by providing a better view of the object, your experience still depends on your eye movement, namely in a noninstrumental way.

The implicit know-how of the sensorimotor dependencies is sometimes referred to in terms of expectancies. The reason is that, when you grasp the

relevant sensorimotor dependencies, you thereby have the implicit knowledge of the sensory consequences that can be expected to result from bodily movement. For example, to see that something is moving is to implicitly know how it may be followed with your gaze, and it is to grasp the sensory consequences that are to be expected from tracking the object with your eyes. And to see that an object is occluding something in the distant background is to have certain expectancies regarding the way in which the occluded parts of the world may be brought into view. Again, the expectancies need not be explicit. For example, consider what it implies when someone sees a thimble – as Ryle writes (also cited in O'Regan & Noë 2001a, p. 945):

> "Knowing how thimbles look, he is ready to anticipate, though he need not actually anticipate, how it will look, if he approaches it, or moves away from it; and when, without having executed any such anticipations, he does approach it, or move away from it, it looks as he was prepared for it to look. When the actual glimpses of it that he gets are got according to the thimble recipe, they satisfy his acquired expectation-propensities; and this is his espying the thimble." (Ryle 1949, p. 218)

Having a perceptual ability implies having an implicit grasp of the relevant sensorimotor dependencies. In order to actually perceive, this ability should be *exercised*. Just as having the capacity to cycle is not yet to cycle, having the capacity to perceive certain features is not yet to perceive these features. Suppose that the ability to perceive the straightness of a line depends on the implicit grasp of characteristic regularities, or sensorimotor laws, such as the fact that the sensory stimulation does not change if one moves one's eye across the line (O'Regan & Noë 2001a). To have the perceptual experience of a straight line this skill then has to be exercised in the encounter with a straight line. Perceptual engagement with the environment consists in the exercising of one's implicit grasp of sensorimotor dependencies.

Note that in this approach, perceptual engagement may but need not involve action. As Helmholtz already appreciated, our spatial perception consists in the implicit acknowledgement of the laws of sensorimotor dependencies. To see that an object stands in front of another object, on this approach, consists in the implicit grasp of the way in which sensory stimulation would change if one were to move. Since we have developed this perceptual capacity, we can perceive the spatial layout of the world: even without overt bodily movement, we are ready to exercise our implicit grasp of the relevant sensorimotor dependencies. In this sense, experiencing is a skillful mode of engagement with the environment.

Subpersonal mechanisms supporting perceptual engagement

Perceptual experience requires an implicit grasp of the patterns of sensorimotor dependencies. Further requirements for full-blown conscious experience shall be discussed in the next subsection. But first I shall relate the preconditions discussed so far to the corresponding subpersonal processes. We already saw how sensorimotor patterns can be given a subpersonal interpretation. But how does the sensorimotor account conceive of the subpersonal processes that enable our implicit grasp of sensorimotor dependencies?

At the personal level, sensorimotor dependencies can only matter for someone with the right perceptual skills: you can only see red objects if you have the implicit grasp of the sensorimotor dependencies pertaining to red objects. At the subpersonal level, a perceiver's possession of perceptual skills implies that the body or the brain is tuned to the laws of sensorimotor dependencies (O'Regan & Noë 2001a,b,c). If the brain is 'tuned to' obtaining sensorimotor regularities, this simply means that there is neural activity that would not be there in case of occurrence of the same sensory stimulation and motor action in a perceiver lacking the sensitivity to the sensorimotor regularities.

To flesh this out, let us first consider cases in which the perceiver is actively exploring the environment. An important class of inner mechanisms that are then at play has already been mentioned in Section 2.3. These are the mechanisms of noninstrumental action-dependence of perception, in particular the mechanisms involving *proprioceptive feedback* and *corollary discharge*. It is clear that, at a neurophysiological level, the development of implicit grasp of sensorimotor dependencies involves the development of corollary discharge circuits and proprioceptive feedback mechanisms.

The way in which the sensorimotor account conceives of these processes can best be understood by drawing a contrast with accounts that are oriented towards inner models. A clear statement of an 'inner model'-oriented interpretation of neural mechanisms is provided by Crapse and Sommer, who characterize the contribution of a type of corollary discharge to perception as follows: "it facilitates the contextual interpretation of sensory information (…) and the construction and maintenance of an internal representation of this information" (Crapse & Sommer 2008, p. 589). Perception, on this view, relies on an inner model or representation of what is perceived, and corollary discharge contributes to perception by facilitating the construction of this inner model; visual stability is analyzed as the stability of the inner model.

The sensorimotor account, as proposed by O'Regan and Noë (2001a,b,c), rejects this view. As said, the account conceives of perception as a way of engaging with the world, rather than as the possession of an inner model. Both in the sensorimotor account, as well as in inner model accounts, neural processes must be *tuned to* the sensorimotor dependencies in order for corollary discharge or proprioceptive feedback to facilitate perceptual experience: to say that the brain is tuned to sensorimotor regularities is just to say that the brain can be engaged with these regularities. The difference is that according to inner model accounts, the tuning to these dependencies serves the construction of an internal representation of aspects of the world, while this is not implied by the sensorimotor account. According to the sensorimotor account, the role of the brain is exactly to facilitate our sensitivity to sensorimotor dependencies. This sensitivity is facilitated by neural activity originating in sensory stimulation in combination with the corollary discharges and proprioceptive feedback.

Of course a person does not have to move to perceive. As a result, the extent to which motor-related signals contribute to experience may vary. When someone does not move, there may be proprioceptive signals in as far as these are not specifically action-related, but there is no corollary discharge. Consider for example the visual experience of a straight line in absence of eye movements. Based on sensory input and perhaps proprioceptive signals related to the position of your eyes, you implicitly grasp the sensorimotor dependencies pertaining to the straight line. That is, you implicitly grasp the sensory consequences that are to be expected if you would move your eyes, e.g. if you were to follow the line with your eyes, the sensory stimulation would not change. That your brain is tuned to the relevant laws of sensorimotor dependencies no doubt has its origin in your history of engagement involving the motor-related signals, but such signals are not presently at work.

There is much action-related activity in the brain and it is clear that this plays a major role in perception (Crapse & Sommer 2008; Guillery & Sherman 2011). But the details of the relevant inner processes remain to be filled in. We know a lot about correlations between neural activity and experience, but much is presently unknown about the precise way in which the implicit grasp of sensorimotor dependencies may be realized (in this sense, the sensorimotor account is in the same position as 'inner model'-based accounts).

Note that on the sensorimotor account, the inner processes do not themselves provide the 'feel' of experience. These processes participate in a characteristic mode of interaction, but for example the corporality and the alerting capacity of perceptual experience are not to be found in the local properties of the brain. Rather, they are properties of our sensorimotor

engagement with the environment. The description of experience is to be found at the level of the person's engagement, not at the level of participating neural processes. When we reflect on the character of visual or tactile experience, we reflect on the patterns of visual of tactile exploration of the environment (O'Regan & Noë 2001a).

In some cases our skillful perceptual engagement may, at the subpersonal level, involve a lack of sensitivity to the sensory changes that result from motor action. For example, due to saccadic movements of the eyes, there are shifts in the patterns of retinal stimulation. The fact that the world does not appear to move during saccadic eye movements may in part be due to our insensitivity to the changes in retinal stimulation (MacKay 1962b; O'Regan 1992; O'Regan & Noë 2001a). This insensitivity should not be viewed as the result of a stabilized inner model, but rather as the insensitivity to potentially attention-grabbing changes in sensory stimulation. Surprising events capture our attention; the unsurprising self-produced changes in sensory stimulation do not. (Note that a role for an insensitivity to change does not imply that a lack of change would serve just as well. After all, if there would be no change of sensory stimulation when you move your eyes across a scene, this would surely be a surprising event, which would come with an experience as if the world moved.)

Much work remains to be done to spell out in detail how the brain facilitates our implicit grasp of sensorimotor dependencies. What the sensorimotor account offers is a framework for bringing neural processes in view within an account of phenomenal experience. It proposes, without relying on the 'inner model'-conception, that the brain facilitates the implicit grasp of sensorimotor dependencies, and it proposes that the sensorimotor dependencies are characteristic of perceptual experiences. Section 4 draws possible consequences of this view for the dual currency ideal. But first we must consider additional requirements that have been associated with perceptual consciousness.

3.3. Further requirements for full-blown conscious experience

Above we saw how conscious perceptual experiences may be characterized in terms of the patterns of sensorimotor engagement with the environment. But this is not to say that all patterns of sensorimotor engagement with the environment are relevant to conscious perceptual experience. What then determines whether conscious experience obtains at all?

When driving a car while completely absorbed in conversation, one may adequately respond to the visible situation, without being fully conscious of the visible environment. Even though one is exercising one's implicit grasp of the visual sensorimotor dependencies, and in that sense may be

considered to be 'aware' of the environment to a certain extent, a full-blown conscious experience of the environment is lacking. In order to do justice to such cases, the sensorimotor approach can formulate additional constraints that must be fulfilled before we may speak of full-blown conscious experience. In the sensorimotor account of O'Regan and Noë (2001a,c), such more specific requirements are proposed. By incorporating further constraints on the sensorimotor account, the characterization of conscious experience should become more complete.

Note that by introducing these further constraints we shall remain within a skill-oriented perspective on perception: we shall not introduce an 'additional layer' consisting of an inner model in the account. A better way to view the matter is this. By introducing additional constraints, we further eliminate from our account the patterns of sensorimotor contingencies which are not pertaining to conscious experience. For example some skillful visually guided actions may draw on sensorimotor contingencies that remain outside the scope of conscious experience. A positive account of conscious experience may ignore such patterns, just as we have ignored the sensorimotor dependencies which a perceiver does not implicitly grasp.

The idea is that conscious experience – at least in the clearest cases – implies that the perceiver is able to exercise one's implicit grasp of sensorimotor dependencies for the *planning or guidance of action*, for *thought*, and in humans for *speech* (O'Regan & Noë 2001a,c). It does not seem right to say that a person is conscious of something, unless the person is able to spend further thought on it, to say something about it, or to otherwise use it to modify your behavior. Conscious experience is conceptually tied not only to phenomenal 'feel', it is also conceptually connected with the fulfilling of a role. When we can deliberately take things into account in our behavior, this means that we are conscious of them, and this typically implies that we can verbally report this.[11]

[11] It is widely accepted, though not uncontroversial, that the difference between conscious experience and its absence must lie in a functional difference, a difference in the fulfilling of a role (e.g. Cohen & Dennett 2011). The wide acceptance is evidenced in the fact that the scientific study of conscious experience uses behavioral measures of consciousness, including subjective report, a practice strongly defended for example by Dennett (1991). The existence of a necessary connection between experience and its possible role is not uncontroversial, for according to some phenomenal experience might fail to play a role, and even the *whole* set of functions usually associated with conscious experience might as well be fulfilled in absence of experience (e.g. Chalmers 1996). The sensorimotor approach follows Dennett in rejecting the latter conception of phenomenal experience,

Not all cases of conscious experience are as clear as that. We may therefore consider the presence of conscious experience as a matter of degree, in the following two ways (cf. O'Regan & Noë 2001a, p. 944). First, you may be conscious of a more or less limited range of aspects of the environment. There are sounds, smells, objects, shadows etc. in abundance, and the more one is engaged with these various aspects of the environment, the more one can be said to be aware of the environment. A second sense in which consciousness is a matter of degree is in that the environmental aspects with which one is engaged may have a larger or smaller impact. If you do make use of perception for deliberate action, while the experience does not engage any subsequent thought, your conscious experience is less extensive than when you also integrate what you perceive in your thinking.

The main focus of the sensorimotor account of phenomenal experience has been on the *phenomenal character* of conscious perceptual experiences. The account tends to start from particular conscious experiences, considering what explains their specific character. Unsurprisingly, then, development of precise criteria for consciousness (or criteria for a lack of consciousness) has not been a priority. But by allowing for specific requirements that must be fulfilled before the mode of sensorimotor engagement pertains to someone's conscious experience, the sensorimotor account can incorporate criteria for conscious experience. Let us briefly consider the way in which such requirements – requirements concerning the role played by one's implicit grasp of sensorimotor dependencies – may be reflected at the subpersonal level, focusing at inner mechanisms in particular.

Subpersonal mechanisms specific for conscious perceptual engagement

Conscious experience is thought to be connected with the selection of action, with thought, and in humans with the potential for verbal report. Thus it is claimed that our grasp of sensorimotor dependencies must play a certain role in our lives to pertain to consciousness. Let me briefly relate this role to its subpersonal processes to emphasize that no subpersonal consciousness is proposed in the sensorimotor account.

In general, the relevant inner mechanisms can be thought of in terms of Dennett's 'multiple drafts' model (Dennett 1991) or his 'fame in the brain' model (Dennett 2001). On Dennett's view, as discussed in the previous

according to which phenomenal experience becomes divorced from everything we say about it (see also Chapter 4).

chapter, at the subpersonal level some processes will have a larger causal influence than other processes, and this difference in *aftermath* makes for the relevant differences at the personal level, such as the person's capacity to say what he or she perceives or to otherwise respond to stimuli (Dennett 1991; 2001; see also Chapter 1 above). Like Dennett, O'Regan and Noë (2001a) reject the idea that consciousness is a matter of the intrinsic properties of subpersonal processes inside the head. (Where the account of O'Regan and Noë differs from Dennett's is in explicitly addressing the phenomenal character of experience, as we saw in Section 3.1 above; see also Chapter 4.) It is proposed that there need not be anything special about the processes underlying conscious experience, except that they happen to have a larger impact on subsequent processes, such as the processes underlying verbal report, thought, or action.

This leaves open the question how the relevant differences in causal impact come about. A more specific view on the relevant neural processes may be provided by 'neural workspace' models, to which we will turn in the next chapter. For example, it has been suggested that the processes underlying consciousness are characterized by "a distributed neural system or 'workspace' with long-distance connectivity that can potentially interconnect multiple specialized brain areas in a coordinated, though variable manner" (Dehaene & Naccache 2001, p. 13). In the next chapter the option of combining such models with a sensorimotor approach will be discussed, as well as the explanatory advantages that such a combination yields (Chapter 3).

In short, perceptual experience is thought to be grounded in patterns of sensorimotor engagement, of which we have now seen some examples. According to the sensorimotor account, experiencing consists in the *exercising* of the *implicit grasp* of *sensorimotor dependencies*, and the extent to which this skillful engagement plays certain roles in our lives determines the extent to which we speak of conscious experience. This is the core of the sensorimotor interpretation of perceptual experience.

4. Approaching the dual currency ideal

We saw how the sensorimotor account draws on the action-dependence of perception to account for the phenomenal character of experience, and how it characterizes perceptual experience as a skillful mode of engagement with the environment. In this concluding section I shall first highlight explanatory advantages of the sensorimotor account by contrasting the account with 'inner model'-accounts (Section 4.1). Next, I shall address the specific advantage of the sensorimotor account regarding the dual currency ideal. I shall explicate the way in which the account approaches this ideal

for the phenomenal character of perceptual experience (Section 4.2), and draw a conclusion (Section 4.3).

4.1. Perception as skillful engagement: descriptive and explanatory advantages

To explain the sensorimotor account I have distinguished personal and subpersonal levels of description. This distinction is important because it allows us to distinguished different claims made by the sensorimotor account. At the personal level, the sensorimotor account offers a description of the phenomenal character of experience. Its claims here concern perceptual phenomenology. At the subpersonal level of description, the sensorimotor approach offers an account of the processes underlying perceptual experience. In the next subsection we shall consider in more detail how these levels of description relate, by focusing on the way in which the sensorimotor account may bring us closer to the dual currency ideal. But let us first keep the levels of description apart in order to consider other advantages of the 'skillful engagement'-oriented sensorimotor approach compared to 'inner model'-accounts.

At the personal level, an advantage of the sensorimotor account is that it helps us to articulate what our experiences are like. For example, among the things that are characteristic for visual experience are the patterns of occlusion and optical expansion that can be found when we move with respect to objects. By focusing on the sensorimotor characteristics of occlusion we can give a more fine-grained description of the experience than when we would just ascertain that one object occludes another.[12] Section 3 above further illustrates how a sensorimotor account helps to describe the phenomenal character of experience (see also Chapter 6 for a more detailed discussion of spatial aspects of vision from this perspective). But even if the sensorimotor account is right in its description of the personal-level explanandum, this does not imply that its description of the subpersonal explanans is correct.

At the subpersonal level, the skill-oriented sensorimotor account has a potential advantage over 'inner model'-accounts. Given that normal perceptual experiences have sensorimotor preconditions – experiences depend not on sensory input only, but on motor action as well – the

[12] The value of a good description should not be underestimated. If we just know that something 'looks red' we have less of a lead for explaining the experience than when we have a description in which the red opject is particularly attention-grabbing as it stands out against a background.

sensorimotor account offers a simpler account of the subpersonal processes than accounts based on inner models. The reason is this. Both the sensorimotor account and 'inner model'-accounts must explain the *implicit grasp* of sensorimotor dependencies involved in perception. For this *both* approaches may appeal to corollary discharges and proprioceptive feedback signals, as pointed out in Section 3.2. However, from the perspective of an 'inner model'-account these processes would be involved in the construction of an inner model, and the challenge would then remain to *also* find the neural basis of these proposed inner models. This difficulty evaporates in the sensorimotor account.

This potential advantage of the sensorimotor account becomes even more salient when we consider its personal-level characterization of experience to be accurate. While the sensorimotor account proposes to explain the phenomenal character of experience by appealing to patterns of perceiver-environment interaction themselves, 'inner model'-accounts propose that in addition to these patterns, inner models are constructed in the brain. 'Inner model'-accounts may then have to explain how the characteristics of sensorimotor engagement apply to the proposed inner models themselves. In other words, the difficulty is then to find a characterization of the proposed inner models that maintains the characteristics of experience: characteristics of the sensorimotor engagement are then *duplicated* inside the brain. For this an 'inner model'-account may postulate not only an inner model of the environment, but an inner model of the perceiver as well (e.g. Metzinger 2003). In such an account, the interaction between perceiver and environment should be captured within the models inside the brain. The challenge then remains to flesh out this inner model in subpersonal terms. Not only patterns of skillful engagement must then be accounted for, but also internal shadow-patterns reflecting this engagement. The sensorimotor approach avoids this second problem by claiming that phenomenal experience consists in the patterns of sensorimotor engagement themselves.

It thus seems that a skill-oriented perspective is more parsimonious than its model-based rival, since it doesn't duplicate the problems that the cognitive neurosciences have to solve. Now of course the choice between inner models and a skill-oriented perspective may not be such a simple matter. At this point it may even be suggested that we are must accept an 'inner model'-account. The reason is that it may seem that an 'inner model'-account is needed to explain such phenomena as dreaming experiences and mental imagery. If that were the case, the most parsimonious account of perceptual experience may be an account appealing to inner models as well. However, while it may be true that the explanation of dreams and mental imagery must appeal to processes inside the brain, it would be a

serious mistake to presuppose that dreaming experiences and mental imagery require inner *models*.

Indeed, we may suppose that during dreaming or the experience of imagery, to a certain extent it is for the perceiver as if he or she is engaged with the environment. A skill-oriented perspective could then suppose that, in as far as similar processes are active as those that are involved in the grasp of actually obtaining sensorimotor dependencies, we may expect the perceiver to have a similar experience. The reason then is not that there is a model inside the brain of the environment or of the sensorimotor interaction. The reason is rather that the implicit grasp of sensorimotor dependencies can be exercised in absence of the usual environmental basis for these sensorimotor dependencies (cf. Thomas 1999). The existence of experiences of imagery or dreaming therefore does not support the 'inner model'-account (whether it supports a focus on inner processes is a different matter). Also in the cases of imagery and dreaming the question arises whether phenomenal experience is best viewed in terms of the specifics of the apparent mode of engagement or whether it should be analyzed in terms of the possession of inner models.[13]

A distinctive explanatory advantage of a sensorimotor account derives from the fact that it does not postulate that experience consists in the possession of inner models. The reason is that the sensorimotor account thereby simplifies our subpersonal account of how perceiving works, compared to traditional accounts that do propose inner models. If the sensorimotor account is right, then compared to inner model accounts a whole layer of subpersonal processes can be eliminated. We are then one step closer to really understanding the processes underlying perception.

In short, a sensorimotor account helps to articulate, at the personal level, what our experiences are like, and it can simplify our view of the relevant subpersonal processes. But there is also a different reason for favoring the 'skillful engagement'-framework as fleshed out by the sensorimotor account. This concerns the way in which the sensorimotor

[13] For an example showing the difference between the two frameworks, consider the different accounts of mental rotation they suggest. On both accounts, the experience of mental rotation is as if one is engaged with a rotating object, and both accounts may suppose that similar neural processes are involved as when one actually perceives a real rotating object. But on one account the reason is that that mental rotation involves an inner model of a rotating object, while the other rejects that hypothesis. According to sensorimotor theorists, mental rotation is similar to viewing the rotation of an object because the apparent mode of engagement is similar. Nothing like inner images or models is required for this.

account links the personal and subpersonal levels of description. This is the topic of the next subsection.

4.2. Sensorimotor engagement as dual currency concept

The explanatory ideal for an account of phenomenal experience is to find concepts or characterizations that are "equally applicable to the mental and the material" (Humphrey 2000, p. 10). Through such *dual currency* concepts or characterizations, personal level descriptions of experience are firmly connected to descriptions of subpersonal processes. Let us consider the way in which the descriptions may be aligned, and investigate how sensorimotor descriptions may apply to both levels of description.

Note that the ambition of dual currency characterizations may vary. Ideally, a dual currency characterization should completely overcome what Joseph Levine called the *explanatory gap* between conscious experience and the natural world (Levine 1983). But, as pointed out in Chapter 1 above, different explanatory challenges come together in the explanatory gap. There is the *absolute gap* regarding the very existence of experience and there are the *comparative gaps* concerning the specific phenomenal character of different experiences (Hurley & Noë 2003; cf. Chalmers 1996).

The question what explains that some sensorimotor dependencies pertain to perception at all – what explains our implicit grasp of sensorimotor dependencies – is relevant to the absolute gap issue: no grasp, no experience. In Section 4.1 I discussed an advantage of the sensorimotor account for approaching the issue of implicit grasp, namely that it offers a simple account of the required subpersonal processes. Still, a basic implicit grasp of sensorimotor dependencies may be necessary for conscious experience, but this does not mean that it is sufficient, and in the next chapter I shall pick up a different question relevant to the absolute gap, concerning the comparison between conscious and not conscious for a perceiver that is presently in a state of conscious awareness (Chapter 3).[14] In what follows I focus on the comparative gap questions regarding the specific phenomenal character of perceptual experience, starting with a few examples.

In Section 3.1 we have seen how personal and subpersonal descriptions can be closely related. For example, the regularities of visual occlusion and

[14] Other issues relevant to the absolute gap are the difference between being conscious and being unconscious for systems that have the capacity for conscious awareness, and the difference between systems that are capable of conscious experience and systems that are not.

optical expansion are personal-level aspects of visual phenomenology that could be given subpersonal interpretations in terms of the characteristic sensorimotor dependencies pertaining to the visual encounter of objects. Indeed, these characterizations are so closely akin that one may propose a dual currency characterization of the experience in terms of the relevant patterns of sensorimotor engagement with the environment. As Hurley and Noë have put it: "When the sensorimotor pattern characteristic of vision is explained, we have an "aha!" reaction; we see *through* the dynamic pattern of sensorimotor contingency to what vision in particular is like" (Hurley & Noë 2003, p. 160).

Other examples discussed in Section 3.1 were the 'corporality' and 'alerting capacity' characteristic of perceptual experience. The patterns of corporality – the fact that bodily movements may have large sensory consequences on sensory stimulation – can be described as subpersonal patterns specific for perception, but they equally can be used to describe what the perceptual experience is like for the person. Also the alerting capacity of perceptual experience – the fact that environmental occurrences may have a particularly strong capacity for capturing cognitive resources – can be interpreted both as subpersonal phenomena and as personal level phenomena. Corporality and alerting capacity can then be considered as specific dual currency concepts, capturing important aspects of our sensorimotor engagement with the environment.

Or consider the patterns of sensorimotor engagement that occur as one visually or tactilely explores a bottle, discussed in O'Regan and Noë (2001a; cf. MacKay 1967). In the one case there are the sensorimotor dependencies characteristic of vision, such as the eye movement-related dependencies, or the perspectival consequences of changes of position. In the other case there are the sensorimotor dependencies characteristic of touch, depending on the way in which the bottle fits in the hand and how it can be manipulated. Again, the relevant patterns may be given a subpersonal interpretation in terms of sensorimotor regularities. But also at the personal level, the sensorimotor account holds, experience is precisely a matter of such patterns. As O'Regan and Noë write: "To reflect, then, on what it is like to see the bottle, or to touch it, is to reflect on just these sorts of facts about the active engagement the perceiver undertakes with the environment" (O'Regan & Noë 2001a, p. 962).

In the examples above, we can see how sensorimotor patterns during active exploration of the environment are characteristic of perceptual experiences. It seems that in all these examples, personal-level characteristics of perceptual experience match with descriptive features at the subpersonal level. The sensorimotor account then provides a dual currency characterization of perception. We can further sharpen our view

of the proposed explanation by considering how the sensorimotor account characterizes more passive cases of perception, in which the perceiver is not currently performing exploratory behavior.

As we saw in Section 3.2, the sensorimotor account proposes that experiencing lies in the exercise of one's implicit grasp of sensorimotor dependencies. At the personal level, the perceiver implicitly grasps how bodily activity would alter the sensory situation. For example, the visual experience of an object occluding its background is proposed to consist in the perceiver's implicit knowledge, or expectancy, that if he or she were to move, occluded parts of the background would come in view. The sensorimotor dependencies pertaining to the situation can be given a subpersonal interpretation in terms of the way in which sensory stimulation would change if one were to move. Thus the sensorimotor situation can be understood at the personal level as well as the subpersonal level, also in cases where no overt bodily movement is involved.

Now it is true that at present we have no worked out neuroscientific explanation of the grasping of the obtaining sensorimotor dependencies. But note that it is not something about a particular way of grasping sensorimotor dependencies that is proposed to explain the specific quality of experience. Instead, the grasped sensorimotor dependencies themselves are proposed to explain the phenomenal character of experience. (Recall that the sensorimotor dependencies relevant to perception are those over which one's implicit grasp is currently exercised, but that the characterization of the experience appeals to the patterns of sensorimotor dependencies themselves.) It is in terms of the characteristic patterns of sensorimotor engagement that personal and subpersonal levels of description are captured within a single sensorimotor account.

To explain *which* aspects of our sensorimotor situation are relevant to perceptual experience, we must consider over which sensorimotor dependencies our implicit grasp is currently exercised. Detailed knowledge of the brain will be crucial to explain how the exercise of the implicit grasp of sensorimotor dependencies is realized. But the properties of neural activity do not in themselves explain the phenomenal character of experience. According to the sensorimotor account, the patterns of skillful sensorimotor engagement with the environment, in which the brain participates, explain what our perceptual experiences are like. Because these patterns of sensorimotor engagement can be given personal as well as subpersonal interpretations, the sensorimotor account then provides a dual currency explanation of the phenomenal character of perceptual experience.

4.3. Conclusion

Perceptual experience relies on sensory input as well as on motor action. Most obviously, our perceptual experience depends *instrumentally* on behavioral skills, such as the capacity to track moving objects with our eyes. But action may also modulate perception *noninstrumentally*, that is, in a way independent of the possible changes of sensory stimulation, presumably through the workings of proprioceptive feedback or corollary discharge signals. Sensorimotor approaches acknowledge both types of action-dependence of perception by characterizing perceptual experience in terms of the specific relation between sensory stimulation and motor action, the sensorimotor dependencies, during perceptual engagement of a perceiver with the environment.

Above I discussed a sensorimotor account which views perceptual experience not as an inner model, but as a skillful *mode of engagement* with the environment. In particular, this sensorimotor account proposes that perceptual experience is a matter of *exercising* one's *implicit grasp* of *sensorimotor dependencies*. The extent to which this skillful engagement plays certain roles in our lives determines the extent to which conscious experience obtains. I explicated this account by emphasizing the way in which it relates personal and subpersonal descriptions of phenomenal experience. I argued that, at the personal level, the account helps to articulate what perceptual experiences are like. A crucial advantage of the sensorimotor account is that it simplifies our view of the processes at the subpersonal level, compared to standard accounts which hypothesize inner models. But the most fundamental contribution of the sensorimotor account to the explanation of the phenomenal character of experience, I suggest, concerns the way in which the account links personal and subpersonal levels of description.

A sensorimotor perspective on perceptual experience allows for a description of subpersonal processes that closely matches the way we experience the world. If the account is right, it offers a *dual currency* explanation of the phenomenal character of experience, a characterization that can be interpreted in personal level as well as in subpersonal terms. This characterization is given in terms of the sensorimotor engagement of perceivers with their environment. More specific dual currency concepts that have been proposed are the corporality and alerting capacity characteristic of perceptual experience.

The explanatory potential of the sensorimotor account is closely connected with its skill-oriented perspective. The reason is that this perspective allows for a characterization of subpersonal processes that remains close to the level at which we appear to live our lives. While it is

hard to see what local properties of neural processes have to do with our experience, a description of the extended sensorimotor dynamics, in which neural processes participate, more readily matches our experience. Indeed, it seems that our sensorimotor engagement with the environment is precisely what we should reflect on when we think about the phenomenal character of our perceptual experience.

Chapter 3
Workspace and sensorimotor theories: complementary approaches to experience

This chapter is joint work with Fred Keijzer:
Degenaar, J. & Keijzer, F. (2009) *Journal of Consciousness Studies* 16 (9), pp. 77-102.

A serious difficulty for theories of consciousness is to go beyond mere correlation between physical processes and experience. Currently, neural workspace and sensorimotor contingency theories are two of the most promising approaches to make any headway here. This chapter explores the relation between these two sets of theories. It is argued that workspace- and sensorimotor theories are complementary rather than competitive. By combining these theories, a number of problems that hamper these individual theories may be overcome and their strengths combined: workspace theories have more to offer for explaining how there can be consciousness in the first place, while sensorimotor theories are strong in making sense of the specific phenomenal character of experiences.

1. Absolute and comparative gaps

There is little doubt that conventional scientific approaches are able to find reliable correlations between neural activity and conscious experiences. The challenge is to provide more than mere correlations. We discuss the merits of two of the most promising proposals for increasing our understanding of experience currently available: neural workspace theories – or workspace theories for short – and sensorimotor contingency theories – or sensorimotor theories for short. Both come in various forms, and are sets of theories rather than specific theories. We will focus on the commonalities within these two sets of theories to assess the potential for combining the proposals.

In this chapter we will assume that the approaches we discuss can increase our understanding of the physical basis of experience in a way that goes beyond mere correlation. It has been argued that conventional

scientific approaches are unable to truly explain why physical processes are accompanied by phenomenal experience, a problem known as the hard problem of consciousness (Chalmers 1995). Rather than making any direct claims on to the hard problem, our focus will be on the increase in explanatory power that a combination of theories may yield compared to the individual sets of theories.

Neural workspace theories provide a set of closely related theories, which seem promising to make significant headway toward a satisfactory empirical theory of the physical bases of experience (Baars 1988; 2002; Tononi and Edelman 1998; Dehaene & Naccache 2001; Varela *et al*. 2001). Neural workspace theories hypothesize that conscious experience depends on a coherent pattern of neural activity that facilitates the availability of information for various processes. The mechanisms that give rise to this pattern of activity could potentially account for important features of experience. The hypothesis may even come to explain why some neural activity is correlated with consciousness while other neural activity is not.

Sensorimotor contingency theories hold that the phenomenal quality of experiences can be understood in terms of the characteristic relations between sensory input and motor action – the 'sensorimotor dependencies' or '-contingencies' (O'Regan & Noë 2001a,b; Hurley & Noë 2003; O'Regan, Myin & Noë 2005; Mossio & Taraborelli 2008). Within these theories, differences between for example visual and auditory experiences are thought of as differences in the sensorimotor dependencies, or differences in the mode of active exploration of the environment. One of the characteristics of visual experience is the way in which movement of the head enables us to look behind objects. While eye movements thereby result in large shifts of the retinal image, we experience the visual world as stable. From findings like these, sensorimotor theorists conclude that our experience is not an inner construct based on input alone, but is rather constituted by displaying the implicit knowledge of the input-output relation.

Much of present-day theorizing about conscious experience, including workspace theory, tends to focus on neural activity. Sensorimotor theories, in contrast, focus on whole patterns of interaction involving brain, body and environment.[1] The idea is that knowledge of neural activity alone may

[1] While sensorimotor theory has affinity with Velmans' idea that the world as-perceived is out-there (e.g. Velmans 1990), we like to emphasize the difference in focus. Velmans' 'reflexive' model of experience is inclined to take a brain-focused approach to the processes underlying experience, taking the contents of some experiences as a 'projection' in space by the brain (Velmans 1990; 2007).

not be enough for an explanation of conscious experience. We will refer to this claim of sensorimotor theories as *externalism* about the processes underlying consciousness and contrast it with the brain-focused *internalism* of workspace theories. For a more elaborate discussion of internalism and externalism with respect to the processes that figure in the explanation, see Hurley (2010). For present purposes, it is important to distinguish the explanatory externalism of sensorimotor theories from the more familiar representationalist commitment about the content of experience, namely that we experience aspects of the external world (e.g. Dretske 1995). The explanatory externalism of sensorimotor theories holds that the *processes* we need to take into account to explain our experiences extend into the world, and this does not necessarily imply that we experience objective features of external objects.

Sensorimotor- and workspace theories are both major players in present-day consciousness research. Still, little work has been done on their relation. The defenders of sensorimotor theories have sometimes stressed the differences with brain-based explanations rather than searching for ways to link sensorimotor theory with explanatory paradigms like workspace theories. This differentiation has been useful to emphasize the particular and independent contribution of sensorimotor theories to the explanation of consciousness. At the same time, workspace theorists may have been disheartened by the externalist tendencies of sensorimotor theories. They may have seen little reason to relate their dominant brain-based theory to this set of relatively new theories. The general differentiation between internalist and externalist explanations (Hurley 2010) – according to which both theories can be categorized – may have given the impression of intrinsic opposition. We claim that this impression is false.

At the basis of our claim lies the distinction between two fundamental problems of experience (Chalmers 1996, p. 5; Hurley & Noë 2003). The first problem is to understand the very *existence* of conscious experience, generally known as the *absolute gap*: why do we have conscious experience at all? The second is to understand the *character* of conscious experience: why do experiences have the specific qualities that they have? This problem is known as the *comparative gap* or *-gaps* (Hurley & Noë 2003). Examples are the problem to explain experiential differences between sensory

---

Sensorimotor theories stress that an understanding of experience requires a focus on whole patterns of interaction with the (third-person identifiable) environment. Whether the physical constitution of experience is purely brain-based is a further question, which we touch upon in Section 4.2 below.

modalities (e.g. seeing versus hearing) and within modalities (e.g. the experience of red versus the experience of blue). Explaining consciousness consists at least in solving both the absolute and the comparative gap problems.

In this chapter we will argue that sensorimotor- and workspace theories can each be cast as best dealing with one of the explanatory gaps. If this is right, a combination of the theories becomes highly desirable. We discuss a potential difficulty for reconciliation of the theories: workspace and sensorimotor explanations have respectively strong internalist and externalist tendencies, which reflect deep differences in theoretical and metaphysical views. As a way to deal with these differences, we sketch three different scenarios for combining workspace- and sensorimotor theories, each reflecting different fundamental outlooks on experience. In all three cases, combining workspace- and sensorimotor theories plausibly leads to an increase of explanatory strength compared to each of the separate theories.

2. Sensorimotor contingency theories of experience

Sensorimotor theories offer an empirical approach of the character of conscious sensory experience. The theories aim to answer the question that Chalmers stated thus: "Given that conscious experience exists, why do individual experiences have their particular nature?" (Chalmers 1996, p. 5). To answer this question, sensorimotor theorists focus on whole loops of interaction involving brain, body and world. Since our intuitions may need some stretching before we can appreciate an explanatory role for processes outside the brain, some preliminary remarks will be useful.

2.1. Not by neural activity alone

Hurley and Noë state that "neural properties are qualitatively inscrutable" (Hurley & Noë 2003). Indeed, it seems doubtful that the character of the experience of red or the taste of coffee can ever be explained in neuroscientific terms (Levine 1983; 1993). Nevertheless, most people are convinced that experience occurs within the brain and that it should be explained in terms of neural activity. Hurley observes conflicting intuitions:

> "If someone really has no conception of how neural or internal functional properties—or indeed any others—could explain phenomenal qualities, then how can he be so confident that *if* phenomenal qualities can be explained, it must be internal factors that do the job?" (Hurley 2010, p. 104)

Given the difficulty to understand the character of experiences in terms of inner states and processes, we may need to challenge the internalist

assumption. After all: "Neural processes are normally in continuous dynamical interaction with external factors; there's nothing magical about the boundary between them" (Hurley 2010, p. 126). As there is no 'magical membrane' to separate the brain from the rest of the world, there is also no reason for seeking the explanation of consciousness in the brain alone, the externalist argues. Thus, the intuition that consciousness must be something within our bodies or even brains can be neutralized by an opposing intuition that consciousness is not intrinsically related to anything within the body.

When intuitions on their own are not sufficiently trustworthy, what kind of evidence do we have for an internalist interpretation of consciousness? Some theorists argue that the existence of rich experiences which are seemingly 'off-line' does provide such evidence (e.g. Koch 2004; Revonsuo 2006; Prinz 2008). As experience is possible in relative isolation from the environment, one can argue that brain processes suffice for consciousness while the environment can only modulate these inner processes. In reply, we will turn to the phenomenon of dreaming, which is a show case example of this line of argument.

While dreaming may seem to provide an obvious example of strictly inner experience, the case is actually not clear cut. First, the existence of rich off-line experience does not imply that the best explanation of on-line sensory experiences must be based on internal processes alone. As Hurley says: "If the enabling role of internal simulations in off-line cases is derivative from their role in extended dynamics, it provides no reason to hold that only internal processes can do quality-enabling work in the primary, on-line cases" (Hurley 2010, p. 142). The explanation of the qualitative character of dreaming experiences may ultimately derive from the explanation of sensory experiences rather than the other way around. Thus a purely neural account can be incomplete even for illusory cases and a weak form of externalism may be needed for dreaming. Second, on-line and off-line experiences may be qualitatively different (Putnam 1999, p. 130; Noë 2004, p. 213-214; Noë & Thompson 2004). Waking experiences do not typically have 'a dream-like quality' (Austin 1962, p. 48-49). As a result, the processes that we need in order to explain the experiences may be different too, even if they partly overlap. Third, even though dreaming takes place without outward action, this does not imply that input from the environment or feedback loops extending in the body play no role. Our brains are never completely off-line as we can be wakened by noise, shaking or other stimuli. Maybe we need to take active external processes into consideration even to account for the phenomenal contents of dreams.

In addition, a positive reason for taking the externalist possibilities seriously comes from the general trend towards more dynamical, embodied

and embedded explanations of cognition (e.g. Clark 1999; Calvo & Gomila 2008; Robbins & Aydede 2008). Instead of decomposing dynamically interacting systems into different parts, more global and environmentally extended patterns of interaction may be involved in the workings of cognitive processes (Clark 1997; 1999) and behavior (Keijzer 2001). Clark and Chalmers (1998) even introduced the notion of an *extended mind* that is literally extending into the environment. If externalism can be accepted for the processes underlying cognition, we see no reasons for a general ban *beforehand* on developing similar ideas for experience. Hurley and Noë suggest: "To find explanations of the qualitative character of experience, our gaze should be extended outward, to the dynamic relations between brain, body, and world" (Hurley & Noë 2003, p. 132). Let us turn to the sensorimotor hypothesis and see what this brings.

2.2. Sensorimotor theory: experience in interaction

Sensorimotor theories of perceptual experience state that experiencing is best characterized as exercising our mastery, or implicit knowledge, of patterns of sensorimotor interaction with the environment (O'Regan & Noë 2001a,b,c; Hurley & Noë 2003; Noë 2004; O'Regan, Myin & Noë 2005). To unpack this idea, we will successively discuss the role of action, the role of the environment, and the appeal to implicit knowledge. For purposes of illustration we will mention evidence from sensory substitution.

A basic idea of sensorimotor theories is that sensory experiences are in general strongly action-dependent. If retinal stimulation shifts in the absence of eye-movements, typically the world would appear to move (or the perceiver would appear to move). However, if the same pattern of retinal stimulation occurs as the result of an eye-movement, this does not impair the apparent stability of the visual world (or perception of the perceiver's own location). Thus, action has more than an instrumental role in experience; it does not 'merely' change the input: it can directly change visual experience (Hurley 1998). Experience is not an inner construct based on input alone.

To capture the qualitative character of sensory experiences an appeal is made to the specific way in which sensory input depends on motor action. A clear example is the experience of the softness of a sponge. According to sensorimotor theories we can understand the character of this experience if we consider the sensory consequences of motor action. For example, when we squeeze a sponge it gives little resistance and this is what its softness consists in. The experience of the hardness and softness of surfaces is not action-neutral; it rather consists in our grasp of the sensory effects of our actions (O'Regan, Myin & Noë 2005). These dynamical

patterns of sensorimotor contingencies are specific to the particular senses; feeling a sponge and seeing one have their own, different contingencies. Sensorimotor theories aim to explain the experiential differences between modalities in this way.

Appealing to sensorimotor dependencies brings the environment into the explanation. The experience of sponge-squeezing is explained by the characteristic ways in which the sponge responds to pressure. Similarly, to understand the visual experience of objects in space, the relevant sensorimotor contingencies obtain as a result of the spatial orientation of objects and the reflective behavior of light. For example, the distance one has to move to look behind an object depends on the relative distances of objects and perceiver. Sensorimotor theories claim that we experience the spatial relations between objects by exercising the implicit grasp of the sensory consequences of movements. For this reason, we need to take the environment into account if we are to understand perceptual experience.

Of course, sensorimotor theory does not require that we must always move in order to experience. We often see at a glance that one object is closer than another. By involving our implicit knowledge, sensorimotor theories can deal with experience in the absence of movement. Once we have the practical familiarity with the sensorimotor contingencies, we make use of this familiarity in our experience of the world. Whether we move or not, we implicitly grasp what sensory consequences are to be expected if we would make a certain move. For example, when we see the spatial orientation of objects, this experience constitutively depends on our mastery over the governing laws of sensorimotor contingency (O'Regan & Noë 2001a). Thus, according to sensorimotor theories, experiencing is a skillful activity of sensorimotor interaction with the world (O'Regan & Noë 2001a; Myin & O'Regan 2002).

A good example that illustrates how sensorimotor theory goes beyond a brain-based focus comes from studies of sensory substitution devices, such as Bach-y-Rita's tactile-visual substitution system (Bach-y-Rita 1984; 2002). This system transforms the image recorded by a camera into a tactile display, e.g. an array of vibrating pins which can be applied to the subjects back. In studies with such a devise, otherwise blind persons report vision-like experiences rather than tactile ones: they experience objects as being at a certain distance and they report experiencing spatial relationships between objects, such as that 'one is partially blocking the view' of another object. In a recent study using a different sensory substitution devise, an auditory-visual substitution system, Auvray *et al*. (2007) tested blindfolded sighted subjects, finding that in some cases subjects reported visual experiences despite the auditory input. Importantly, when subjects have no control over the camera – when

someone else controls the camera, or when it is stationary – the change from tactile to semi-visual experience does not occur (Bach-y-Rita 1984). This fits very well with sensorimotor theory as in this case there are no new sensorimotor contingencies to be mastered by the subject, but only passive sensory stimulation. The subjects will not acquire the practical mastery of the sensorimotor contingencies relevant to 'seeing' with the device (Hurley & Noë 2003).

Bach-y-Rita concluded from his findings that "we do not see with our eyes, but with our brain" (Bach-y-Rita 2002, p. 497). However, referring to the brain does not explain the differences enabled by the device, as we can also be said to hear, feel or taste with our brain. In contrast, the change in the sensorimotor contingencies does explain why a tactile modality acquires vision-like experiential features (Hurley & Noë 2003).

Sensory substitution provides an example of how sensorimotor contingencies can help to explain differences in the qualitative feel of sensory modalities, and possibly even how new modalities can arise (Auvray & Myin 2009). The sensorimotor interactions between an agent and its environment provide a systematic constraint on experience. Brain functioning is shaped by our active encounters with the environment, and sensorimotor theorists stress that it should be considered in the context of these temporally extended patterns of interaction.

2.3. Explanatory promise and limitations

Given this short description of sensorimotor theories, what can be said about their strengths and weaknesses? In particular, how does sensorimotor theory relate to the absolute and comparative gaps of consciousness?

Sensorimotor theories seem particularly strong on comparative gap issues, such as the experiential differences between different sensory modalities (Hurley & Noë 2003). Sensorimotor theories may even provide a handle on dealing with new sensory modalities (Auvray & Myin 2009), and some of the aspects of differences within modalities (Hurley & Noë 2003), such as between colors (Philipona & O'Regan 2006). In addition, a sensorimotor perspective has been applied to differences between conscious thought and sensory experiences (O'Regan, Myin & Noë 2005). Sensorimotor theory arguably provides explanations that are more than 'mere correlation'. From a sensorimotor perspective one can understand why, e.g., visual experiences differ from tactile experiences in the way they do. In contrast, neural activity in a certain area may be reliably correlated with a particular experience, but this correlation would not explain why a particular phenomenal experience is associated with this activity. Rather

than accepting intrinsically qualitative properties of neural activity with all its difficulties, sensorimotor theories characterize differences in experiences as differences in the dynamical patterns of agent-world interactions. In this way, sensorimotor theory also opens up the possibility of an evolutionary understanding of the origins and differences between sensory modalities.

Sensorimotor theorists have occasionally made claims concerning the absolute gap, stating that they explain the *presence* of experience (see especially O'Regan & Noë 2001a, pp. 1011-1012). However, as also Noë (2004, pp. 228-231) recognizes, sensorimotor theory seems less convincing as an account of the switch from non-conscious processes to conscious ones. Sensorimotor contingencies are used by widely different systems, ranging from organisms like insects to robots, not all of them plausibly interpreted as experiencing beings. In addition, sometimes full-blown conscious experience is not present in humans despite the exercise of mastery of sensorimotor interaction with the environment. For example, when driving, absorbed in thought or conversation with a friend, you may hardly experience aspects of the environment that are used to guide your behavior. Thus, making use of sensorimotor contingencies in itself does not seem to suffice for the presence of experience.

In their reply to the 'unconscious' driving problem, O'Regan and Noë agree that an extra ingredient is required. They write:

> "A driver (…) would be said to be aware of a red traffic light if, in addition to the mastery of sensorimotor contingencies associated with the red light, his attunement to these sensorimotor contingencies is integrated into his planning, rational thought or linguistic behavior. Depending on the extent to which the seeing of the red light is incorporated into his planning or thought, the driver would be said to be aware of the red light to varying degrees." (O'Regan & Noë 2001c, p. 94)

It may be that sensorimotor dependencies are only relevant to consciousness if they play a certain role in planning, thought or speech. However, this rather seems to describe the presence of consciousness than to explain it. As a theory of the presence of experience, sensorimotor theory seems too descriptive. As also Rowlands (2003) and Noë (2004) suggest, sensorimotor theory may not in itself be able to explain the presence of experience, but it rather has to presuppose it.

Thus, it appears that sensorimotor theory is well equipped to deal with comparative gap problems that arise once consciousness is present and enables one to make sense of various qualitative differences. However, it seems more difficult to address the absolute gap with the means provided by sensorimotor theory.

3. Workspace theories of consciousness

Conscious experience implies the availability of information: you will be able to tell when you are conscious of a stimulus, at least if you have the capacity to speak. Neural workspace theories form a set of theories that aim to identify the underlying neural mechanisms that can explain the conscious availability of information.

The central idea of neural workspace theories is that consciousness-correlated neural activity forms *a coherent pattern of neural activity that makes information globally available throughout a neuronal workspace* (Baars 1988; 2002; Tononi & Edelman 1998; Dehaene & Naccache 2001; Varela *et al*. 2001; Metzinger 2003).[2] This 'workspace' is 'a central information exchange that allows many different specialized processors to interact' (Baars 1988, p. 43). By hypothesis, information that reaches this workspace will influence the processing in large parts of the brain, a bit like the information on a blackboard being available for the whole class.

Workspace theories claim that a stimulus will influence conscious experience if and only if it modulates the activity in the neural workspace. Based on criteria of availability of information for the person, typically the availability for verbal report, it offers a theory of the neural basis of experience. This theory is usually formulated in terms of the availability of information to subsystems. It builds on the subpersonal availability of information that is implied in workspace activity and, closely related, the subsequent *role* that the activity plays. This sits well with Dennett's philosophical views on consciousness. As in Dennett's (1991) theory, workspace theories approach consciousness not as an intrinsic feature of neural activity: rather it is because of the use that is made of information that it classifies as conscious. Indeed, a neural workspace can flesh out Dennett's idea of consciousness as 'fame in the brain' (Dennett 2001).[3]

[2] Neural workspace theories differ in their stress on neural or informational aspects, some being almost exclusively formulated in informational terms, like Baars' original formulation, others in neural terms, as in Varela *et al*. Accounts that are focused on neural processes are often given an informational interpretation, as in Tononi and Edelman's models. The similarities are strong as the neurally oriented theories remain committed to some form of information processing, while those in informational terms presume neural information processing mechanisms for their implementation.

[3] In as far as Dennett's positive theory of consciousness occupies the same explanatory niche as workspace theories, it may have similar strengths and weaknesses.

However, note that workspace theory is not necessarily committed to such an interpretation. Although the theory is based on behavioral criteria of personal-level availability, this by no means excludes the possibility to acknowledge intrinsic experiential features of neural activity. The neural workspace can be – and sometimes is – seen as 'the place where consciousness happens'.

Note that availability of information for perceptual report is not a simple criterion to judge the presence of experience. As Metzinger (2003, p. 75) points out, consciousness may come in degrees, depending on the extent to which information becomes available. He also differentiates between three dispositional properties that can exemplify availability; information can be available for guided attention, for cognitive processing, and for behavioral control. Thus, to study the neural basis of consciousness we may have to ask exactly in what sense information becomes available (Metzinger 2003, p. 124). Another difficulty is that availability is not always sufficient for experience: information may in some cases be available only unconsciously – as in the case when you respond adequately but unconsciously to a stimulus. However, when we put problematic cases aside there will remain enough reasonably uncontroversial cases that can be used in this empirical approach of experience.

3.1. Consciousness as global cortical activity

There is ample evidence that cognitive processes often occur without associated conscious experience (Dehaene & Naccache 2001; Merikle & Daneman 1999). The challenge is to determine whether there is a systematic difference between consciousness-correlated and not consciousness-correlated processing (Dehaene & Naccache 2001). Some neural activity will be specifically associated with experience and the question is how this activity differs from the activity that isn't.[4] Neural workspace theories offer a possible answer.

A good example of a workspace hypothesis is provided by Tononi and Edelman's (1998) 'dynamic core hypothesis'. Interestingly, they start from the *character* of experience, to hypothesize on the nature of the neural processes that underlie consciousness:

[4] Some theorists, in contrast, have suggested a link between life and experience (e.g. Noë 2004; Thompson 2007). For example Noë speculates that "living beings are already, by dint of being alive, *potentially conscious*" (Noë 2004, p. 230). Workspace theories address the contrast between conscious and unconscious processes *within* living beings.

> "(…) our strategy is to characterize the kinds of neural processes that might account for key properties of conscious experience. We emphasize two properties: conscious experience is integrated (each conscious scene is unified) and at the same time it is highly differentiated (within a short time, one can experience any of a huge number of different conscious states)." (Tononi & Edelman 1998, p. 1846)

The *dynamic core hypothesis* proposes that which neurons are part of the 'dynamic core' can rapidly change and that the 'dynamic core' is the neural activity that correlates with consciousness. Tononi and Edelman hypothesize that the unity of experience can be explained by the 'functional integration' of the relevant neuronal activity: "at a given timescale, these elements interact more strongly among themselves than with the rest of the system" (Tononi & Edelman 1998, p. 1848). The differentiated character of experience is proposed to be reflected in the 'complexity' of the activity of the dynamic core, which is a function of the amount of mutual information that subsets of the dynamic core share with the rest of the core (Tononi & Edelman 1998).

The intensive 'cross-talk' between the neurons within the dynamic core, or more in general within a neural workspace, should ensure that each part of this workspace is influenced by the other parts. If a part of the workspace activity carries information about a certain aspect of the environment, this results in the global influence of this environmental feature throughout the workspace. According to a representationalist analysis (e.g. Metzinger 2003), it results in the availability of this information for other subsystems. The workspace activity could potentially stretch out to areas devoted to speech so that perceptual reports can come under the influence of the environmental feature that modulates workspace activity.

Presumably, neurons throughout large parts of the brain can be part of the workspace. There are various hypotheses on the specifics of the workspace. For example, Dehaene and Naccache (2001) propose that specific 'workspace neurons' with long-distance connectivity form a neural workspace. If such neurons are sufficiently activated, they will result in brain-scale coherent activity that makes information available throughout the workspace. Tononi and Edelman (1998), in contrast, offer a more dynamical view in which it is possible that at one moment a neuron is strongly activated without being part of dynamic core, while at another moment it is part of the dynamic core. Several authors have argued that in addition to widespread cortical neurons also neurons in the thalamus may be involved (which is consistent with the re-entrant connections in the

thalamocortical system) (Tononi & Edelman 1998; Dehaene & Naccache 2001).

Much has been written elsewhere on the evidence that is in agreement with workspace hypotheses (Dehaene & Naccache 2001; Baars 2002; Varela *et al*. 2001). An important finding is that *neural synchrony* is correlated with conscious experience (Engel *et al*. 1999; Engel & Singer 2001; Varela *et al*. 2001). When something is consciously experienced, an associated increase of the synchronous firing of neurons may be found throughout large parts of the brain. This indicates a high level of interaction between neurons, which is exactly what workspace theory predicts.

3.2. Explanatory promise and limitations

How well does workspace theory fare with respect to the absolute and comparative gaps of consciousness? We submit that since workspace theories aim to clarify the preconditions for conscious experience, they are chiefly focused on issues relating to the absolute gap.

Workspace theories aim to account for the presence of consciousness by explaining how the neural activity that underlies consciousness differs from activity that is not directly involved in consciousness. The basic idea is that certain ways of responding to a stimulus implicate awareness of it: when you can report that you heard a noise, you have experienced it. Neural workspace theories aim to explain features that are descriptive of experience. They do so in terms of underlying neural activity.

In their account of the mechanisms that enable persons to consciously perceive the world, workspace theories tend to use a subpersonal notion of information and they speak of the availability of information for parts of the brain. In particular, it is proposed that the workspace forms part of the subpersonal mechanisms by means of which information about the environment can become available for the subject. A way to construe an informational interpretation of subpersonal processes is as a third-person, correlation-based ascription of information. Note that such an ascription of information plausibly depends on more than neural activity alone, indeed it may only make sense in the larger sensorimotor context.

Workspace activity could make perceptual information globally available for thought and action. If a stimulus influences workspace activity, this can directly influence behavior and information related to the stimulus will be available for report (Dehaene & Nacacche 2001 pp. 21-22). The neural processes involved in the workspace will have a special influence on further lines of thought, action, and speech, as a direct result of the physical/functional properties of the workspace. Unconscious workspace activity is impossible, because (above a certain level) workspace activity

just leads to global availability of information for neural subsystems, and this implies that the information is available to modulate behavior such as verbal report. (One way to construe the subpersonal 'availability of information' is as a shorthand for e.g. parts of the brain coming under the influence of a certain aspect of the environment to which the person is sensitive.) Thus, workspace theories clearly aim to go beyond mere correlation: they potentially provide an explanation in which the presence of particular brain processes *implies* key psychological features signifying conscious availability.

At the same time, a detailed account is still lacking of how workspace activity can have the specific effect it has. For example, it remains to be clarified how 'speech centers' are influenced by visual processes correlating to the presence of a butterfly in such a way that it enables the person to tell others that there is a butterfly. While such an account is missing, we submit that the most important explanatory promise of workspace theory concerns the contrast between consciousness-correlated and not consciousness-correlated processes. This contrast is addressed in terms of differences in subpersonal availability, or 'fame in the brain'.

Workspace theorists do have also aimed to address the character of experience – the comparative gaps. For example, Tononi and Edelman (1998) suggested that properties of the dynamic core could account for the differentiated character of experience – the higher the complexity of the workspace activity, the more differentiated the experience. In addition, more daring attempts have been made to address the specific character of sensory experiences (Edelman & Tononi 2000; Tononi 2004). To approach comparative gap problems such as the differences between visual and auditory experiences, these authors appeal to the 'discriminations' made within the dynamic core. One difficulty with this attempt concerns the required understanding of the neural activity in terms of 'discriminations'. Suppose this understanding is dependent upon the theorist's knowledge of the larger pattern of interaction with the environment within which the discrimination plays its role. In that case, the understanding of experience is based in this larger pattern rather than in the workspace activity. Another difficulty is that there seems to be no intelligible link between the specific phenomenal character of experience and a set of discriminations. It is far from clear why a certain set of discriminations would result in a visual experience rather than an auditory experience, and how the experience of blue can be characterized by the way blue stimuli can be discriminated from other stimuli. Focusing on neural activity alone makes it hard to see why this activity is associated with particular experiences.

This problem becomes aggravated when workspace theories are compared to sensorimotor theories on this count. The latter theories help

to understand why and even how vision and touch constitute different forms of experience in a way that is not available to workspace theories.

To conclude, workspace theories are well-equipped to address the neural mechanisms that underlie conscious experience. They can potentially explain the difference between consciousness-correlated and not consciousness-correlated processes – a difference that is important to deal with the absolute gap of consciousness. However, the specific character of experience seems to be less approachable in terms of workspace activity.

4. Compatibility of the theories

Workspace- and sensorimotor theories thus have complementary strengths and weaknesses with respect to the absolute and comparative gaps. Given this circumstance it would seem to be a good strategy to combine the two theories, turning them into a unified framework that keeps the strong aspects of both theories and applies them to the separate gaps. However, workspace theories and sensorimotor theories are usually seen as competitors that aim to provide *different* explanations for human experience.

The division of labor suggested here is not a standard interpretation. Some defenders of both theories have claimed to address both the absolute and the comparative gaps. For example, O'Regan and Noë (2001a) enter the natural territory of workspace theory when they claim that sensorimotor theory explains the existence of experience. – They suggest that sensorimotor theory solves the problem of the absolute gap, in as far as experience is constituted by its qualities (O'Regan & Noë 2001b, pp. 1011-1012). – The opposite also holds true: workspace theory is sometimes used in an attempt to make sense of the specific character of experiences. For example Edelman and Tononi (2000) and Tononi (2004) suggest that the quality of experiences can be thought of as the discriminations that are made within the dynamic core.

However, against these claims, it should be noted that there is no intrinsic theoretical need to apply either of these theories to both gaps. As we argued above that in both cases the weak aspects of the one theory coincide with the strong points of the other, dividing up the territory in a combined effort is beneficial. In our view, a more fundamental difficulty for a profitable combination of workspace- and sensorimotor theories comes from deep differences in theoretical outlook involved and the very interpretation of what an explanation of consciousness amounts to. This is particularly so for the issue of localization: the question which processes, if any, are constitutive of consciousness.

In the following section, we will first discuss the issues at stake in combining internalist and externalist explanations. In section 4.2, we will turn to the issue of localization and we will sketch three different scenarios for a combination of workspace- or sensorimotor theories.

4.1. Internalist and externalist explanations

Are the ways in which workspace- and sensorimotor theories aim to explain different aspects of experience compatible? The first issue at stake concerns the commitments to respectively internalist and externalist forms of explanation. Should we approach conscious experience as something that takes place inside the head or not?

From the perspective of a workspace theorist, the discrepancy between internalist and externalist methodologies may at first not be so obvious: no one ever denied that workspace activity is embedded in a body interacting with an environment and that this impinges on consciousness, even if only via sensory input. However, sensorimotor theories envision a much more important role for the interaction with an environment than merely 'impinging'. Indeed, sensorimotor dependencies are cast as the *key* feature behind consciousness. The sensorimotor dependencies that arise from the interaction with the environment are ongoing shaping factors for brain processes. Without these shaping factors, there would not be consciousness as it occurs in normal human beings. Sensorimotor theorists claim that the explanation of consciousness needs to build on the dynamical patterns of sensorimotor interaction with the environment, and that experiencing is best seen as an activity (O'Regan & Noë 2001a). From this perspective, a purely brain-focused theory does not even address the problems that should be addressed. Can these internalist and externalist forms of explanation be reconciled?

A first positive reflection is that ultimately workspace activity is of course part of a larger pattern of interaction. And as sensorimotor explanations appeal to more extensive patterns of sensorimotor interaction, workspace theory can be cast as a subset of *neural* processes that co-constitute this interaction. Influencing the sensorimotor aspects of experience may be closely coordinated with workspace activity. The latter could even be cast as the very mechanism that makes sensorimotor dependencies relevant to planning, thought and language, as required by O'Regan and Noë (2001a). Thus, if both theories are correct, then those patterns of sensorimotor contingencies that underlie our experience – that is those over which we are actively exercising our mastery – are those patterns that involve workspace activity.

Against this positive reading, workspace theorists sometimes claim to characterize the processes that are *directly* involved in consciousness, while external influences work only to the extent that they impinge on the workspace. This may be read as contradicting the idea that the whole pattern of interaction with the environment is intrinsically relevant for the character of consciousness. However, we think this tension can be eased. Instead of casting the workspace as a central inner conscious domain, separated from peripheral processes that are not directly involved in consciousness (Figure 3.a), workspace theory allows a different interpretation in which the connections of the workspace with certain sensorimotor processes are co-constitutive of the activity directly relevant to consciousness (Figure 3.b).[5] Note that our point right now is only that this is a *possible* reading of workspace theory, which would lead to a possible combination of the explanatory means available to both sets of theories.

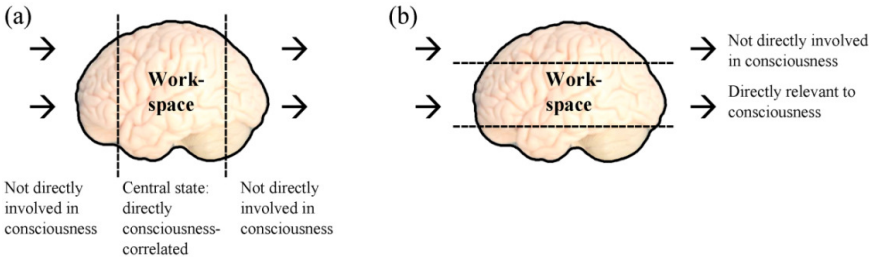

Figure 3. Schematic representation of two interpretations of the neural workspace. (a) Consciousness is interpreted as a central state. Peripheral and environmental processes are seen as not directly involved in consciousness. (b) The processes directly involved in consciousness can include processes outside the workspace. Patterns of input-output coordination that do not involve the workspace are not directly involved in experience.

The plausibility of such a combination of the theories is reinforced by the fact that a workspace approach can easily accommodate the basic *action-dependence* of experience that is stressed by sensorimotor theory. Workspace theories are informational theories of neuronal dynamics, and as such they are not committed to a particular conception of the origins of

[5] In effect, the distinction is between a vertically modular 'sandwich model' reading and a dynamical reading in terms of horizontal modularity of the processes directly involved in consciousness, similar to Hurley's (1998) horizontal/vertical modularity distinction.

the integrated patterns of neural activity. Nothing requires workspace activity to be under influence of input alone. Indeed, some work has recently been done on the integration of sensorimotor dynamics with a workspace perspective. For example, Shanahan (2006) proposes a model in which internal simulation of sensorimotor interaction with the environment is implemented within a workspace framework. Functional integration of a workspace with sensorimotor dynamics therefore seems not particularly problematic.

Some defenders of workspace theories have already stressed the need for a strong link between neural dynamics and sensorimotor coupling. Thompson and Varela (2001), for example, propose that the neural substrate of consciousness consists of 'large scale dynamical patterns' of neural activity, and that these should be considered in a broader context of sensorimotor coupling with the environment. They even suggest that "the processes crucial for consciousness cut across brain-body-world divisions, rather than being brain-bound neural events" (Thompson & Varela 2001). Even though Thompson and Varela did not specifically develop a sensorimotor theory, this general view is very congenial to the combination of workspace and sensorimotor theories that we propose. To conclude, from a practical point of view, there do not seem to be any intrinsic problems that forbid a combination of workspace- and sensorimotor explanations.

However, while it seems that there are possibilities for combining the explanatory focus of both sets of theories, so far we have glossed over more fundamental disagreements concerning the different interpretations of experience. For example, Deheane and Naccache (2001) identify experience with neural (workspace) activity. O'Regan and Noë (2001a,b,c) reject this identification, insisting that experience rather consists in a way of exploring the environment. These are deep differences in viewpoint that are difficult to reconcile. At the level of such fundamental theoretical commitments a common ground has to be found for a successful reconciliation of the theories.

4.2. Fundamental differences sorted into three scenarios

Workspace theories tend to – but are not necessarily committed to – the identification of consciousness with neural processes. Sensorimotor theories on the other hand tend to oppose this identification. A way to highlight the fundamental differences in outlook between theorists is by casting them as ideas on the localization of the processes that constitute consciousness. We will discuss three different positions on this issue, namely internal localization, external localization and no localization. We

do not aim to take position here on this fundamental issue of the constitution of consciousness. Instead, we will argue that from all three positions there are systematic benefits to be had by combining both sets of theories, even when the ways in which these unions take form can be very different under these three interpretations.

The first option for a combination comes under the assumption of *internal localization*, the standard background assumption of many workspace theorists. Experience is here conceptualized as a neural process, or more specifically workspace activity. Under this assumption, a combination with sensorimotor theory would involve the explicit articulation of the systematic links between workspace activity and sensorimotor contingencies. The latter can be interpreted in terms of their direct and indirect impact on workspace activity, the place where experience comes about in this interpretation. If sensorimotor theory correctly identifies differences in the character of experiences, the processes that the internalist claims to constitute consciousness had better make appropriate contact with the patterns of sensorimotor dependencies. The systematic differences in workspace activity brought about by the contingencies of different sensorimotor modalities would allow the explanation of comparative gaps in a way that goes beyond workspace theory. We could call this the *sensorimotor workspace hypothesis*, because ideas of sensorimotor theory are assimilated in a workspace framework, which is improved upon by this combination.

This option should be congenial for those who think that consciousness must ultimately be a brain-process. Others will see important drawbacks: even if one accepts that the experiential relevance of sensorimotor contingencies operates through their influence on neural processes, this does not make these contingencies external to experience. Sensorimotor theorists Hurley and Noë say:

> "Qualitative character may supervene on neural properties even if the qualitative expression of neural activity is determined, as we have argued, by dynamical sensorimotor considerations. (…) But if both claims are true, we hold that our account is explanatory in a way that the neural supervenience claim is not." (Hurley & Noë 2003, p. 161)

One way to elaborate this point in a slightly stronger way is by drawing an analogy between *flying* a plane and *being in* a flight simulator. Even when the pilot is not aware which condition she is in, it is only actually flying a plane that makes her fly a plane. Being in a flight simulator does not. The point here is that similarly conscious experience involves doing things and cannot be dissociated from such doings without changing the phenomenon. Even though one may preserve certain aspects of experience in a

dissociated brain and body that maintains the local representatives of normal sensorimotor contingencies – as in the flight simulator – the result is not the same natural phenomenon that one initially set out to explain. Being an experiencing individual *includes* interaction with the world.

For those who see fundamental problems with inner localization there are two alternative scenarios. One of these is the option of *external localization*. External localization focuses on the realizers of conscious experience, claiming that these involve both internal and external processes. In this interpretation, consciousness is located partly outside the head in the sense that, next to brain processes, the reciprocal sensorimotor links make processes in the environment co-constitutive of conscious experience. External localization has been defended for cognitive processes by Clark and Chalmers (1998), Clark (2008), and Keijzer and Schouten (2007). Although many find external localization highly counterintuitive for consciousness, it is explored in a positive spirit by Rowlands (2003) and Hurley (1998; 2010).[6]

An external localization scenario sets the contingencies of sensorimotor theory center stage. Within this scenario, a combination with workspace theory would be highly beneficial. The whole set of ongoing dynamical sensorimotor interaction loops, including workspace activity, could together constitute the experiential state. While sensorimotor contingencies would enable the explanations of differences in phenomenal quality, workspace theory would help to explain which interaction loops are constitutive of ongoing experience.

Both internal and external localization are subject to criticism. It is regularly argued that the criteria for applying the concept 'experience' are absent at subpersonal levels. In this view, it is a fundamental mistake to speak of physical processes as being experiences (Putnam 1999; Bennett & Hacker 2003). Of course, there are localized processes that are necessary for experience, but it is denied that any of these processes themselves constitute consciousness – at the subpersonal level there is no place 'where consciousness happens' (Dennett 1991). Even if physical processes are

[6] Rowlands (2003) defends the view that consciousness is partly externally located. He explicitly subscribes to a literal localization of consciousness in contrast with the possibility of having no localization of consciousness. For Hurley (2010), who does not discuss the issue of no localization, drawing a boundary between the processes that are 'merely' causally involved in consciousness and the processes that constitute consciousness rests on explanatory issues.

interpreted as the vehicles of the content of experience,[7] this does not imply that consciousness can be *identified* with these processes. After all, the ascription of content to vehicles may derive from the way these vehicles are functionally integrated with the activity of the organism as a whole, so that experience is conceptually tied to the activity of the organism rather than to the necessary subpersonal preconditions. This conceptual background brings us to our final, *no localization* scenario.

Sensorimotor theory provides examples of such a non-localization interpretation, e.g. where it is said that experiencing is something we skillfully do rather than any of the underlying physical processes (O'Regan & Noë 2001a,b,c; Myin & O'Regan 2002). In this view, sensorimotor contingencies are relevant to the contrasts within experience since they characterize our perceptual engagement with the environment. While the conscious/not conscious contrast does not figure within the subject's experience, we suggest that workspace theory adds to the picture by providing a subpersonal theory of the neuronal dynamics that form a precondition for experience. Workspace theory can provide a framework to understand the mechanisms by which processes become integrated to enable thought, speech and further action, thus helping to explain some of the preconditions on which a sensorimotor theory depends.

These three scenarios provide different starting points and directions for possible combinations of workspace and sensorimotor theories. In all three cases, it is beneficial to combine both workspace and sensorimotor theories. Thus without making prior commitments to any of the three scenarios, we can hold that workspace and sensorimotor theories should be combined as it will lead to an increase in explanatory potential compared to both sets of theories separately.

5. Conclusion

We have argued that sensorimotor- and workspace theories of conscious experience can be fruitfully combined. First, they are no rivals since they have different domains of application. Sensorimotor theories can best be cast as addressing the specific quality of experiences. Workspace theories on the other hand are best seen as addressing the differences between those processes that are- and those that are not directly correlated to experience. Second, even though proponents of both sets of theories work

[7] Note that such an interpretation does not imply internalism, as some have argued that vehicles of content need not be limited to processes within the head (Hurley 1998; Rowlands 2003; Noë 2004; Thompson 2007).

from fundamentally different presuppositions concerning experience, it is possible to formulate different scenarios under which a combination can take place, reflecting the different background assumptions: consciousness can be interpreted as internally localized, as partly externally localized, or as not being localized at all. In all three scenarios, advantages are to be expected from the combination of the theories.

It should be obvious that these three scenarios have different implications for the separate theories, some of which will not be acceptable for current defenders of workspace- or sensorimotor theories. In the inner localization scenario, a workspace forms the dominant basis of consciousness, while sensorimotor dependencies become relevant as factors modulating this neural activity. This relatively modest role is certainly not what sensorimotor theorists have in mind. In the case of external localization, some workspace theorists may be dissatisfied with the externalism which reduces the role of the brain to being merely part of the relevant interaction loops. Nevertheless, this scenario has the potential to be developed further such that it integrates neural, bodily and environmental processes in a way that is common practice in embodied cognition. From a sensorimotor perspective it could be questioned to what extent this is about experience, rather than merely the subpersonal preconditions of experience. The scenario that rejects localization may appeal to those who are strongly committed to the sensorimotor theory as an account of experience itself. In this scenario, workspace theory will remain necessary to explain features that are descriptive of experience. In particular it could account for differences between consciousness-correlated and not consciousness-correlated processes within conscious organisms.

The three scenarios thus have different implications for the way and extent in which the sets of theories are to be integrated. In the first two cases, integration will be an important issue, requiring both sets of theories to become adapted to one another. In the no localization scenario, this need may be felt less strongly and both sets of theories can remain comparatively independent from one another. Combining workspace and sensorimotor theories is thus not a unitary affair but an enterprise that may enfold in very different ways, depending on the scenario chosen. Eventually it will become an issue which of the possible scenarios provides the most grip on the absolute and comparative gaps of consciousness.

Finally, an important issue remains how much of conscious experience will be covered if the combination succeeds, irrespective of the way the combination eventually takes shape. The sensorimotor theories we discussed focus mainly on perceptual experience and it remains to be seen to what extent for example the experiential aspects of emotions can be

integrated in the approach. Workspace theories are typically developed based on *reports* of experience, leading to the question whether they are necessarily limited to the domain of reportable experiences. Other criteria besides reportability may be added here, such as availability for action and attention. Still, workspace theory remains strongly oriented on high-level human experience. It remains to be seen whether either workspace- or sensorimotor theory or their combination can be used to develop ideas on other forms of experience. While it remains an open question to what extent the range of application of the theories can be extended, we suggest that the procedure of *combining* these different kinds of theories will prove beneficial.

Chapter 4
Perception from the phenomenal stance

A shortened version will appear in *Logique et Analyse*.

The explanation of experience must link to the workings of our perceptual systems. To take up the challenge of the phenomenal character of experience, we should not commit ourselves to the existence of 'something else as well', over and above the processes that allow us to have access to environmental features. But how does phenomenal experience relate to perceptual access to the environment? Drawing on the notion of a phenomenal stance, it is argued that the phenomenal character of experience does not reduce to what the experience is an experience of, i.e. to the intentional content of experience.

1. Introduction

The distance from the display of my computer to the window is less than a meter. I can see this, I can judge that it is true, but it seems to me that there is also a sense in which I can 'feel' the distance. It is not merely that I have discriminatory abilities: it is that experiencing the distance has a specific phenomenal character. There is 'something it is like' to experience the spatial relations between the two objects, and this differs in hard-to-express ways from what it is like to see colors or to hear sounds. Perception guides our behavior, and through perceptual experience we acquire knowledge of our environment. But perceptual experience also comes with a specific phenomenal 'feel'. I will argue that doing justice to the phenomenal character of perceptual experience requires a distinctive approach towards perceivers, that I will call the *phenomenal stance*.[1] I will

[1] I had been thinking about the phenomenal stance for some time when Erik Myin brought a paper by Robbins and Jack (2006) to my attention. The idea of the phenomenal stance discussed here is essentially the same as theirs, but while they focus on the adoption of the stance, my main concern is the phenomenon at which the stance is aimed, as explained in Section 3.

explicate this approach by drawing on the sensorimotor account of perceptual experience (e.g. O'Regan & Noë 2001).

At present, we can distinguish two broad and well-known perspectives on consciousness. Borrowing from Joseph Levine (1994), we could say that some theorists are 'qualophilic' while others are 'qualophobic'. Qualophiles argue that standard theories of cognitive functioning fail to capture the qualitative character of experience, insisting that there is more to explain than cognitive functions or discriminatory abilities (e.g. Levine 1994; Block 1996; Chalmers 1996). Some conclude that there must exist something in the world over and above the processes underlying our cognitive functioning, proposing natural laws connecting information with experience (Chalmers 1996), or neural processes that are not involved in cognitive access (Block 2007). Qualophobes such as Daniel Dennett point out that these 'extra ingredients' could be different without any difference in perceptual judgments, arguing that it is unclear what reason we could have to posit experiential differences that make no such difference. On this view, no extra ingredient is required to ensure that the functional processes of cognitive access are accompanied by phenomenal 'feel' (e.g. Dennett 1991; Dretske 1995; Tye 1995).

Of course the qualophobe does not deny that we have perceptual abilities. One way or another we have access to our environment. Disagreement concerns whether there is more to phenomenal experience than that. Is there more to perceptual experience than access to *intentional content*, to the (apparent) object of experience? Qualophobic forms of *intentionalism* claim that that there is not (e.g. Byrne 2001; Dretske 1995; Tye 1995). But I will argue that this leaves part of experience unaddressed.

In this chapter I aim for a third road between the extremes of qualophilia and qualophobia. I share the qualophilic view that the phenomenal character of experience calls for an explanation, and that this explanation is not given by accounts of discriminatory abilities or cognitive access. At the same time, I share the qualophobe's suspicion that a qualophilic appeal to 'extra ingredients' will not be very helpful.

The tool to develop a third road will be a stances-approach: I will claim that to focus on cognitive access is to take a particular stance towards a perceiver, and that to focus on the phenomenal character of perceptual experience is to take a very different stance. While different stances towards perceivers can single out different aspects of their perceptual experience, this does not imply that there must be additional underlying processes involved, or different fundamental laws or properties. An interest in phenomenal character differs from an interest in behavior and cognitive access, but from this we should not conclude that what enables the

phenomenal character of experience involves anything over and above the processes enabling behavior and cognitive access.

The roots of the stances-approach are of course in Daniel Dennett's idea of the *intentional stance* (Dennett 1987). Taking the intentional stance towards a system (e.g. towards a perceiver) involves interpreting this system as having intentional states such as beliefs and desires. The idea is that only by taking the appropriate stance can one identify certain objective patterns in human behavior: intentional behavior and mental states simply do not figure in detached descriptions in terms of physical processes (see also MacKay 1962a).

As this chapter is focused on *perceptual* experience, I will focus on perceptual states, a subset of intentional states. Since perceptual states can be judged as true or false in virtue of their content, we could say that with this focus, the intentional stance is directed at the 'perceptual beliefs' or (potential) 'perceptual knowledge' of the perceiver. When applied to perception, the intentional stance takes an interest in the intentional content of a system's perceptual experience, in what its experience is an experience of.

When we adopt the *phenomenal stance*, we treat perceivers as subjects having phenomenal experience; from the phenomenal stance we take an interest in what their experience is like (Robbins & Jack 2006).[2] I will aim to characterize perceptual experience from this perspective by drawing on the sensorimotor account, as proposed by Kevin O'Regan and Alva Noë (2001).

My focus on the perceiver's phenomenal experience should be sufficient to avoid a possible misreading of the idea of a stance: the concept of a phenomenal stance should not be taken to imply that the phenomenon of interest is merely 'in the eye of the beholder', as if phenomenal experience does not really exist. The idea of a stance is not opposed to the reality of phenomenal experience; rather the stances-approach should be seen in contrast to the above-mentioned appeal to 'extra ingredients'.

This chapter is structured as follows. First I will regard perceptual experience from the intentional stance, focusing on intentional content (Section 2). Then I introduce the phenomenal stance (Section 3), which will be fleshed out for perceptual experience by means of the sensorimotor

[2] In this chapter, the phrase 'what an experience is like' is used as synonymous to 'what the phenomenal character of an experience is'. I bypass some of the worries from Nagel's famous 'what it is like to be a bat?' (Nagel 1974), for I will not require that an account of the phenomenal character of experience produces this very experience in the one who understands the account.

account (Section 4). This account is contrasted with the intentional stance account (Section 5), leading to the conclusion (Section 6) that the concept of a phenomenal stance allows us to do justice to the phenomenal character of experience in a way that opens a third road between qualophobia and qualophilia.

2. Perception from the intentional stance

What is the intentional content of our (veridical or non-veridical) perceptual experience? In this section we will regard conscious human perceivers from the intentional stance as familiar from Daniel Dennett (1987). The intentional stance takes an interest in such mental phenomena as beliefs, desires, preferences, and aims, which simply do not figure in detached physical descriptions of the world. We will single out a subset of these intentional states: by restricting the intentional stance to perceptual experience, we will focus on the content of experience. What we take ourselves or others to perceive can be the case, or it can fail to be the case. The intentional stance, when aimed at perception, therefore takes an interest in epistemically evaluable content, in 'perceptual belief' or (potential) 'perceptual knowledge'.

Note that the issue here is not by means of which underlying mechanisms we come to have the perceptual knowledge that we have. Neither do we need to get into the controversy concerning direct perception versus indirect perception: my discussion will be neutral with respect to these issues. The question is purely a matter of the perceptual content that can reasonably be ascribed to a perceiver. Let us consider such an issue of content to see what can be discerned when we adopt the intentional stance.

Say you are looking at two similarly-sized trees standing at different distances from you. You will often be able to see that the trees are of similar sizes, and of course that they are standing at different distances. But there is more to your visual experience than that. As Christopher Peacocke notes, "the nearer tree occupies more of your visual field than the more distant tree", and "this is as much a feature of your experience itself as is its representing the trees as being the same height" (Peacocke 1983, p. 12). I think this is quite right. In fact, equal sizes at different distances could in principle be represented in experience *without* the representation of a

difference in apparent size.[3] But in normal visual experience, often the more distant tree *will* appear smaller. A theory of perceptual experience should acknowledge this.

Now, according to Peacocke, "no veridical experience can represent one tree as larger than another and also as the same size as the other" (Peacocke 1983, p. 12). From this he concludes that one of the above-mentioned properties of experience must be a non-representational property. He proposes to distinguish 'sensational' properties of experience from their representational properties. Representational properties of experience are the way the experience represents the world to be; they are the intentional content of experience. Sensational properties are "properties an experience has in virtue of some aspect – other than its representational content – of what it is like to have that experience" (Peacocke 1983, p. 5). While the sizes of the trees are represented as being similar, Peacocke argues that the sensational properties of the experience of the trees differ. The *sizes* of the trees are perceived as similar, while their 'sensational' *apparent sizes* are at variance.

To generalize the issue, we could distinguish *nonperspectival properties* from *perspectival properties* (e.g. Noë 2004). Nonperspectival properties are properties of the object of perception that do not depend on the perceiver's perspectival positioning, such as the size of a tree, the roundness of a coin, or the glossiness of a metallic surface. Perspectival properties are properties that do depend on the perspective of the perceiver, such as the size a tree occupies in your visual field, the elliptical appearance of a coin when viewed under an angle, or the specific place where reflections appear on a glossy surface. An account of perceptual experience should do justice both to nonperspectival properties and to perspectival properties. Are perspectival properties then, as Peacocke's discussion suggests, *non-representational* properties of perceptual experience? Answering this question is the business of the intentional stance, for it concerns what we perceive.

Peacocke's reason to deny the status of intentional content to perspectival properties was the apparent conflict that would otherwise result within perception: how can the intentional content of our experience contain the sameness as well as the difference of size of the two trees? How can we perceive a coin as round but at the same time as elliptical? However, while these combinations of intentional content may seem problematic, it

[3] The following sentence presents a linguistic analogue of this possibility, in which no difference in apparent size is mentioned: 'There is a ten meter tall tree at a hundred meter distance, and a ten meter tall tree at two hundred meter distance.'

has been argued that the conflict is only apparent. For as Noë (2002) points out, the property of occupying a certain extent of your visual field is a different property from the property of having certain physical dimensions. As a result, Peacocke's problem of incompatible veridical content does not seem to arise: there is nothing contradictory in the idea that an experience represents two objects as being the same size, while also representing that they have a different perspectival size from here.[4]

Now it is generally accepted that we can perceive nonperspectival properties. For example, we can often see that a coin is round: that is how, in experiencing, we take the world to be. Such nonperspectival properties can clearly figure in the intentional content of perceptual experience. But there is no elliptical object to be perceived when you look at a round coin, and you do not experience the coin as actually being elliptical. Can we nevertheless regard the perspectival appearance of a coin as part of the intentional content of perceptual experience?

What speaks in favor of this is the fact that there is a clear objective side to how things appear from a certain perspective. Where reflections appear on an object, how much of a scene is occluded by an object, or how much noise can be heard, these are all objective properties of the world as perceived from the bodily position of a perceiver. In this sense perspectival properties can be thought of as perfectly objective environmental properties (e.g. Harman 1990; Noë 2002; 2004). They relate to the properties of objects by precise mathematical laws, such as the laws of linear perspective (Noë 2004). From this perspective, it is an objective property of the stimulus that the 'perspectival shape' of a coin is elliptical when the coin is viewed from an angle, or that the perspectival view of a flat object reduces to a line when it is viewed from aside. This objectivity may allow us to think of these properties as objects of perception.

---

[4] Irvin Rock (1977) speaks of 'proximal mode experiences' to refer to experiences of perspectival properties, arguing that: "Proximal mode experiences are better thought of as perceptions rather than as sensations" (Rock 1977, p. 349). His proposal avoids incompatible content by suggesting that we perceive perspectival size based on visual angle. Peacocke objects to this proposal suggesting that representational content depends on the possession of concepts (Peacocke 1983, p. 19-20). He argues that the experience of someone who does not possess the concept of visual angle cannot represent this property. Against this, it can be argued that the possession of concepts is not required for intentional content (e.g. Jacob & Jeannerod 2003). Alternatively, one may suppose that the relevant concepts are in fact in place when we perceive perspectival properties, even if these concepts may not normally include the concept of visual angle.

A good way to assess whether perspectival properties figure in the intentional content of perception is by asking whether we can be wrong about them. For in as far as it is possible to be mistaken about these properties, or to be right about them, it makes sense to speak of illusory or veridical experience of these properties.

An example can show how we can be mistaken about perspectival properties. Consider the shiny spots that can be experienced on glossy surfaces, known as specular highlights or specular reflections. The place where these reflections appear on objects depends on the perspective of the perceiver, which makes them perspectival properties (Madary 2008). At the subpersonal level, the highlights function as 'evidence' for the nonperspectival glossiness of objects, as well as for the object's shape. But it seems clear that specular highlights can also be perceived themselves. Indeed, the possibility of misperception of specular highlights seems assured by the following three facts: first, specular highlights can evidently be experienced, second, it is a perfectly objective feature of the environment where the specular highlights appear on a surface when looked at from a certain perspective, and third, our visual systems are fallible. This suggests that specular highlights are as (mis)perceivable as nonperspectival properties such as gloss or shape.

It seems plausible that the possibility for misperception applies to other perspectival properties as well. If a perceiver were to judge that the farthest tree of Peacocke's two same-sized trees occupies the largest part of his visual field, he would be wrong. Such a mistake may not typically occur, but the point is that this type of mistaken judgment seems possible, given that the perspectival size of a tree is an objective feature of the world as encountered from a certain perspective. The point may be generalized, on the basis of the idea that the objectivity of perspectival properties ensures the possibility for misperception.

But should we indeed attribute this possibility for misperception to the perspectival properties? Against this interpretation, one might suggest that the mistake would be 'parasitic' on a mistaken perception of the nonperspectival properties. Suppose for example that the apparent misperception of the location of a specular highlight boils down to the erroneous experience of the (nonperspectival) shape of the object. Or that a mistaken judgment concerning the perspectival sizes of two trees is in fact based in a misperception of the actual sizes of the trees. The mistaken judgment on the highlights or perspectival sizes would then derive from the mistaken perception of shape or size, and this would weaken the case for placing these perspectival properties at the side of the intentional content of experience. The question, then, is whether perspectival properties provide a *distinctive* possibility for making mistakes.

As Susanna Schellenberg (2007) has argued, building on Noë (2004) and others, the perception of nonperspectival properties depends on the sensitivity to perspectival properties.[5] Interestingly, she provides an explicitly *epistemic* interpretation of this dependency: given the idea that perspectival properties are epistemically primary, she aims to characterize the requirements for the perception of nonperspectival properties. Given this notion of epistemic primacy, the question becomes whether there is *already* room for misperception (and thereby for genuine perception) at the basic level of perspectival properties.

According to Schellenberg, *perceiving* nonperspectival properties epistemically depends on *perceiving* the perspectival properties. To make sense of this double use of the term 'perceiving', let us suppose that at least in the second instance of 'perceiving', conscious experience need not be involved. On this view, the (conscious) perception of nonperspectival properties depends on the (subpersonal) contact with perspectival properties. In this sense, of course, also the conscious experience of perspectival properties depends on perceiving the perspectival properties.

At the level of conscious experience, neither the perception of perspectival properties nor the perception of nonperspectival properties has to be epistemically foundational of the other. After all, it is perfectly conceivable that some perceivers would be more certain about nonperspectival properties than about perspectival properties, and the nonperspectival properties could even figure as evidence justifying claims on perspectival properties. The experience of the roundness of a coin, together with the experience of the tilt of the coin, could lead to the conclusion that the coin must appear elliptical from here. The reverse is possible as well: a perceiver could conclude that a coin is round, based on conscious awareness of the perspectival shape of the coin as viewed under a certain angle. The conscious experience of perspectival properties and of nonperspectival properties may in fact be based on overlapping underlying processes, in a way that precludes an analysis according to which either of the experiences forms the basis for the other. On the level of conscious experience, neither perspectival properties nor nonperspectival properties need to be basic.

The reliability of the conscious experience of these properties is of course an empirical matter. It cannot be decided a priori whether, at the level of conscious experience, a perceiver can be more certain of the

[5] Schellenberg (2007) speaks of 'intrinsic properties' and 'situation-dependent properties' to refer to what I call nonperspectival and perspectival properties respectively.

perception of perspectival properties or of the perception of nonperspectival properties. With training the reliability may shift, as when someone gets more sensitive to perspectival properties due to experience with painting. If this is right, conscious perception of perspectival properties is not necessarily the starting-point for personal-level epistemic practices. But neither are non-perspectival properties.

These considerations suggest that, at the level of conscious experience, perspectival properties may be robustly on the same side as nonperspectival properties, that is, they may be on the side of intentional content. What Peacocke classified as 'sensational properties' may then figure in the intentional characterization of perceptual experience. The intentional stance, when restricted to perceptual experience, can reveal both perspectival properties and nonperspectival properties as potential objects of perceptual experience.

Given that the intentional stance can do justice to perspectival properties as well as to nonperspectival properties, it yields a relatively rich notion of perceptual content. This richness may raise the impression that nothing is missing from the intentional stance account. But let us now approach perception from a different entry point, and see where that might lead us.

3. The phenomenal stance

The question in this section is no longer what experience would tell the perceiver about his environment if the experience were veridical. Rather, we take the *phenomenal stance* towards a system, which is to regard it as a locus of phenomenal experience, and to take an interest in the character of this experience (Robbins & Jack 2006). Here I will introduce this stance, and I will highlight a difference between the resulting view on experience and the account in terms of intentional content. In Section 4 and 5, we will see whether we can identify aspects of experience that become salient from the phenomenal stance, which are overlooked by the intentional stance.

Philip Robbins and Anthony Jack (2006) argue that different cognitive capacities are involved in our thinking about physical processes as compared to our thinking about intentionality and phenomenal experience. Regarding a system as a purely physical mechanism is one thing, regarding others as intentional systems is another, and regarding them as phenomenal beings is to take yet another stance. The context in which their discussion is framed is the psychology of the observer: if these different stances are competing psychological processes, Robbins and Jack argue, this may help to explain why many people find it difficult to understand how a system can be capable of conscious experience, while it is at the same

time a physical system. At least part of the intuition that there is an explanatory gap between physical processes and conscious experience, they suggest, "results from more or less hard-wired restrictions on information flow across competing neural networks" (Robbins & Jack 2006, p. 78). The idea is that the difficulty to fit physical processes and phenomenal experiences together in one's mind exceeds the difficulty of combining a merely physical interpretation and an intentional interpretation of a system. They discuss evidence indicating that regarding others as having phenomenal experience reflects in fact a different capacity than adopting the intentional stance.

Dennett already pointed out that there is more to normal human interaction than can be captured within the intentional stance. A reason for this is that the intentional stance lacks a moral dimension. We can usefully regard robots and chess computers from the intentional stance, but even if we do so, "one is guilty of no monstrosities if one dismembers the computer with whom one plays chess, or even the robot with whom one has long conversations" (Dennett 1978, p. 240). When we treat other human beings as the object of moral concern, Dennett proposed, we take a moral or personal stance towards them (Dennett 1978). But when do we adopt this moral stance? Robbins and Jack argue that this has its basis in our everyday psychological capacity to treat others as possessing experiential states (Robbins & Jack 2006). For the question then is not so much what beliefs or other propositional attitudes others can have, as in the intentional stance, but rather whether for example they can suffer.[6]

When we encounter someone who is suffering, we naturally take the phenomenal stance towards this person. But perhaps adopting the phenomenal stance towards someone's *perceptual* experience is less common. For example, when we encounter someone who is looking at a colorful scene, we typically will not be concerned with the way the colors appear to that person. Given that we are often more interested in what the world is like than in what the experience of the world is like, taking the phenomenal stance towards others *as perceivers* may not be our natural attitude. But common or not, we certainly can take such a stance. For example, we may at times wonder what a certain scene is like for someone with color blindness or someone with macular degeneration. In fact, we can take such a stance not only towards others, but also towards ourselves: this is exactly what I do when I reflect on my experience of the distance between the display of my computer and the window.

---

[6] Cf. Bentham (1823, second edn.), footnote to Chapter 17.

Taking a stance is always taking a stance towards someone or something. While Robbins and Jack focus on the side of *adopting* the phenomenal stance, I shall be mainly concerned here with the phenomenon at which the stance is aimed. Just as the physical stance and the intentional stance single out different patterns in the world (Dennett 1991b), so the intentional stance and the phenomenal stance may pick out different phenomena as well. When applied to perceptual experience, the phenomenal stance is aimed at the phenomenal character of the experience, and not necessarily at intentional content.

I suggest that the phenomena towards which the stances are directed do in fact differ. For while the intentional stance analyzes perceptual experience in terms of the objects of experience, the phenomenal stance is aimed at the way the perceiver is affected by these objects. In this sense, phenomenal character is a perceiver-centered notion, which contrasts with object-centered notions such as perceptual knowledge, which figure in the intentional stance. I take this difference in focus to be a key difference between the phenomenal stance and the intentional stance.

It should be emphasized that adopting the phenomenal stance towards someone need not produce the target experience for the one who adopts this stance: one need not suffer in order to recognize that others are in pain, although one may need to have experienced the feeling before.[7] A characterization from the phenomenal stance of someone's perceptual experience should provide an intuitive or theoretical grasp of the phenomenal character of someone's experience, but it need not put us in the other's first-person perspective. One should not suppose that adopting the phenomenal stance towards a bat allows us to experience what it is like to be a bat in the same way as the bat itself does.

Various authors have provided accounts of phenomenal experience. A recent example is Evan Thompson's (2007) 'enactive' approach to the mind, which draws on the phenomenological tradition. From a different starting-point, Thomas Metzinger (2003) provides a representational analysis of the structure of subjective conscious experience. Within the domain of *perceptual* experience, however, the sensorimotor account is perhaps the most detailed and promising proposal, and I shall turn to it in the following section. On the basis of this proposal I aim to flesh out the phenomenal stance for perceptual experience.

---

[7] Robbins and Jack (2006) note that full-blooded empathy often involves sharing the phenomenal feel of the one we engage with. This means that our common empathic engagement goes beyond the core of the phenomenal stance.

4. Perceptual phenomenology: a sensorimotor characterization

When we adopt the phenomenal stance towards others, we see them as subjects of phenomenal experience. But how can we formulate more precisely what we can learn about others when we adopt this stance towards them? What do we focus on when we take this stance towards ourselves? The *sensorimotor account*, as proposed by Kevin O'Regan and Alva Noë, aims to account for the phenomenal character of perceptual experience (O'Regan & Noë 2001). It forms a plausible candidate to get grasp on the target of the phenomenal stance.[8] Let me explain the approach by means of a few examples.

By way of introducing the sensorimotor approach to perceptual experience, consider the following. When we do not move, changes in sensory input are required for the visual perception of movement in the environment. However, if we move our eyes in such a manner that they track the movement of an object against a neutral background, it is precisely because the input from the stimulus does *not* change that we can see that something moves. We can conclude that our ability to perceive movements in the environment is based on the specific relation between sensory input and motor action, the 'sensorimotor contingencies' or 'sensorimotor dependencies'. For only the combination of the signals relating to possible self-movement and signals relating to possible retinal changes provides the information regarding environmental movement. The sensorimotor account of perceptual experience claims that perception more generally builds on such sensorimotor dependencies. For example,

[8] We saw in the previous section that adoption of the intentional stance need not involve any commitment to hypotheses about the underlying mechanisms of perception. The definition of the phenomenal stance need not contain any reference to precise underlying mechanisms either: to adopt the phenomenal stance is simply to approach others as subject of phenomenal experience, taking an interest in the phenomenal character of this experience. Still, we need some terms in which the phenomenal character can be framed. When we conceive of experience as an embodied phenomenon, we may expect that the – perceiver-centered – phenomenal stance characterization of perception in one way or another appeals to the causal processes underlying perception. From this perspective, one might suggest that for perceptual experience a phenomenal stance proposal generates a sensorimotor account in much the same way as an intentional stance proposal generates an account in terms of beliefs, desires, and the like.

also the perceptual experience of different colors, sounds, and shapes, all involve different sensorimotor dependencies (O'Regan & Noë 2001).

I suggested in the section above that an account of the phenomenal character of experience will be an account in perceiver-centered terms. The sensorimotor approach offers such an account. Indeed, the key move of sensorimotor theory is to focus on the characteristic patterns of sensorimotor engagement that enable perceptual experience, rather than on the object of perception such as a perceived environmental movement. It does so not merely with the aim of addressing the causal preconditions of perceptual experience. By appealing to the dynamic patterns of sensorimotor engagement with the environment, the theory explicitly intends to characterize and explain the phenomenal character of experience (O'Regan & Noë 2001; Hurley & Noë 2003; O'Regan, Myin & Noë 2005).

Now suppose that you look at an object at some distance from a background, say the display of your computer. You will be able to see that it is located at a certain distance from the wall behind it, for example. You can estimate this distance, perhaps indicating it with your thumb and index finger, in which case your estimate could be quantified as approximately a certain number of centimeters. So far, your experience is treated from the object-centered intentional stance, in terms of the information provided by perception. But intuitively it seems that there is more to the experience of distance than the mere possession of this information. What does the phenomenal 'feel' of this distance consist in? How can we characterize this experience from the phenomenal stance?

According to the sensorimotor account, the experience of the distance between an object and its background derives from the specific way in which this distance is encountered. New information becomes available to the perceiver by the way in which sensory input changes as he or she moves, depending on the distances between object, perceiver, and background. The phenomenal 'feel' of a distance, which seems insufficiently addressed by merely mentioning the distance or the possession of information, may be better addressed in terms of the characteristic sensorimotor dependencies implied by the spatial situation. For the phenomenal character of experience, it is claimed, the patterns of sensorimotor engagement are what matter.[9]

[9] This appeal to sensorimotor dependencies does not imply that one has to move in order to see. For example, binocular vision enables the experience of the distance between figure and background, presumably without involving movement, but it does so only after binocular vision has been developed. It is because former

Reflection on the feeling of softness may help to get an intuitive grasp of the matter. We clearly experience the softness of a sponge by the exploration of it, in which sensory stimulation changes in a characteristic way as a function of motor action. The sensorimotor approach focuses on the ways in which the perceiver is engaged with the soft object, rather than exclusively addressing the epistemically evaluable upshot of these encounters. As Erik Myin writes:

> "Though softness is clearly grounded in material properties of objects, the *experience* of softness can only be understood by reflecting on how softness is apprehended. In the sensorimotor account, the experience of softness comes about through a specific pattern encountered in a sensorimotor exploration, including facts as that if one pushes on a soft object, it yields." (Myin 2003, p. 43)

The sensorimotor account similarly characterizes experiential differences between perceptual modalities (such as vision and touch) in terms of the characteristic differences in the patterns of sensorimotor dependencies (O'Regan & Noë 2001; Hurley & Noë 2003; O'Regan, Myin & Noë 2005). On its account, what it is like to see differs from what it is like to feel to the extent that the patterns of sensorimotor dependencies differ. Characterizing these patterns is characterizing phenomenal experience, conceived as the perceiver's embodied engagement with its environment.

The sensorimotor account offers a natural way to characterize the object of the phenomenal stance. We saw that its perceiver-centered analysis captures aspects of perceptual experience that do not figure in the intentional stance. This becomes even clearer when, in the next section, we relate this view to the analysis of perception in terms of intentional content.

5. Character and content

A focus on the perceiver's embodied engagement with her environment may flow naturally from the concerns of the phenomenal stance. But how exactly does this perspective relate to the analysis in terms of intentional content? I will distinguish two views of this relation. On either view, a phenomenal stance allows us to address issues of perceptual experience that are untouched by intentional characterizations, such as the characterization in terms of perspectival properties and nonperspectival

---

encounters enable the implicit grasp of the distance-specific sensorimotor contingencies that one need not move to experience the distance.

properties. I close the section by contrasting the analysis with intentionalism.

Consider the differences between visual and tactile perception, the difference between seeing and touching. How does the phenomenal stance towards these experiences differ from the way in which the intentional stance treats them? Alva Noë writes:

> "From the side of the object, what differentiates seeing and touching are their different objects (looks as opposed to feels, say). But from the side of the perceiver what differentiates seeing from touching are the different patterns of activity in which seeing and touching respectively consist." (Noë 2002, pp. 66-67)

Perceptual experiences can be identified both in terms of the *objects* of perception and in terms of the different *patterns of activity* they involve. When we take the intentional stance, we focus on the objects of perceptual experience, on what the experience is an experience *of*. The sensorimotor theory adds an analysis of experiencing in terms of the way in which the perceiver is engaged with his or her environment. Clearly this enhances the richness of our account of perceptual experience. But we may wonder whether the perceiver-centered analysis of the sensorimotor account on its own can do sufficient justice to the phenomenal 'feel' of having meaningful experiences, such as that we see a chair *as* a chair. May not the intentional content somehow be part of the phenomenal character of experience? According to Noë:

> "What determines the quality of experience (...) is two-fold. First, there is *what* you experience (the representational content). And second, there is, roughly, what happens to you while you experience." (Noë 2002, p. 67)

There are two ways to read this claim. One is to take it as suggesting that there are two factors, which together compose the phenomenal character of experience: one factor is the intentional (or 'representational') content; another factor is an *additional* bit of phenomenal character. On this reading, only part of the phenomenal character is provided by the perceiver-centered approach. The result is a hybrid conception of phenomenal character as consisting of intentional content plus something else (e.g. non-representational 'sensational' properties of experience). This would conform to Ned Block's attribution of intentional content to phenomenal character, while claiming that phenomenal character 'outruns' this intentional content (Block 1996).

The second reading is as follows. To claim that the phenomenal character of perceptual experience is partially 'determined' by what you

experience is just to say that it *depends on* what there is to perceive. This is of course trivially true. But it is not to say that phenomenal character of experience partly *consists of* intentional content. Thus Noë's quote above is fully compatible with the idea that the intentional content of experience does not figure in the phenomenal stance at all, at least not as co-constitutive part of the phenomenal character of experience.[10]

It seems to me that phenomenal character can plausibly be analyzed first and foremost as a matter of the process of encountering the environment, rather than as an amalgam of content and engagement. To account for the fact that we phenomenally experience chairs as chairs, we may appeal to our former encounters with chairs, which shape our present way of engagement with our environments, providing it with significance. For the phenomenal character of experience, it is the mode of embodied engagement that counts, rather than the intentional object which we encounter.

For present purposes, however, the important thing is that either reading implies that the phenomenal character of experience cannot be reduced to intentional content. Consider for example the experience of *perspectival properties* and *nonperspectival properties*, discussed in Section 2. The phenomenal stance analysis can easily be applied to the experience of these properties. While what we perceive can be for example the elliptical appearance of a coin or its roundness, what it is like to perceive these properties is at least partly a matter of the perceiver's embodied engagement with the perceived object. The phenomenal character of experience lies at least in part in the specific way in which this information is acquired, rather than purely in the information itself.

Now let me contrast this analysis with intentionalism, which claims that phenomenal character supervenes on intentional content (e.g. Byrne 2001; Dretske 1995; Tye 1995). According to this supervenience claim, a creature that acquires exactly the same intentional content by means of vision as by means of touch would thereby have experiences with the same phenomenal character. An interesting feature of the sensorimotor account, in contrast, is that it allows that the phenomenal character of experience may differ even

[10] Note that a perceiver-centered analysis may involve processes outside the body of the perceiver, insofar as they have impact on the perceiver. Cf. Clark (2008, p. 139): organism-centered does not imply organism-bound. But while we may identify patterns of engagement as characteristic for interaction *with certain objects*, this provides no reason to include intentional content (that is, the intentional upshot of this encounter) within the phenomenal character of the perceptual experience.

if there is no difference in the intentional content of experience. Imagine a creature that acquires no extra information by means of touch than it would acquire by means of vision. For example it can see that there is a square of a certain size and orientation, but it could also acquire the very same knowledge by means of touch. Since the sensorimotor account appeals to the processes through which the perceiver obtains perceptual knowledge, and given that vision and touch imply different sensorimotor patterns, the sensorimotor account allows that the phenomenal character of these two experiences may differ. They may differ despite the fact that there is no difference in the object of awareness, i.e. in the perceptual knowledge the perceiver possesses.

6. Conclusion

Focusing on perceptual experience, I aimed to explicate a third way between the extremes of qualophobia and qualophilia. I share the fundamental qualophilic conviction that the phenomenal character of experience is something that needs to be accounted for. I discussed the notion of the phenomenal stance, as introduced by Robbins and Jack (2006), arguing that the phenomenal character of experience forms a different explanatory target as compared with the intentional content of perceptual experience.

To adopt the phenomenal stance towards a system is to regard it as having phenomenal experience, and to take an interest in the phenomenal character of the system's (supposed) experiences. There are two ways to approach the phenomenal stance. One is focused at the adoption of the stance. From this perspective, Robbins and Jack (2006) argue that the dualistic tendencies in our thinking may partly result from the observers' 'Balkanized' brains, in which the adoption of different stances exclude each other. I approached the phenomenal stance in a complementary way, namely by focusing instead on the phenomenal character of experience at which the stance is aimed.

I suggested that this focus requires a perceiver-centered analysis of the process of experiencing. The sensorimotor account provides such an analysis. Rather than zooming in on the epistemically evaluable upshot of perceptual encounters, it focuses on the patterns of dependency of sensory input on motor action that are causal preconditions for perception. It appeals to these patterns to provide a positive characterization of experiencing.

Although born from qualophilic worries, this phenomenal stance proposal at the same time respects one of the qualophobe's core beliefs. For the acknowledgement that phenomenal character and intentional content

are different phenomena is accompanied by shunning 'extra ingredients' in the world over and above the processes that enable behavior and perceptual access. The phenomenal stance can thus be regarded as a further development of Dennett's qualophobic thoughts on the intentional stance (1987), the personal stance (1978) and consciousness (1990b).

The kinship of the phenomenal stance proposal with Dennett's views is further underlined by considering its relation with Dennett's *heterophenomenology*. The aim of heterophenomenology is to connect experience and the natural sciences, accommodating the first-person point of view within a third-person framework (Dennett 1991; 2003). The heterophenomenologist communicates with the subject in order to get a descriptive account of the subjective experience of the person, the subject's 'heterophenomenological world', before attempting to explain this heterophenomenological world.[11] Given the commitment to taking the subject's reports seriously, without granting him or her absolute authority on the explanation of the experience, it is not surprising that O'Regan and Noë write: "It may be – indeed, it is likely – that our phenomenological analysis can be accommodated by heterophenomenology" (O'Regan & Noë 2001b, p. 1014).

The key point is that according to the phenomenal stance proposal, heterophenomenology should not be focused exclusively on intentional content. For beliefs about the phenomenal character of experience are not the same as beliefs about the intentional content of experience. This may not make it easier to apply the approach to naïve subjects, who may be naturally inclined to be directed at what the world is like rather than at what their experience is like. But this is not to say that it cannot be done. Indeed, as Dennett says, "The policy of training subjects (…) might in some circumstances heighten the powers of subjects to articulate or otherwise manifest their subjectivity to investigators" (Dennett 2003, p. 29).

While the intentional stance characterizes perceptual experience in object-terms, the phenomenal stance – as exemplified in the sensorimotor approach – provides a characterization in terms of the perceiver's embodied engagement with her environment. On this account, the phenomenal character of perceptual experience does not reduce to the

[11] As Dennett emphasizes: "Notice that when you are put in the heterophenomenologist's clutches, *you get the last word*. You get to edit, revise, and disavow *ad lib*, and so long as you avoid presumptuous *theorizing* about the causes or the metaphysical status of the items you report, whatever you insist upon is granted constitutive authority to determine what happens in your heterophenomenological world" (Dennett 1991, p. 96).

possession of intentional content. This allows the phenomenal stance to reveal aspects of experience that are left out of intentional characterizations such as characterizations in terms perspectival and nonperspectival properties. In fact this proposal leaves open the possibility that experiences may differ in phenomenal character, even if they are alike in intentional content. Importantly, the phenomenal character of experience does not require *pukkah* ingredients over and above the processes that enable behavior and perceptual access. I suggest that the rejection of these 'extra ingredients' is acceptable exactly because it turns out that they are not necessary in an account of the phenomenal character of experience.

Chapter 5
The structure of color experience and the existence of surface colors

This chapter is joint work with Erik Myin:
Degenaar, J. & Myin, E. (s*ubmitted*)

Color experience is structured. Some 'unique' colors (red, green, yellow and blue) appear as 'pure', or containing no trace of any other color. Others can be considered as a mixture of these colors, or as 'binary colors'. According to a widespread assumption, this unique/binary structure of color experience is to be explained in terms of neurophysiological structuring (e.g. by opponent processes). The argument from structure builds on this assumption to argue that colors are not properties of surfaces, and that color experiences are neural processes without environmental counterpart. We reconsider the argument, and we discuss recent models in vision science which point at environment-involving patterns that may be at the basis of the unique/binary structure of color experience. We conclude that the argument from structure should be rejected.

1. Introduction

Not all colors come equally. Cross-cultural studies of color naming showed that four 'focal' colors are widely singled out by color terms across languages (Berlin & Kay 1969; Regier *et al*. 2005): red, green, yellow and blue. There is a considerable psychophysical literature showing that normal perceivers can identify 'unique' shades of red, green, yellow and blue that appear not to be tinted by the other hues (Hurvich 1981; Valberg 2001). That is, red, green, yellow and blue lights can appear in a 'pure' or 'unique' form, that is as experienced without any trace of another hue. In contrast, 'mixed' or 'binary' hues, such as orange and indigo, appear as mixtures of two other (unique) hues. For example, magenta can be described as a slightly bluish red. In contrast, it does not seem natural to describe red as a purplish orange. Four focal colors correspond closely to the unique hues (Miyahara 2003; Kuehni 2005). Thus it seems natural to attribute a

perceptually special status to particular shades of red, green, yellow, and blue, which we will call the 'unique colors'.[1]

The unique/binary structure of color experience has figured prominently in arguments to the effect that colors are not properties of surfaces, and that color experiences are neural processes lacking an environmental basis. We will focus on a particular argument, which we will call *the argument from structure*.[2] It has two premises (concerning structure) and two conclusions (concerning colors), and the argument for each conclusion can be said to form a strand. The first premise is this:

> *Anti-realism about physical structure*. This is the idea that there is nothing analogous to the unique/binary structure to be found in the stimulus.

The idea here is that the unique/binary structure of colors has no genuine basis in the non-neural. The second premise is this:

> *Realism about neural structure*. This is the idea that the structure of color is matched by a structure in the human visual system.

From the two premises two conclusions are then drawn. The first premise is used to argue for the first conclusion:

> *Anti-realism about physical colors*. This amounts to the thesis that colors are not physical properties of the extracranial world.

From this conclusion, together with the second premise, a second conclusion is drawn:

> *Realism about neural colors or color experiences*. The idea here is that colors or color experiences, which presumably lack a physical basis in the environment, are a property (Hardin 1988/1993) or a construct (Palmer 1999; Zeki 1983) of the brain.

---

[1] For critique on the idea that there are four primitive unique hues, see Saunders and Van Brakel (1997). Our use of the term 'unique/binary structure' is only meant to capture the widely accepted (although not completely uncontroversial) special status of four focal colors. We certainly do not deny the existence of further focal colors that also have a somewhat special status, and we are not committed to the idea that binary colors are truly *composed of* the unique colors.

[2] 'Arguments from structure' can take somewhat different forms. An important feature of our analysis is that we show how ideas about the explanation of the structure of colors figure in the argument. We argue that assumptions regarding this explanation, although widely shared by proponents and opponents of the argument, are in fact mistaken.

As Hardin puts it:

> "so much about the colors – particularly the division between unique and binary hues – is manifestly bound up with the peculiar characteristics of the human visual system that it seems a vain pursuit indeed to search for a parallel set of structures in the general order of nature." (Hardin 1993, p. 67)

Building on a selection and interpretation of explanations given in the science of color vision, the argument from structure leads to the conclusions that "We are to be eliminativists with respect to color as a property of objects, but reductivists with respect to color experience" (Hardin 1993, p. 112).

In this chapter we critically examine this argument from structure, through scrutinizing both of its premises in relation to its two conclusions. In Section 2 we address the first strand of the argument, the claim that colors are not in the environment because the environmental stimulus lacks the required structure. We will address this strand as coupled to a requirement that any candidate environmental correlate of the relevant structural properties (unique versus binary) be explanatory with respect to the structure of experience. If this requirement is justified, it is not sufficient for refuting the argument from structure to point out that there is *some* possible description of the environment in terms of the unique/binary structure, for non-explanatory environmental properties are rejected. This explanatory requirement has not been met in recent defenses of color realism (Byrne & Hilbert 2003; Churchland 2007).

In Section 3 we examine the second strand in the argument from structure, in which it is concluded from the conclusion of the first strand, plus the idea that the unique/binary structure is to be explained fundamentally in terms of neural processes, that colors or color experiences inhere in the brain. We question both the evidence on which this reasoning builds, especially its interpretation of opponent processing models, as well as its logic, especially the problematic reliance on *a priori* excluding contingent explanatory possibilities. Alternative explanations have been unavailable for a long time. This may have raised the impression that both strands of the argument from structure have color science at their side.

Recent analyses, however, offer a different perspective. In Section 4, we shall discuss the possible role of the environment in the explanation of the unique/binary structure of color experience. We shall focus in particular on an analysis of the stimulus as probed through our photoreceptors, which reveals systematic differences between unique and binary colors (Philipona & O'Regan 2006). In this approach, environmental factors come to carry

much of the explanatory weight regarding the structure in color experience that is carried by neurophysiological factors in opponent channel models .

Section 5 discusses the consequences for the argument from structure as an argument for anti-realism about physical colors and neural reductionism about color experience. It shouldn't be taken for granted that key facts of vision science support the argument from structure. If the empirical approaches to which we draw attention are on the right track, the explanation of the structure of color essentially involves 'extradermal' factors. We claim that the conclusions of both strands of the argument become seriously discredited by these results. Contrary to the conclusion of the argument from structure, there very probably is both organismic and environmental reality to the structure of color.

2. The first strand: from lacking physical structure to anti-realism about physical colors

Minimally, the unique/binary structure of color experience consists in the fact that some colors have a special place in human color vision, and that these colors are natural points of reference for describing color experiences. Stronger notions of the unique/binary distinction could invoke the idea that binary colors must be composed of unique colors, and that unique colors are not composed of any other colors. Such stronger notions may not be obvious to those familiar with mixing paint, especially in the case of green which after all lies between yellow and blue and can be seen as a mixture of these (for those familiar with mixing light, the 'uniqueness' of yellow in this respect might not be obvious). The argument from structure as we shall construe it need not appeal to such a strong notion of uniqueness. It needs only to refer to the less controversial fact that certain shades of red, green, yellow and blue appear to be special points of reference within color vision. The resulting structure of color experience is thought to be fundamental to our concept of colors (see and cf. Hardin 1993, p. 66).

According to the first premise of the argument from structure there is no physical structure in the environment analogous to the unique/binary structure in experience. From this, it is concluded that there are no environmental colors. The absence of a physical structure intuitively comes down to the following. There are reflectances on the environmental side, and experiences on the experiential side. The issue is not that some (sets of) reflectances reliably cause types of experiences. Rather, the relevant point for the defender of the argument from structure is that the structure discernible in experience is not genuinely present in the environmental properties. In particular, there is no properly physical property of the

particular reflectance patterns for unique colors, that distinguishes them from the reflectance patterns for binary colors.

According to the proponent of the argument from structure, no scrutinizing from a purely physical angle will deliver the required difference. Based on the idea that the unique/binary structure is fundamental to what colors are, an inference is drawn from the absence of appropriate environmental structure to the claim that there are no environmental colors. As Hardin puts it:

> "If hues are physical complexes, those physical complexes must admit of a division into unique and binary complexes. No matter how gerrymandered the physical complex that is to be identical with the hues, it must have this fourfold structure, and (…) it must be possible to characterize that structure on the basis of physical predicates alone." (Hardin 1993, p. 66)

We think the argument from structure should be understood in an explanatory context. That is, we propose that the question of there being a 'genuine basis' or a 'match' for the structure of experience in the environmental realm is to be decided by whether or not reference to independent structure in the environmental realm carries explanatory weight regarding the structure of experience. 'Independent' here means that the environmental properties in the explanation are not defined in terms of experience itself, or in terms of higher level, in particular neural, properties. This is not to say that a genuine environmental structural basis cannot refer to any neural properties in its explanation of the structure of experience. It can do so, as long as the neural properties do not themselves explain the unique/binary structure of color experience – in Section 4 we shall describe a model in which the environment plays such a role. Interestingly, the requirement of explanatory independence can be applied both to the physical (or environmental) as to the neural level. For a neural explanation of the structure of color experience is only independent of the environment if it stands apart from such structural properties as might be found in the environment. That is, if the neural structure itself can be explained in terms of the environmental structure, it loses its independent status.

There are several reasons for interpreting the argument from structure in this explanation based way. First, it provides a version of the argument that is harder to refute. Under the proposed interpretation, the first strand of the argument from structure, leading to the denial of the physicalist thesis that colors exist 'out there', cannot be simply refuted by pointing at some non-explanatory, dependent physical properties. At the same time, the explanatory requirement ensures that the second strand of the

argument from structure yields a substantial claim, because the explanatory requirement applies to the neural level as well. Thus, by an explanatory reading, the argument against physical colors becomes more robust, and the assertion of the neural reality of color experiences becomes substantial. Second, this taking of the argument from structure is responsive to considerations given in favor of a 'naturalization' of ontology (Ladyman & Ross 2007; Hurley 2010, and of course Quine 1960). For a question about existence of a genuine basis or match is decided by reference to what plays a role in explanations. Third, the proposed interpretation seems to accommodate existing treatments of the argument from structure. For example, when Hardin writes that, in order to establish that hues are 'physical complexes', "it must be possible to characterize that structure on the basis of physical predicates alone" (Hardin 1993, p. 66), this seems to fit very well with our requirement of independence for physical predicates.

The proposed interpretation of the argument from structure implies that two physicalist responses to the first strand of the argument from structure no longer can be mounted against it. To reiterate, in this strand, it is concluded that there are no physical colors, based on the premise that there is no physical property in the environment that corresponds to, or explains, the structure in experience. A first physicalist response is to concede that the unique/binary structure may be fundamental to our experience of colors, but to deny that this structure is fundamental to colors themselves. This opens up the possibility to accept the idea that the unique/binary structure lacks an environmental reality without accepting anti-realism regarding environmental colors (e.g. Churchland 2007). If the environment does not contain the unique/binary structure, one might prefer to downplay the importance of this structure rather than the environmental reality of colors. This type of color physicalism is being put aside by the assumption that physical properties, in order to count as colors, must have a unique/binary structure which contributes to the explanation of the structure of experience.

A second physicalist way of answering is to accept that colors are structured into unique and binary hues, but to claim that a dependent physical structure is sufficient for the reality of surface colors. Byrne and Hilbert (2003), following this path, define the property of "hue quantity" associated with a certain reflectance profile (a related proposal can be found in Bradley & Tye 2001). A hue quantity of a reflectance profile corresponds to how much of one of the hues red, green, blue and yellow is estimated to be present in the profile, when seen by some observers. Surfaces that are colored unique red are then the set of surfaces with

reflectance profiles that have a maximal value on the red hue quantity dimension, and a minimal value on the other hue quantity dimensions.

The identification of hue quantities clearly derives from properties of observers. Hilbert and Byrne make clear that they think the relevant observer properties can be neural properties, related to the organization of the neurophysiology of color in terms of opponent channels (to which we shall turn in Section 3.1). Thus they grant that the characterization of having a certain hue quantity may have recourse to exactly those high level neural properties of observers which are assumed to directly explain the structure of color experience in a neural model. They think that this does not imply that the property of being unique red itself is not physical. But such newfound object properties are not independently characterized, and therefore they remain explanatorily impotent. In Byrne and Hilbert's hue quantities proposal it would be wrong to claim that the instantiation of hue quantities explains why human color experience has a unique/binary structure. That would turn matters on their head, as it is rather the case that the properties of experience explain the instantiation of hue quantities by surfaces.

The 'hue quantity' approach does grant that the explanatory burden for the binary structure of experience can be carried by neural structures, in particular opponent neurophysiology. As will become clear in the next section, we think that there are both logical and empirical reasons against relying on opponent neurophysiology to explain the unique/binary structure of experience. It is a live possibility that the environment *does* play a stronger explanatory role then is presumed by both anti-realists about physical color, like Hardin, and realists about physical colors, like Churchland or Byrne and Hilbert. If an independent structure could be established, of course, this would directly refute the first strand in the argument from structure and open up the prospects of a realism about physical colors that does live up to the standards flowing from the explanatory form of the argument from structure. We will explore this niche of possibilities in Section 4.

3. The second strand: from neural structure to realism about neural colors

The second strand of the argument from structure takes two premises. The first one is the conclusion of the first strand, namely that environmental colors do not exist. From this, plus the additional premise that the unique/binary structure of colors is explained by the structure of neurophysiological processes, it is concluded that color experiences are neurophysiological properties without a physical basis in the environment.

Is the hypothesis forming the second premise of this strand justified? Can the unique/binary structure of color experience be explained by the peculiarities of our neurophysiology? Philosophers writing on the subject have expressed widespread confidence that *opponent processes* in the visual system can explain the structure of color experience (Hardin 1988/1993; Clark 1993; Churchland 2005; 2007). However, as we point out below, the evidence base for this claim is less robust than these philosophers – unlike many color scientists – assume.

3.1. The opponent processes hypothesis

It is clear that the difference between unique and binary colors cannot be explained by the sensitivities of our photoreceptors: while there are presumably four unique colors, there are photoreceptors of three sensitivity types. Many vision scientists have therefore searched in 'subsequent stages' in the visual system for processes that could account for the unique/binary structure of color experience. The most popular neurophysiological account appeals to *opponent processes* in the brain (Hering 1920; Hurvich 1981; Hardin 1993; Clark 1993). The idea, derived from psychophysical studies, is that the visual system is organized in pairs of antagonistic physiological processes: one process corresponding to the black/white dimension; one for red and green; and one for yellow and blue. The processing of one color in an opponent process would then exclude the processing of the other, so that when the red/green opponent channel signals red, it cannot signal green. If the experience of unique colors correlates with the extremes of such antagonistic processes, these channels would be the neural process underlying the structure of color experience. While the experience of unique colors would involve activity in one color-sensitive channel only with the other channel at equilibrium, the experience of binary colors would have to result from the integration of signals from both channels. The opponent processes could then be interpreted as *color opponent* processes, which would be the physiological correlate of the unique/binary distinction (Hardin 1988/1993; Clark 1993; Churchland 2005).

For example, in a philosophical discussion of color experience it has been suggested that the color opponent processes hypothesis generates predictions regarding negative after-images, and that these predictions are borne out (Churchland 2005). The reasoning is as follows. Visual pathways are fatigued by elongated stimulation with the same color. As a result, we become temporally less sensitive to this color. If the extremes of the opponent processes are signaling the unique colors, we may expect that a red/green channel that is fatigued with a red stimulus favors the signaling

of green when a neutral stimulus is presented. Against a neutral background, the resulting after-image of unique red may then be expected to be unique green – assuming that the fatiguing takes place in the alleged color opponent stage in the visual system. But the trouble here is that there is evidence that for normal human perceivers the after-image of red is cyan and of green it is magenta (Wilson & Brocklebank 1955; Pridmore 2008). Cyan and magenta can indeed be classified as 'green' and 'red' respectively, but surely not as unique green and unique red. So after-images do not support the hypothesis that opponent processes match with the unique/binary structure of color experience.

Psychophysical experiments have supported the idea that there are chromatically opponent channels in two 'cardinal directions' (Krauskopf *et al*. 1982), and in the extensive literature on color vision the existence of *physiological* opponent processes is widely accepted. The neural activity originating in the different photoreceptor types is recombined into opponent processes that inherit a wavelength-dependency from the photoreceptors at their basis (Valberg 2001; Gegenfurtner 2003; Wuerger *et al*. 2005). For example, neurons with opponent properties have been found in the lateral geniculate nucleus in macaques, which have similar color vision as humans (Sandell *et al*. 1979), (De Valois *et al*. 1966; Derrington *et al*. 1984). The activity of these neurons displays a wavelength-dependency combining the photoreceptor inputs into chromatically opponent processes.

Importantly, however, the extremes of the activation of these physiological opponent processes in the lateral geniculate nucleus do not generally coincide with the unique hues. For example, the yellow-blue dimension of color experience is not located at the equilibrium of any of these opponent processes (De Valois *et al*. 1966; Derrington *et al*. 1984). There is a clear mismatch between the physiological opponent processes in the lateral geniculate nucleus and the experience of the unique hues (Mollon & Jordan 1997; Valberg 2001; Wuerger *et al*. 2005).

In response to this mismatch, some vision scientists have hypothesized a further stage in the visual system where neural activity does reflect the unique/binary structure of color experience (e.g. De Valois & De Valois 1993; Wuerger *et al*. 2005). While many neurons in the lateral geniculate nucleus are tuned in the two cardinal opponent directions, it has been found that neurons in visual areas in the cortex 'differ widely in their chromatic preferences' (Gegenfurtner 2003). In the visual cortex of monkeys, for example, some cells respond to the wavelength of the stimuli, while the activity of other cells correlates more closely with surface colors as perceived by humans (Zeki 1983). Instead of containing opponent channels, the macaque primary visual cortex may contain 'hue maps', in

which perceptually more similar chromatic stimuli activate adjacent neural substrate (Xiao *et al*. 2007). Also neuroimaging research of human visual cortex has revealed neural activity in visual cortex correlating with chromatic stimulation (e.g. Parkes *et al*. 2009; Brouwer & Heeger 2009). There are correlations between neural activity and various aspects of color, and these seem not particularly restricted to correlations to the unique colors.

While there is ample evidence of neural activity correlating to hues (Gegenfurtner 2003), no opponent processes have been discovered to correlate with the unique hues, and no special selectivity to the unique hues has been found. Thus the interpretation of opponent processes as neurophysiological correlates to the unique/binary distinction is certainly not to be taken for granted. According to present knowledge, the extremes of the activation of known opponent channels do not match with the unique colors, and a well-supported neurophysiological account of the unique/binary distinction is presently unavailable.

3.2. Logical issues with neural correlations

Suppose for the sake of argument that a neurophysiological correlate of the unique/binary difference will be found, within or without an opponent processing framework. Would that provide support for the hypothesis that the structure of color experience has a neurophysiological origin?

There is a fundamental difficulty with the proposed tight link between neurophysiology and the explanation of the structure of color experience. The reason is that, if we find neural activity that mirrors the unique/binary structure of color experience, there is a further question regarding what gives rise to and therefore explains this structure. For an unknown factor in the environment may structure neurophysiology, in which case the structure of colors may be thought of as deriving from the environment. As a result, neural correlates themselves have limited value for explaining the unique/binary structure of color experience as long as it remains obscure what structures the neural activity. If environmental factors, which are not themselves defined in terms of neurophysiology, are determining the structure of neurophysiology, then the explanation for the structure of color experience does not derive exclusively from the peculiarities of neurophysiology. The hypothesis that the structure of colors derives solely from the brain depends on excluding other potentially relevant factors.

As a result, it is hard to find sufficient empirical support for the second strand of the argument from structure. For a long time, a promising stimulus-involving account of the unique/binary structure of color experience has not been available. This might have reinforced the level of

confidence in an exclusive neurophysiological account. But if an environment-involving explanation can be given this disqualifies the claim that the structure of colors derives exclusively from arbitrary peculiarities of neurophysiology. The still popular hypothesis of color opponent processes underlying to the unique/binary distinction may not only be empirically flawed; also the reliance on the explanatory status of the hypothesized color opponent processes may be misconceived.

4. An environmental contribution to the unique/binary structure

Models in which the structure in the environment plays a crucial role in accounting for the unique hues can, for the purposes of this chapter, be taken as being of two sorts. We shall speak of pure and mixed models. Though the two kinds of models are related in their orientation towards the environment, they differ in the degree in which they make reference to properties of the visual system. Pure models do so only most minimally. Besides being concerned with *visible* light as it reflects off surfaces, such models don't rely on further properties such as the individual sensitivity spectra of the human cones, in their account of uniqueness. A mixed model more extensively refers to the human visual system, stopping short, however, of neural factors which themselves explain uniqueness. Such a model could refer, besides to environmental factors, to the specific sensitivities of the human photoreceptors. Importantly, both types of models meet the requirement of explanatory independence, as set out in Section 2. This is obviously the case for pure models, as they relate uniqueness directly to environmental properties, such as surface reflectances. But it applies to mixed models as well, even if they rely in their explanation of uniqueness on the receptor sensitivities. The reason is that the cone sensitivities do not independently explain the unique/binary structure of color experience. Indeed, the mismatch between the three photoreceptors and four unique hues is an important reason why theorists have proposed that the explanation for uniqueness lies at a post-receptor level.

The contours of a pure model can be found in Broackes (2011). Building on observations and suggestions of Shepard (1992), and Mollon (2006), Broackes points to the fact that variations in direct sunlight and skylight lie along a line between unique yellow and unique blue. Roughly, direct sunlight and skylight objectively fix unique yellow and unique blue and provide an easily available reference for calibrating color vision (Shepard 1992; Mollon 2006). Broackes indicates how an account of unique yellow and blue surfaces can be built from this. On this account, unique yellow

surfaces are those that reflect most strongly under direct sunlight, and that gain most in luminance when the light changes from skylight to sunlight (Broackes 2011, p. 622). Blue surfaces would then be those that, compared to other surfaces, are particularly strongly reflecting skylight, and that thus gain most in relative luminance when the light changes from sunlight to skylight.

Broackes suggests that these special characteristics of yellow and blue may form the basis, not just for these colors to form natural points of reference for color vision, but also for these colors to be experienced as particularly 'pure', and suitable primary elements for other colors. The idea here is that these colors, given that they match with and minimally darken generally available light, are likely to be taken as neutral. At the same time, these colors may be salient, or characterful, for yellow and blue objects are often more saturated than the sunlight or skylight. As Broackes puts it: "the colors of these illuminants give us *both* a kind of neutrality (in the relatively unsaturated cases) and untingedness and purity, *and* (in the more saturated cases of the same hue) a kind of *primary* characterfulness" (Broackes 2011, p. 625). He proposes that the minimal way in which these colors darken the light under common conditions may give them a simplicity that makes them suitable elements from which other colors would seem to be composed. The account sketched by Broackes may then apply to a strong notion of uniqueness, according to which other colors appear composed of elementary unique colors. As he admits, more exact development of these ideas is desirable, and a lot more is needed to extend the account so as to encompass red and green.

A particularly interesting aspect of this account is that it is independent of the details of our photoreceptor sensitivities. For even with different cone sensitivities, blue surfaces will be sampled as reflecting more strongly under increasingly blue illumination (Broackes 2011, p. 622). In such an account, the special status of yellow and blue would then come to the fore independently of the details of our sensory apparatus.

A mixed model is provided by Philipona and O'Regan, who do take into consideration the sensitivities of the human photoreceptor pigments (see Philipona & O'Regan 2006, and Vazquez-Corral *et al*. for a further elaborated version). That is, Philipona and O'Regan studied how surfaces reflect light, but they were concerned specifically with how the light has effects on the three types of cones in the human eye. Thus, instead of the physicists 'reflectance function' of a certain surface, which shows what proportion of incoming light is reflected at any wavelength, they constructed a more biologically constrained reflectance measure, which takes into account only the light, both illuminating and reflected, as it affects the human photoreceptors.

Philipona and O'Regan plotted this biological reflectance measure for a wide variety of surfaces. They found that it was possible to characterize the biological reflectance function for a particular surface by a 3x3 matrix. Such a matrix can be seen as specifying how a particular surface transforms any incoming light into a specific activation pattern in a receptor space. For most surfaces the receptor space is three-dimensional – which is what one would expect, given that there are three receptors whose responses can vary independently. Philipona and O'Regan came to the interesting observation, however, that some restricted classes of surfaces stood out, in that the activation patterns in receptor space were simpler for the lights they reflected. While the description of the illuminant requires three dimensions, some surfaces transform the light into a light that only needs one or two dimensions to be accurately described. In other words:

> "certain surfaces have the mathematical property of being "singular". What this means is that these surfaces take incoming light, which usually can vary in a 3-dimensional space defined by L, M and S cone activations, and transform it into light which varies only in either a 2- or in a 1-dimensional subspace of the LMS activation space. Because singular surfaces reduce variability from three dimensions to two or one dimension, they can be said to display a simpler behavior as concerns how they affect incoming light than the majority of surfaces." (Vazques-Corral *et al*. *forthcoming*)

The real interest of this finding was revealed when it turned out that the most strongly singular surfaces were quite exactly those surfaces picked out as 'focal red' and 'focal yellow' in research on focal colors, while 'focal blue' and 'focal green' corresponded to two other less strongly singular surfaces (see the plot on p. 335 of Philipona & O'Regan 2006). Moreover, when surfaces with those focal hues were illuminated by the light source standardly taken as approximating daylight (D65), the light reflected off them showed accurate correspondence with light associated in psychophysical experiments with unique hues.

The Philipona and O'Regan approach shows how an explanation for the unique/binary structure of color experience is possible by appealing to nothing more than the patterns of sensory stimulation as available to the photoreceptors in the human eye. No opponent processes or other neurophysiological structures that match the structure of color experience have to be presupposed. In the words of Philipona and O'Regan:

> "these facts, not previously noted, have been obtained without appealing to neural mechanisms that underlie opponent channels in the visual system. They are merely a consequence of the asymmetries in surface reflecting properties induced by the limited way biological

> photoreceptors sample physical spectra. It could thus be argued that the reason the colors "red", "yellow", "green", and "blue" are so often singled out among all colors as being worth giving a name, is that surfaces of these colors have the particularity that they alter incoming light in a simpler way than other surfaces (…)." (Philipona & O'Regan 2006, p. 336)

It is tempting to make sense of these findings in a framework which emphasizes the temporal extendedness of color perception. In this framework color perception is conceived of as becoming sensitive to color as a permanent property of the surfaces by means of sensitivity to the way the surface behaves over time under different illumination conditions. Such changes can be brought about by different means. Sometimes they might involve motion of the object, which may or may not be brought about by the perceiver. Alternatively, they might involve changes in the direct illumination, or in the indirect illumination, if surrounding objects move. Over time, one could become sensitive to the underlying reflectance profile of a surface on the basis of minimal cues, so that one could detect the profile by just noticing the behavior of the surface in a minimal set of lighting conditions. Such a minimal set, moreover, is almost always available, as it exists when different nearby objects reflect differently on a surface (Ruppertsberg & Bloj 2007).

In such a framework it is plausible that the singular surfaces of the Philipona and O'Regan analysis would be perceptually salient, for they would stand out as bringing about less change – and therefore as being simpler – whenever illumination conditions change. Such simple colors would be obvious points of reference for color vision and provide structure to color experience.

To sum up, the analysis of Philipona and O'Regan (2006) shows that there are systematic differences between focal red, green, yellow, and blue, compared to the other colors, and between the unique and binary hues. It can be derived from objective measures that surfaces with unique colors provide simpler patterns of sensory stimulation, compared to other colors. Rather than offering an *ad hoc* re-description of the stimulus, the analysis is based on independent evidence of surface properties and our retinal sensitivities, and it does not involve idiosyncrasies of human neurophysiology beyond the receptor sensitivities. The analysis strongly suggests that there is a crucial environmental contribution to the unique/binary distinction in color perception, for the environment contributes to the relevant patterns of sensory activation.

5. Consequences for the argument from structure

According to an assumption that is widely shared in philosophical writings on color, the unique/binary structure of color experience is due to the structuring of experience by the peculiarities of neurophysiological processes, in particular the opponent channels. We have seen that, in contrast with this assumption, the explanation for the structure of color experience may crucially involve environmental factors. Below we discuss the consequences of this finding for the two strands of the argument from structure.

5.1. Consequences for the first strand

The first strand of the argument from structure leads to the conclusion that strictly speaking physical colors do not exist. Does it follow from the alternative accounts considered that this conclusion is false, and that physical colors exist after all? It should be clear that only a pure model can establish what would count as a proof of the physical reality of surface colors according to the standards set by the argument from structure. For according to the argument from structure, such a proof would only be obtained if it would be established that there is some distal physical structure, completely independently specifiable of "the peculiar characteristics of the human visual system" (Hardin 1993, p. 67), which is isomorphic to the structure in experience. A pure model comes as close to this as possible. Of course, even a pure account remains tacitly committed to a number of assumptions linked to human vision. For example, it is only concerned with *humanly visible* light, without even considering the realms of the ultraviolet or the infrared. But a demand for total independence is not reasonable, whenever the topic of theorizing is perception. Evidently, any perceivable property can only be perceived by a creature that is sensitive to it – so there will always be organismic dependence in an account of perception.

Therefore, on any reasonable assessment, it must be granted that, in providing properties of lights and surfaces that single out those that are seen as uniquely yellow and uniquely blue, in a way that is independent of but the most general characteristics of the visual system, pure models establish that, counter to the premise in the first strand of the argument form structure, there is an environmental basis for the unique/binary distinction.

A moment's reflection shows that not only pure, but also mixed models refute the premise of the first strand of the argument from structure. For, although mixed models appeal to more fine-grained aspects of sensory sensitivities, they are not thereby committed to a vicious kind of

explanatory dependence as regards uniqueness. The characterization of the unique hues provided by Philipona and O'Regan takes into account how light affects our photopigments, but it will be difficult to resurrect in the light of it the idea that the explanation of uniqueness is not a matter of environmental physics but rather of neural processes. The reason is that the photopigment sensitivities do not themselves provide an explanation of uniqueness – unlike what color opponent physiological channels are hypothesized to do. The latter point is crucial in this context, because the exclusive appeal to neural structures such as color opponent channels is essential to the argument from structure. It is precisely the alleged neat mapping between the structure of neurophysiology and the structure of experience, together with the absence of such mapping between the structure of experience and the structure of the physical stimulus (if properly, independently characterized), that is supposed to lead to the conclusion that colors or color experiences reduce to neurophysiological properties. The appeal to the photoreceptor properties in the Philipona and O'Regan model does not assume such a direct mapping between these and neural properties.

Although the Philiponna and O'Regan theory does not invalidate the opponent channel framework, it undeniably does shift the explanatory weight away from opponent channels and it provides a better explanation of the loci of the unique hues.[3] This is not to say that the Philipona and O'Regan approach is incompatible with the possibility of post-receptoral structuring. But in providing an alternative account of the loci of the unique hues, it does constrain the explanatory scope such structure might have. If the model is correct, singling out a surface as unique or not unique does not necessarily require processes in neurophysiological opponent channels. For without taking into account any fact about opponent channels at all, it can be predicted by the model which surfaces will have a special status for perceivers like us. The Philipona and O'Regan approach thus allows for a characterization of the unique/binary distinction by means of a description which does not refer to hypothetical color opponent channels.

---

[3] The fact that it now appears that the structuring effects of opponent processes have been overestimated is fully consistent with the proposal that opponent processes optimize information transfer from the eye to the brain (Buchsbaum & Gottschalk 1983).

5.2. Consequences for the second strand

The second strand in the argument from structure requires that the explanation of the structure of experience be fundamentally or (in our terms) independently neural. It is quite clear that it is difficult to bring this strand to its conclusion if one of the models discussed in Section 4 is true. For these models provide explanations of uniqueness which are in the required sense independently environment-involving. If true, they show that there *is* a genuine environmental explanatory basis for the structure of color experience. If so, the crucial premise that such environmental explanations do not exist should be rejected. The structure of color experience then does not provide support for the idea that color experiences are purely neural properties. In other words, the second strand of the argument from structure depends on the first strand. It is concluded that color experience belongs exclusively to the neural realm because of both the supposed absence of a genuine basis for structure in the environment and the supposed presence of such a genuine basis in the neural realm. Clearly, whenever there is an external structure, the conclusion of this line of argument no longer holds.

As we have seen, there are serious problems with the neural model that traditionally has been assumed to bring the explanatory goods, namely a physiological opponent process model. Still it remains possible that, even if pure or mixed models are correct regarding an environmental basis for uniqueness, some neural structures might be found which might have 'internalized' the structure of the environment in the organism's sensory physiology.[4] It may then be argued that such neural structure does, after all, carry the explanatory load regarding color experience. By our own explanation bound criteria for 'genuine existence', this might seem to imply that color experiences are purely neural after all.

However, there are reasons to be skeptical about presupposing such a shift of the explanatory basis towards the brain. In the first place, it remains to be seen whether such 'internalization' of the unique/binary structure of hues is more than a possibility. Indeed, the more successful the pure or mixed environmental models for this structure, the less need there seems to be to let the brain do the structuring. Powerful if controversial arguments have been given in favor of the 'offloading' of structure from the brain into the environment concerning other explananda in vision science such as visual stability or the filling-in of the blind spot (O'Regan 1992; 2011). Moreover, even if an internalization of structure were to be found, it

[4] We thank a reviewer for raising this issue, and part of its formulation.

would still itself have to be explained by the external structure. The latter is what is internalized. It comes first, in the temporal and in the explanatory order. This explanatory dependence remains even if, after the internalization has occurred, neural factors alone might suffice to bring about color experience (see also Hurley 2010).

5.3. Conclusion

The argument from structure should be rejected: the unique/binary structure of color experience does not support the claim that colors have no environmental basis at all, and neither does it legitimate the conclusion that color experiences only genuinely exist as neurophysiological phenomena. We highlighted analyses of structural aspects of color focused on the way the human eye is affected by the light reflected by surfaces, which indicate an objective explanatory basis for the unique/binary distinction without reference to peculiarities of neurophysiological processes beyond the basic sensitivities of our photoreceptors. From such analyses, it can be concluded that, while perceptual experience must of course depend on the perceiving organism, the explanation for the structure of color experience is at the same time crucially world-involving. The ontology of color and color experience should respect this wide explanatory basis.

Chapter 6
Through the inverting glass

Experience with inverting glasses reveals key factors of spatial vision. Drawing on my experience with wearing left/right inverting glasses, I show how a sensorimotor analysis helps to describe visual experience, and how a puzzle raised by the literature resolves. I further report my findings on mental imagery, supporting the idea that imagery is grounded in sensorimotor engagement with the environment.

1. Introduction

There is a window on my left, a desk in front of me, and there are books to my right. When I look around, I effortlessly see where things are, and I perceive the world as stable while my eyes sweep across the scene. We are so familiar with sight that we tend to take for granted the experience of the stability and the spatial layout of the visual world. But when you think about it, it is not so evident why things should appear as they do. As we look around, the images sweep across our retinas, but at the same time, the world does not appear to move. What is this phenomenon of visual stability? And what does appearing to be at the right consist in? In this chapter I take up these questions, drawing on my experiences with wearing left/right inverting glasses. I offer a sensorimotor characterization of key aspects of the phenomenology of spatial vision and I report my findings on visual imagery.

In normal circumstances, several factors work together to yield a coherent visual experience. We cannot normally differentiate between these factors, so that first-person reflection on vision is hampered by the perfection of our visual skills. Inverting glasses provide a partial remedy. By means of lenses, mirrors, or prisms, the light reaching the eyes can be altered such that left and right and/or above and below are inverted. Such an experimental manipulation of vision may help to address questions of visual orientation and stability, for it removes the veil of familiarity and it helps to disentangle what is usually tied together. Since one may get used to inverting glasses, wearing them for extended periods of time allows us to investigate, say, whether the typical orientation of the retinal image is required for upright vision (normally the retinal image is inverted with respect to the distal stimulus), or whether the typical stimulation of the hemispheres of the brain is required for the experience of objects as being

at the left or the right side (normally the experience of stimuli at one side of the visual field involves the visual areas of the opposite hemisphere). Inverting glasses introduce systematic changes in the relation of retinal stimulation to the environment, to action and to the other senses. With inverting glasses we can investigate how the introduced changes affect visual experience. In consequence, vision with inverting glasses may provide a window into the determinants of spatial visual experience.

The experiential consequences of wearing such glasses are of particular interest from the perspective of the sensorimotor account. According to this account, as we have seen, the way in which a person subjectively experiences the environment can be described in terms of the relation between sensory stimulation and motor action (e.g. O'Regan & Noë 2001; Hurley & Noë 2003; O'Regan 2011; Chapter 2 above). This gives rise to characteristic patterns of sensorimotor engagement with the environment, such as that in vision you can sample the scene by looking around, while tactile exploration gives very different patterns of dependency of sensory stimulation on bodily movement. Inverting glasses systematically transform the sensorimotor patterns of visual experience, and therefore allow us to explore how such changes affect experience. Wearing these glasses can thus help to identify the sensorimotor patterns that are characteristic of spatial vision.

In this chapter I report on my experience with wearing left/right inverting glasses. By refection on my visual phenomenology, that is, on what my experience with inverting glasses was like, I aim to yield insights into spatial vision. A disclaimer seems fitting. While I aimed to contribute to a characterization of spatial vision, the purpose of my enterprise was not so much to generate new data on the precise course of adaptation in one subject under one training regime. Rather, I set out to wear inverting glasses to provide an occasion for first-person reflection on spatial vision, after reading some reports on experiments with inverting glasses. It further must be noted that, being exposed to visual theory for some years, I am not a naive subject. But this can be considered as an advantage rather than as a disadvantage, for it may help to keep phenomenological descriptions in touch with cognitive science. I have in particular been engaged with the sensorimotor account of perceptual experience, which in my view may help to articulate what the experience is like in a way that makes contact with the underlying processes (see especially Chapter 2 above, and Chapter 4) Here I shall show how a sensorimotor analysis can contribute to the personal level description of visual experience.

Previous studies have raised the question whether visual experience may turn back to normal after adaptation to inverting glasses. While some reports, when carefully read, suggest it does not (e.g. Stratton 1897), others

seem to suggest it does (e.g. Taylor 1962; Kohler 1964). Using my own experiences as the starting-point for reflection, I shall argue that a more fine-grained sensorimotor analysis could resolve the issue. Below I shall first discuss some classical studies on vision with inverting glasses (Section 2). I shall not provide a review of the literature, but merely point towards a question raised by previous reports. To sketch the background for my findings I will then describe the glasses I used and the general course of the experiment (Section 3). After this I provide a description and interpretation of the main findings, relating to visual stability (Section 4), experienced location and left/right orientation (Section 5), and visual memory or imagery (Section 6). I conclude that inverting glasses introduce a conflict at the very heart of spatial vision, that a sensorimotor description is crucial for an adequate characterization of the resulting visual phenomenology, and that such a characterization may resolve a question raised by the literature (Section 7).

2. A puzzle from earlier studies

In a pioneering study on 'vision without inversion of the retinal image', George M. Stratton (1896; 1897) used lenses that inverted both the left/right and the up/down orientation. He wore these lenses full-time for several days: 3 days in the first experiment, 8 days in the second (he used the lenses for one eye; the other eye was covered). In the beginning of the experiment, the hand that would feel as on the lower right, where it was, would visually appear as if it were at the upper left. There thus was a conflict between vision and touch, and the visual appearance of things no longer conformed to their actual location. Frequent inadequate behavior was the predictable result, and Stratton increasingly learned to cope with the glasses. But the most interesting thing was that as the experiment progressed, experience itself started to change. Stratton reports that later in the experiment "the limbs began actually to feel in the place where the new visual perception reported them to be" (Stratton 1896, p. 615).

This may suggest that the conflict between touch and sight was resolved as a result of the adaptation of the felt position of the limbs to visual experience. If touch adapts to vision, we may suppose that visual experience was still inverted, while the experience of the body had adapted to the new vision. But Stratton reports that in a sense 'upright vision' was in fact restored, as vision and touch were once again experienced as harmonious, and he concludes that upright vision is possible without the usual inversion of the retinal image:

> "The inverted position of the retinal image is, therefore, not essential to 'upright vision,' for it is not essential to a harmony between touch and

> sight, which in the final analysis, is the real meaning of upright vision. For some visual objects may be inverted with respect to other visual objects, but the whole system of visual objects can never by itself be either inverted or upright. It could be inverted or upright only with respect to certain non-visual experiences with which I might compare my visual system—in other words, with respect to my tactual or motor perceptions." (Stratton 1897, pp. 475-476)

We should not conclude that the spatial phenomenology of the reported 'upright vision' is similar to the pre-experimental upright vision. The reason is that 'upright vision', for Stratton, means that vision and touch are in agreement – that you see things where you feel them. Moreover, Stratton suggests that perhaps neither sight adapts exclusively to touch, nor touch exclusively to sight (Stratton 1897, p. 472). In order to say more about the resulting visual phenomenology we must consider a third crucial factor in Stratton's report, besides touch and sight, namely *visual memory*. Stratton notes that the 'memory images' from before the experiment "preserve a spatial arrangement whose lines of direction were opposed to those of the actual field of view" (Stratton 1897, p. 472). This testifies to nontrivial differences between the 'upright vision' before wearing the glasses and the reported 'upright vision' after adaptation to the glasses. Although Stratton characterized his experience with inverting glasses as eventually 'upright', his comparison of vision with visual memory suggests that visual experience remained importantly different from normal upright vision (see also Harris 1965).

Some subsequent findings, in contrast, seem to go against the idea that the spatial phenomenology of vision with inverting devises remains different after adaptation. In *The Behavioral Basis of Perception*, James G. Taylor (1962) reports an experiment in which a subject was wearing left/right inverting glasses part-time. At the eighth day of the experiment, the subject visually experienced a chair "as being both on the side where it was in contact with his body and on the opposite side", and the experience of the chair, which in reality stood at the right, was "like the simultaneous perception of an object and its mirror image, although in this case the chair on the right was rather ghost-like" (Taylor 1962, p. 202).[1] The occurrence of these curious experiences did not last long, as visual experience adapted

[1] A similar report can be found in the work of Kohler, who writes that a subject, after wearing left/right inverting glasses, "was capable of seeing two points of light when only one was presented (and this happened even monocularly)" (Kohler 1964b, p. 161).

towards accordance with the actual location of objects. Such adaptation suggests that the ghostly experiences strengthen, and gradually lose their ghostly appearance until they eventually take over visual experience. Given that these experiences are the 're-inversion' of the artificially inverted vision, we may suspect that this would leave the subject with an experience similar to the normal visual experience.

If we were to conceive of visual phenomenology in terms of having 'visual images', there would appear to be a conflict between the reports of Stratton and Taylor. Stratton's report would then suggest that the subjective 'visual images' remain inverted after adaptation to inverting glasses, while Taylor's report would rather suggest that the 'visual images' turn back to normal. This apparent conflict may of course be due to differences between the experiments or they may reflect inter-individual differences in visual phenomenology. In some circumstances or in some people visual experience may be more prone to change during adaptation to inverting glasses. But it is also possible that a more fine-grained analysis can help to resolve the conflict.

The need for a more subtle analysis is suggested by the finding that different aspects of visual experience may adapt independently. For example, it has been reported that after adaptation to inverting glasses, when much of the world was reported to be seen as having its actual orientation again, other parts, such as letters, still seemed to be inverted (Kohler 1964b, p. 155). If we were to conceive of visual experience in terms of images, we may wonder what kind of an image it would be with such exotic properties, allowing for partial inversion (Taylor 1964, p. 73). For example, how can the words on a signpost appear inverted if they start at one side of the signpost and end at the other, if we believe that the image of the signpost is not inverted? Instead it seems more parsimonious to accept that visual experience is not like a unitary image that may or may not be inverted. Perhaps visual experience may then better be viewed in terms of the various subsystems or perceptual skills that are brought to bear on the environment, as various authors have argued (e.g. Taylor 1962, p. 207; O'Regan 2011).

However, if we view perceptual experience in terms of perceptual skills rather than in terms of having 'images', the question still remains how we can reconcile the findings of Stratton and Taylor. After all, it is not the case that the one reports, say, inverted letters while the other reports a non-inverted rest of the world. As I shall demonstrate below, a sensorimotor interpretation of perceptual experience offers a possible solution. The reason is that within such an interpretation, there is room for ambiguity at the very basis of perceptual judgments. There may then be aspects of visual experience that become normal after adaptation to inverting glasses, while

at the same time differences in experience remain. While one author may have been more impressed by the similarities between normal vision and vision with inverting glasses, the other may have been more taken by the differences, resulting in the apparent conflict between the reports. Through careful reflection on vision with inverting glasses we may then resolve the conflict, and get a better view on spatial visual phenomenology.

3. Donning the glasses

The glasses I used were a simple device with a prism in front of each eye. This resulted in a left/right inversion of the light coming from the visual field, so that on first use of the glasses, distal stimuli at the left side within the field of vision subjectively appear as being on the right and vice versa. At the same time, moving the head to the right brings in view the objects at the right, as during normal vision, although of course they appear on the left. Looking straight ahead, what visually appears as being at the right therefore goes out of view first when turning or shifting the head to the right. When one is not used to the inverting glasses, one has to shift one's head leftwards in order to look behind the side of the object that visually appears as being on the right.

While inverting glasses do not alter the relation between head movements and the part of the world that is seen (the distal field of vision), they do alter the relation between head movements and the proximal stimulation. In this respect, the consequences of wearing inverting glasses are the opposite for eye movements. When defined in relation to the distal stimulus, inverting glasses do alter the relation between eye movements and focal vision – the place of highest resolution in the center of the view of the eyes. For the proximal stimulus, in contrast, inverting glasses leave unaltered how eye movements relate to retinal stimulation: whether or not one is wearing inverting glasses, the light falling on the left side of the retina can be brought to focal vision by turning one's eyes towards the right. As long as the head remains stationary, visual exploration of objects by means of eye movements therefore provides no difficulties. For a schematic illustration of prominent consequences of using the glasses, see Figure 4.

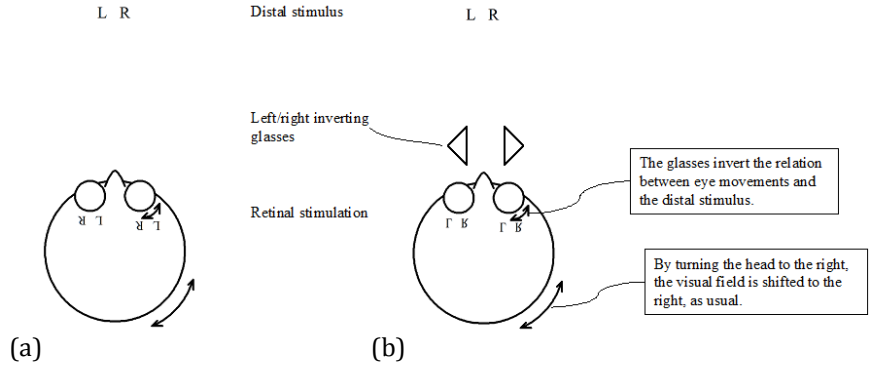

Figure 4. Schema illustrating consequences of the left/right inverting glasses (b), compared to vision without inverting glasses (a). Note that the glasses invert the relation between the distal stimulus and eye movement (e.g. with inverting glasses objects at the left can be brought into focal vision by an eye movement towards the right), but that with eyes aimed forwards, the direction of the head determines the center of the distal visual field as normal (e.g. head movement towards the right brings the right side of the visual field in central view). Further note that, while the relation between retinal stimulation and the distal stimulus is altered, the relation between retinal stimulation and eye movements is unaltered (see text).

I wore the glasses superimposed on my normal glasses, wearing a hood against false light. The glasses strongly restricted my visual field. At arm's length, the scope was about two hand's broad, with the width of one hand of stereoscopic vision. The vertical reach of sight was much more generous. As a result of my limited scope, I had to make scanning movements with my head to acquire a reasonably rich impression of my surroundings. When outside I used a white stick to signal my self-imposed visual handicap.

I had no fixed scheme for wearing the glasses. At some days I did not have the opportunity to wear them at all, but most evenings I did. In the thirty-one days of using the glasses, I wore them on average 4 hours and 8 minutes per day, resulting on a total of over 128 hours at the 31st day of wearing the glasses (43 days after starting the experiment). But part of this time my activity was rather limited, especially when I watched movies, which required eye movements and minor head movements only. Other activities I engaged in included typewriting, cooking and doing the dishes, playing the board game *go* with a friend, and going for a walk. I started walking outside only after 15 days of wearing the glasses. As in the experiments of Stratton (1896; 1897) and Kohler (1964b), I used no systematic training program. With systematic training, quicker adaptation could have been expected (Taylor 1962).

On the first day of wearing the glasses I got sick. Not only did I experience a lack of visual stability, I also saw double, except at arm's

length. After a few hours I suddenly got so sick that I threw up, and felt weak for the rest of the evening. I therefore decided to take my time to get used to the glasses. This I did, when I got the opportunity to wear the glasses again some days later, by engaging primarily in simple activities, such as watching movies. Quickly I could also walk around (though clumsily) and I haven't been sick at any further moment of the experiment. The third day, the reach of proper stereovision was already much larger, and the ability to accommodate my eyes increased.

I often reached in the wrong direction, even when I knew where objects in my room were located. Vision tends to overrule knowledge, and to a certain extent habits are cancelled or transformed. A notable behavioral impairment expressed itself when I attempted to replace a cup that was standing too close to the edge of a table. It was almost impossible to find the right direction. Trying to correct the movement, I instead altered it in the wrong way. Even days later, cutting tomatoes still had a similar effect: with eyes wide open, the appropriate orientation of the knife was almost impossible to bring about.

The visual effect of a sideward inclination of my head was that the world appeared tilted; when I did this while standing on one leg, it was difficult to keep my balance (while I can easily do this without glasses or with eyes closed). Walking down the street, at some point I ended up at a slant. It took me a while before I could recover my balance, no doubt due to the erroneous alteration of my posture based on unmastered vision. Clearly, the visual modality is quite dominant, at any rate for me, not just in case of object-oriented action, but also for keeping my balance.

When I took off the inverting glasses after I had worn them for some hours, head movements disrupted visual stability. These after-effects, which can be considered clear signs of perceptual reorganization, often lasted for over half an hour. At a larger timescale, some of the mornings after wearing the glasses my head felt heavy. During the course of the experiment, in general the after-effects decreased in strength, as I became used to the alternation of vision with and without inverting glasses.

A few times, I also experienced more striking after-effects. At the 8th day, while I had not yet donned the glasses, I noticed that I sometimes moved the mouse of my computer in the wrong direction. Also on the afternoon of the 10th day I had not been wearing the glasses yet. While I was reading an article, I moved my hand to turn the page and I found myself seriously surprised at the sight. I had not anticipated that turning the page with my right hand would look like that; I had rather expected the other hand to turn. I did not experience such somewhat disturbing after-effects during the rest of the experiment.

When I was wearing the glasses, throughout the experiment my behavioral skills increased. On the 4th day I cooked a simple meal, which I wouldn't even have tried at the beginning of the experiment. I also developed strategies for walking: by looking in the direction where I had to go, I somehow managed to automatically correct my way. This way I zigzagged through the hall, and throughout the experiment my path got straighter and straighter. At the 1st and the 12th day I tested my skills of tracing a circle with a pencil, confirming a serious increase of skill approaching the level reached without inverting glasses.

After the 15th day of wearing the glasses (67 hours), I started walking outside, using my white stick for safety. The first time it took me about an hour time to walk home from my office at the faculty of philosophy (without the glasses it takes me less than half an hour). After eight more days of practice it took me less than 35 minutes. But even then I had remarkably limited awareness of the side of the road I was on. At narrow sidewalks this could be slightly disturbing at first, with the nearby noise when cars whiz by at the side where I wouldn't expect it. By attending to the direction of my head I could perceive which side was which, but without such deliberate attention this was at first certainly not evident in my experience, and even later such awareness often remained limited. In general, awareness of the direction of my gaze was better while seated or standing still than while walking, and it increased over the days.

With this general context in place, let me now turn to some themes that are of particular interest, concerning visual perception and imagery.

4. Visual stability

Normally the world appears as something stable through which we can move and within which we can look around. One of the most prominent effects of wearing the left/right inverting glasses was the initial break-down of this visual stability; head movements resulted in an apparent movement of the scene.

Without movement, the consequences of wearing inverting glasses for the experience of the environment are quite limited. Of course, things at the left seem to be located at the right and vice versa, and this does have some curious consequences. For example, I experienced that with letters appearing mirror-inverted, and with the inverted direction of reading, words lost their familiarity so that I often could read them only with effort. But in terms of more general aspects of spatial experience, things get more exciting only when one begins to move. For example, moving one's right hand visually appears as if one moves one's left hand, resulting in a strange conflict between vision and touch. In my experiment I have been mainly

concerned with my visual experience of the environment (rather than with experiences of the outside world on the basis of other sense modalities, or with experiences of my body). The movements that make things interesting are then the movements of my head and the movements of my eyes.

Not being used to inverting glasses, the effect of turning my head was a serious disruption of visual stability: when I moved my head in the horizontal plane, the scene appeared to sweep in front of my eyes in the same direction as the head movement, but at higher speed.

This experience of the sweeping of the scene was not like watching a movie that is shot with a sideways moving camera, in which the images sweep across the movie screen. In case of the movie, one may effortlessly follow objects with one's eyes and head to get a good view of them. Precisely this ability to track objects was reduced when I moved my head while wearing the glasses. If I would attempt to track the sweeping scene by moving my head, the result would only be that the scene swept even stronger. What I was lacking was, as it were, a firm visual grasp of the scene. Although my experience of location was impaired, I did experience objects as spatially located. But with even the slightest head movement, my experience of location did not amount to the robust sense of direct contact with external objects or projections, as characteristic of normal vision.

While it is hard to describe my experience at this point, it must be noted that a description of my experience in terms of 'visual images' would be incomplete at best. As said, the lack of grasp on the visual scene was a prominent aspect of my visual experience; this aspect is not captured in the 'images' metaphor of vision. For this reason it is more accurate to characterize the experience as one in which the whole visual scene appeared to sweep untrackable in front of my eyes.

The experience of the sweeping of the visual scene could not be the result of the narrow scope of my field of view. For when I merely restrict my field of vision this leads only to a subtle instability of visual experience, far less pronounced than with the inverting glasses. No doubt the main factor responsible for the disruption of visual stability was that the relation between retinal stimulation and the movement of my head had been altered. The breakdown of visual stability depends both on sensory stimulation and on movement-related factors. In consequence, the re-acquisition of visual stability must involve adaptation to the transformed sensory-motor relation.

During the course of the experiment, the experience as if the scene swept across my eyes gradually decreased. It was on the 13th day of wearing the glasses, after over 57 hours of wearing them, that I reached visual stability. Visual stability was quite absolute at this point, roughly as good as it is for me during normal vision. Interestingly, the onset of visual

stability seemed quite abrupt, and it was with no little excitement that I noted this 'perceptual breakthrough'. At this point, I could look around without having a sense of sweeping of the scene, although too fast movements of the head would still disrupt visual stability. Also walking disrupted stability, perhaps due to the fact that it resulted in head movements that were not the consequence of actively turning my head. But when I was standing still I could move my head, resulting in the feeling of looking around in a stable world.

Although the world appeared stable, I still experienced the scene as left/right inverted. I'll address the strange consequences of this in Section 5 below.

The experience as of a stable world was still fragile. Not only did it break down when I moved too much or too quickly, there were also more specific causes for disturbance. One of these was tilting my head, which still resulted in the experience as if the visual world rotated. Another condition that frequently led to breakdown of visual stability was when I attempted to look over my shoulder. Although I experienced visual stability when I looked around in different directions in front of me, sudden attempts to catch a glimpse of something behind me still resulted in the experience as if the world moved in front of my eyes.

It turned out that, when attempting to look over my shoulder, I tended to direct my eyes in the way appropriate only to perception without inverting glasses. Without inverting glasses, moving my head to look over my shoulder automatically engages eye movements towards the same side. This tendency was still in place, but with inverting glasses this led to turning my eyes away from the parts of the scene that newly entered into view. As a result, my attempt to look over my shoulder often failed: I tended to turn my eyes in the wrong direction, so that a view of the objects behind me was not so quickly and automatically obtained as it usually is. A crucial consequence of the misdirection of my eyes was that objects at focal vision could not be tracked with my gaze – after all, they immediately disappeared from view. Looking over my shoulder then still came with a breakdown of visual stability, even though visual stability was already becoming the standard when looking around at the scene in front of me.

In fact, as I later figured out, the relation between the direction of the eyes – as well as the possibility to keep track of objects – and visual stability can be confirmed without any inverting lenses. To try this, aim your head and eyes completely to the right. If you now turn your head to the left, while keeping your eyes turned completely to the right (this may not be easy, but it is crucial), you will be able to notice that visual stability breaks down. In my case, this happens even at low speed, while much higher speeds of

turning my head would normally still result in the experience of a stable visual world.

From a first-person perspective, the return of visual stability may be described as a return of normal visual skills, such as the capacity to look over my shoulder, to track objects, or to keep objects in view while turning my head. The experience of visual stability cannot be captured in terms of visual images: it is rather a matter of having a firm visual grasp on the scene; or having the right exploratory skills in place.

5. The experience of left and right

As long as I looked straight ahead, without moving my head, the effect of wearing the glasses could initially be described as a left/right inversion of visual experience. But as soon as I moved my head this description failed. One reason for this was that visual stability broke down when I moved my head, resulting in an experience notably different from a simple left/right inversion. After reaching visual stability, it became more natural to describe my transformed experience as a dynamic left/right inversion. However, this description applied only to the experience of the relation between objects within my visual field. But there is more to visual experience than can be captured in terms of the apparent spatial relations between objects within the part of the environment that is seen. The experience of the spatial location of objects can differ depending on the orientation of the visual field in relation to the body, even if the objects are in the same position within the visual field. For example, an object in the middle of one's visual field can appear to be located at one's right or at one's left, depending of the direction of one's gaze. As a result, also after reaching visual stability, a description of vision with inverting glasses in terms of a simple left/right inversion should be rejected, as it fails to capture a crucial aspect of spatial vision.

Already from the beginning I often had a clear notion of the direction in which I looked when I turned my head towards the left or the right. But this certainly was not always the case. In the first week of wearing the glasses I noticed that I was not always sure which hand I used in activities such as typewriting, during which I had to make scanning movements with my head. I could use the hand I saw, anchoring my action in sight, without a definite sense of using my left or right hand, and without explicit awareness of the direction of my view. While typing was far from fluent (and remained so during the experiment, although it got better), the general position of the keys – whether left or right – did not seem to play too large a role in my performance.

When I didn't have my hands in view I needed to rely on my judgment of direction, for example when I tried to use my nearest hand to place stones at the selected location in the board game *go*. The sense of left or right became more prominent then. I developed a simple way to cope: by deliberately turning my head I could ground my action in the direction or the movement of my head, or in the general direction of my visual field, in that I could just use the hand on the side to which my head turned. With calm movements of my head I was often aware, without a need to think about it, of the general direction of my gaze. Thus I could correctly perceive the location of the objects (or board positions) that were in the center of my visual field, and use the most appropriate hand to reach out to the objects.

When I could look around without breakdown of visual stability (from the 13th day onward), the left/right inverted visual field appeared quite natural to me. This was so even though the objects that seemed to be at the right in my visual field disappeared from sight first when I turned my head to the right. While seemingly natural, this was a curious inconsistency in visual experience, for what I visually experienced as being at the right, immediately disappeared from sight when I moved my head towards what I (correctly) experienced as the right.

My spatial experience at this stage is well described by the following event. On the 21st day (after 86 hours of wearing the glasses) I had dinner in a house where I had never been before, entering the room wearing the inverting glasses. We talked, wined and dined, and by looking around I got an impression of the place. But when I first took off my glasses after dinner, I noticed that I had to adjust my idea of the room. For example, the couch that had appeared in front of me, moderately to the right, now turned out to be located next to me at my far right. This indicates awareness of the general direction of my visual field, but inversion within my visual field. I could have figured out the actual position of the couch if I had tried, but clearly I had not unreflectively registered its position.

Later, wearing the glasses again, at some point I looked at my legs while walking and noticed a curious phenomenon. The experience of the step I saw and the step I felt corresponded, but the curious thing was that, roughly speaking, vision was dominant. My right leg felt left. Apparently, even bodily feeling can be firmly anchored in sight. It seems that experience tended towards coherence: my feeling of bodily position at this point conformed to my – eye movement dominated – visual experience. During

most of the experiment no such altered bodily awareness occurred.[2] But it is worthwhile to describe the experience in some more detail.

To convey what my experience at this point was like, it will help to contrast my experience with the 'harmony between touch and sight' mentioned by Stratton (1897), which he considered a sufficient condition for veridical (or 'upright') visual experience. I must stress that the experience described above could not be described as veridical, not even on Stratton's terms. Given that I experienced the movement and direction of my head as usual, I had no difficulty in judging the incorrect nature of the tactile experience of my legs and the experience of their location within my visual field. While there was a considerable harmony between sight and the felt position of my legs, there certainly was no harmony between the sight of my legs and the felt head movements. Crucially, my awareness of the general direction of my view was still in conflict with the sight of the position of my legs within my view. For example, when I moved my head to the right, I was well aware that I looked to the right, but I thereby brought in view the leg that had appeared to be on the left (where it was felt to be as well). Thus only eye movement-based visual direction was in agreement with the (incorrect) proprioceptive experience of my legs, and my visual experience in as far as it was grounded in head movements was still conflicting with these experiences. For this reason, it would be more precise to speak, not of the dominance of vision over felt bodily position, but of the dominance of eye-movement dominated visual experience over felt bodily position. My experience could not be described as veridical because my spatial vision itself was still lacking the coherence required for veridicality.

On the 30th day of wearing the glasses I decided to wear the glasses all day. So I donned the glasses before opening my eyes. I had an active day, walking on the beach and on the narrow paths winding through the dunes, and in the evening I settled to engage in careful observation. As I had already noticed before, calmly looking around by moving my head helped

[2] By contrast, it has been argued that adaptation to inverting glasses may primarily involve changes of proprioceptive experience rather than visuomotor adaptation or changes of visual experience (e.g. Harris 1965). The experimental conditions are probably crucial to explain when primarily proprioceptive changes occur, and when visual changes. I would expect that where head movements are counteracted (as in most studies on which Harris relies), visual adaptation is counteracted, which may be favorable for proprioceptive adaptation. Furthermore, a high degree of visual attention to one's body may plausibly facilitate a change of proprioceptive experience.

to see how objects are located in space: I could ground visual experience in the movement or direction of my head, or in the general direction of my view. I therefore looked around the room attentively. This is when I noticed, after about 123 hours of wearing the glasses, that objects finally and definitely appeared visually where they actually were. A few moments later, when I had not moved my head for a while, I fell back in the other way of experiencing the visual field again, so that the objects once again appeared where they actually were not. But when I continued looking around again, by slowly moving my head, I could now see objects where they were. This was no longer just the case for the objects in central vision, where the experience of direction can be based on the direction of the head, but also for the objects to the left and to the right within my visual field.

In the room there were two sculptures of birds, and I could now see that their beaks were pointing to the upper left. I could follow this direction with my head and gaze, and I could accurately indicate the direction with my hand. Even when I let my gaze rest on an object for a while, avoiding head movements, I could now see what the left side was and what the right side. Let me stress that this was not merely the deliberate judgment of left and right – I could do that from the onset of the experiment – but the location as visually experienced. I could unthinkingly anticipate how movements would change my experience. The experience of the orientation of objects within my visual field was now in accordance with the experience of the general direction of my visual field during head movements; the inconsistency between the experience of the direction of my view and the apparent orientation within my visual field was gone. A good way to describe this is by saying that the seen objects were now included in the larger stream of vision: the appearance of the position of objects within my visual field fitted to the larger dynamics of visual experience during head movements.

In my case, this perceptual breakthrough was less spectacular than I had anticipated on basis of the description of Taylor (1962) mentioned above. No curious experiences were apparent as of the simultaneous perception of objects and their ghost-like mirror image. In fact I had been more excited by the recovery of visual stability than from the distinctive change in the experience of left/right orientation. Only a few days earlier I had been doubtful about my chances at success. My behavior remained clumsy but I could cope, and I had feared that I would merely acquire competence in making do with visual impairment. But at the 30th day, when I could see where objects were, this only seemed perfectly natural. It is significant that I could still switch to seeing the scene as inverted, in a way similar to the switching between different ways of seeing a Necker cube. The resulting different ways of seeing the same scene may be described as a different

perceptual interpretations of the scene, but, like in the Gestalt-switches of a Necker cube, this was nothing like a pictorial inversion. Even if I would not move my eyes or head, the two alternative ways of experiencing the scene were subjectively different.

With some effort I could still feel to which side I moved my eyes, but the visual significance of the eye movement in relation to the environment had changed. Rightward movements of the eyes no longer gave the false impression that my view traced a rightward path through the environment: I now used them to look towards the left. Although my experience now appeared surprisingly natural, this did not mean that my visual experience of the location and orientation of objects was now the same as my experience without inverting glasses. One way in which I could contrast the novel experience with the experience without inverting glasses was by considering the eye movements involved in tracing a rightwards path through the scene. But also without considering movement or the felt direction of my eyes, my experience was clearly different from visual experience without inverting glasses. Without head movements, initially experience is transformed by donning inverting glasses in the same way as by a left/right inversion of the environment. After adaptation to inverting glasses, I did not experience a pictorial flipping back of the scene. As said, my sudden change of experience was more like a Gestalt switch which led to a different way of experiencing the world. The fact that I did not experience an inversion of the environment or a pictorial inversion of experience shows that there remained crucial differences between normal vision and visual experience after adaptation with inverting glasses.

To understand my experience at this point it will be useful to consider the new contingencies to which I had become accustomed (see Figure 4). With inverting glasses, an eye movement towards the left traced a path through the environment that could be continued by moving my head to the right. An object that in reality stood on the right side of central vision could thus be brought in central view either by moving my eyes to the left, or by moving my head to the right. Without inverting glasses, of course, an object standing to the right of the focus point can be brought into central vision by moving either my eyes or my head to the right. Suppose I would be drawing on my normal visual skills, acquired during years of looking without inverting glasses. Following a line by moving my eyes to the right would then appear indicative of a trajectory towards the right. If I would be implicitly relying on these normal regularities while wearing inverting glasses, the result would be a different way to experience the same visual scene. It seems that this is exactly what happened when I occasionally fell back to the pre-adaptation-way of seeing the scene.

When first wearing inverting glasses one can correctly judge the direction of objects by saying, when one's head is directed to the front, that an object is at the right when it visually appears to be at the left. In that case, one deliberately infers the correct direction. During adaptation, non-inferential perceptual judgments become correct again. As Gibson puts it, adaptation to inverting glasses can be described as the veridicalizing of perception (Gibson 1964). Now it is true that one becomes used to veridicalized experience. But it is more accurate to describe my transformed experiences as the accurate or veridicalized visual experience of spatial location, rather than as the recovery of normal vision.[3] The fact that, without head movements, veridicalized vision with inverting glasses was not like a pictorial inversion of vision with inverting glasses before veridicalization, testifies to experiential differences between normal visual experience and veridicalized vision with inverting glasses. In other words, I found that an accurate visual experience of location can differ phenomenally from the normal visual experience of location. This vindicates the distinction between the apparent object of perception and the phenomenal character of perceptual experience, discussed in Chapter 4 of this thesis.

6. Mental imagery

On the 30th day, when I could see objects at their true location, my experience was like a bi-stable percept such as the experience of the Necker cube, or of the famous picture that can be seen as a duck or as a rabbit, in that I could see the scene in different ways. I could deliberately imagine what the effect of head movements would have been, had I not worn inverting glasses – for example, I could imagine which parts of the world would then have come into view by turning my head to the right. When I imagined such effects, objects that were on my left visually appeared as if they were on my right. My subsequent findings can best be understood against this background.

---

[3] See also the report of Hubert Dolezal, who writes: "If the question, 'Does anything re-reverse?' means 'Are the new appearances *indistinguishable* by any criterion from the remembered appearances of pre-spectacle days?' then the answer is an unequivocal 'no'." (Dolezal 1982, p. 228). Kohler quotes a subject saying that "the picture remains the same, but it is experienced differently" (Kohler 1964b, p. 155), and he even reports a comparison with a particular multistable picture (Schröder's stair illusion) (Kohler 1964a, p. 33). This confirms experiential differences between normal vision and vision after adaptation to inverting glasses.

By considering the sculptures of birds mentioned above, whose beaks were pointing to the upper left, I discovered that memory or mental imagery played tricks on me. When I looked at these sculptures for a while and then closed my eyes, I could as it were retain the sight for a few seconds: I could vividly imagine what it would be like to see the objects. Keeping my eyes closed, I then indicated with my hand the direction of their beaks as I experienced them. To my surprise, I consistently pointed in the wrong direction. There was nothing wrong with my bodily feeling of left or right, but somehow my visual memory failed in an unexpected way.

The next morning I repeated the test. Again I pointed in the wrong direction when I based my hand direction on visual memory. But I then quietly kept moving my head, while imagining what the sight of the objects would be like. With eyes closed, I had no trouble to vividly imagine the sight of the scene that I would have encountered if I had kept my eyes open, with my inverting glasses on. When I now indicated the direction of the beaks as I imagined them, the direction of my hand was in accordance with their real direction. I tried the same for other objects, with the same results.

Next I took the test one step further: no longer moving my head, but merely imagining how my experience would vary with movements of my head. Even this way I could retain the proper left/right orientation: by indicating with my hand the direction of objects as visually remembered, I consistently made correct judgments.

I believe these findings tell us something important about the basis of mental imagery. Note that, when I still pointed in the wrong direction based on mental imagery, the incorrect direction of my hand was consistent with the pre-experimental relation between eye movement and the spatial position of objects: I had to move my eyes to the upper right to follow the birds' beaks pointing to the left. Thus my findings suggest that my memory skills were still drawing on the normal significance of eye movements, rather than on the new significance brought about by wearing inverting glasses.

On reflection, this should not be surprising. After all, the last 30 days of training I had focused on looking rather than on memorizing, and now I was testing my visual memory rather than my visual skills. Moreover, I now imagined the statues as I had seen them without moving my head, so that eye movements became particularly relevant. On the assumption that visual imagery is embodied in a way that is closely related to the processes underlying visual experience, we may even have expected the memory failure. After all, imagery in which imagined eye movements in a lower

left/upper right direction would trace the direction of the birds' beaks would normally conform with beaks in a lower left/upper right direction, rather than to beaks in a lower right/upper left direction.[4]

My subsequent findings support the assumption that the processes underlying visual memory or mental imagery are closely akin to the processes underlying perceptual experience. It is clear that mental imagery was under influence of real or imagined bodily movements, as evidenced by my report of imagery with the orientation of my hand – we may say that short-term visual memory 'veridicalized'. Imagery or visual memory has its perceptual counterpart (to a certain extent it is to the subject as if he or she perceives), and as described above, this counterpart is not just a matter of sensory stimulation, but rather of sensorimotor engagement with the environment. My findings indicate that sensorimotor differences in the perceptual case are reflected in differences in the case of mental imagery or visual memory. Thus it seems that the embodiment of imagery is closely akin to – and perhaps overlapping with – the embodiment of perceptual engagement with the environment.

I have not investigated memory effects during a longer period. But it would be interesting to further test the ways in which visual memory is grounded in perceptual interaction. For example, when a movie is seen on a large screen with inverting glasses with a restricted scope, so that movement of the head plays a serious role in watching the movie, will memory be based in head movements? If so, one would expect that memory of orientation will be more likely to be correct in cases where head movements play a crucial role, compared to cases where only eye movements are involved, for the relation between head movements and the field of view remains unaltered by inverting glasses. Further research in these matters may throw light on the way in which certain forms of memory are embodied.

7. Conclusion

Experience with inverting glasses reveals key factors in spatial vision. Above I have reported my findings on visual stability, the experience of the

[4] I suppose that there is an analogue in imagery for eye movements, which might be called 'imagined eye movement'. This could be interpreted as shifts of attention or 'mental scanning', in which 'subjects covertly go through the motions of such scanning' (Thomas 1999). By gently touching your closed eyelids while imagining looking around at a familiar scene, you may be able to ascertain that imagery sometimes comes with overt eye movements as well.

location of objects, and visual imagery. We have seen how a description of visual experience needs to refer to sensorimotor patterns rather than to patterns of sensory stimulation only, and I argued that such a description does not reduce to a description in terms of the possession of visual images. Also the spatial phenomenology of mental imagery may best be described in terms of apparent sensorimotor patterns.

Previous studies have raised the question whether visual experience may turn back to normal after adaptation to inverting glasses. Stratton has reported the return of 'upright vision' after several days of wearing inverting glasses, asserting that upright vision is a matter of the harmony between sight and touch (Stratton 1897). Still his report suggests that differences remain in the phenomenology of vision with inverting glasses, even if the harmony between sight and touch is restored: memory of normal upright vision appears to testify to a difference with experience of 'upright' vision after adaptation to inverting glasses. By contrast, reports of Taylor (1962) and Kohler (1964b) seemed to suggest that visual experience itself turns back to normal: after a stage with double vision, the re-inverted part of the visual experience remains. If we were to conceive of vision in terms of the possession of inverted or non-inverted 'visual images', there would appear to be a conflict between these reports.

Above I offered a more fine-grained analysis of my experience with inverting glasses. As the breakdown of visual stability shows, visual experience cannot be fully described in terms of 'images'. Visual stability requires a visual 'grasp' of the scene that goes beyond being exposed to sweeping images, for in the case of sweeping images one may still be able to track parts of the image with one's eyes. In my description of visual experience I stressed the patterns of interaction with the environment. While some of these patterns remain unaltered by inverting glasses (e.g. to look to the right one has to turn one's head to the right), others are systematically altered (e.g. to look to the right one has to turn one's eyes to the left). As a result, inverting glasses introduce a conflict at the very heart of spatial vision. Before full adaptation to inverting glasses, judgments of visual direction grounded in head movements differ from judgments grounded in eye movements. The effect of wearing inverting glasses therefore cannot be analyzed as an inversion of 'visual images'. Even before a state of full adaptation is reached, the experience of visual direction based in head movements is often correct.

As described above, perceptual adaptation to inverting glasses does not cancel out the differences with normal vision. In static cases, the initial effect of inverting glasses could be described as an inversion within the field of vision. But when accurate experience of the position of objects is restored, the resulting experience can neither be described as the full re-

inversion of experience nor as the remaining of inverted vision. Differences remain between 'upright' vision with inverting glasses and normal upright vision, as suggested by Stratton's report. But it would be a mistake to suppose that these differences consist in the visual experience remaining fully 'inverted'. First of all, visual experience was never fully inverted in the first place. As said, only some patterns of visual experience are transformed, and experience during active exploration of the environment cannot be fully described in terms of (invertible) images. Second, there was a clear difference between experience with inverting glasses before and after perceptual adaptation, also in cases without head movements. This difference is comparable to the different ways in which one may experience bi-stable stimuli such as the duck-rabbit picture. The fact that no image-inversion was apparent indicates that adaptation to inverting glasses did not fully counteract the change of visual experience brought about by inverting glasses.

At the same time, there are ways in which visual phenomenology with inverting glasses becomes normal again, as suggested by the reports of Taylor and Kohler. As I described, the visual field becomes once again integrated in the larger stream of vision; the conflict between the visual consequences of eye movements and head movements resolves. In the meantime, veridical perception of the position of objects is re-acquired. However, this does not imply that there is no difference between normal experience and 'veridicalized' experience with inverting glasses. That there are such differences is evidenced by the different ways in which one may see the scene after reaching perceptual adaptation. Even if the object-oriented judgments are the same (as also found by Taylor and Kohler), my findings indicate that the experiences differ. This vindicates the distinction between on the one hand regarding perception from an object-oriented or perceptual knowledge-oriented stance, and on the other hand regarding perception from a perceiver-oriented or mode of engagement-oriented phenomenal stance, as discussed in Chapter 4 above.

A difference remains between my experience and the experience reported by Taylor and Kohler. It seems that two ways of seeing the scene could be *simultaneously* grasped by Taylor's subject, who saw a chair both in its actual place and on the opposite side (Taylor 1962, p. 202). The same goes for Kohler's subject, who apparently had the experience of seeing two lights when only one was presented (Kohler 1964b, p. 161). In my case the different ways of seeing were rivals: as if I saw the scene either as drawing on my newly acquired expectancies, or according to the old expectancies, but never both at the same time. Perhaps more extensive use of the glasses could yield different results.

I further found that visual adaptation need not coincide with adaptation of visual memory or mental imagery. Under certain circumstances, visual memory or imagery of objects tends to be grounded in eye movements, according to the significance they have in the absence of inverting glasses. By deliberately giving a larger role to (real or imagined) movements of the head also short-term visual memory veridicalized, and the scene was once again imagined to be as it was actually experienced. This supports the view that the processes underlying mental imagery or visual memory are at the very least related to the processes underlying perceptual engagement with the environment, and it suggests that they may be partly overlapping. The role of real or imagined movements in visual imagery, and the fact that imagery can adapt under influence of wearing inverting glasses, provide strong indications that imagery is grounded in actual sensorimotor engagement with the environment.

Let me close this undertaking in experimental phenomenology by recalling Stratton's (1897) report, where he pointed out that the inverted position of the retinal image is not essential to 'upright vision', for the whole system of visual objects can never by itself be either inverted or upright. To this we might now add that there is more to visual phenomenology than the harmony between touch and sight. As long as we remain sensitive to eye movements and head movements, we can differentiate between normal upright vision, and upright vision with inverting glasses. I would suggest that this is not just because we can compare vision with tactual or motor perceptions, but because these eye movements and head movements are part of what we must reflect on when we reflect on visual phenomenology.

Summary and conclusion

In this thesis I have discussed the phenomenal character of experience. I have argued that a sensorimotor approach can shed considerable light on the matter. After a synopsis I shall briefly point towards further issues, and I shall tentatively relate the presented view to a recent discussion in the philosophy of mind concerning the location of the processes underlying consciousness.

Synopsis

How can we understand conscious experience as a genuine part of the natural world? What explains the specific *phenomenal character* of experience, i.e. what the experience is like for the person? In recent decades, attempts to understand conscious experience have often appealed almost exclusively to the brain. In the case of perception, conscious experience has then typically been thought of as an inner model of the environment in the head of the perceiver. At the same time, however, the behavioral and cognitive sciences have become more sensitive to a broader range of processes, cutting across organism/environment divides. The actual manipulation of the environment can contribute to problem-solving, for example when someone moves around pieces of a puzzle to see where they might fit. The fact that elaborate inner operations may not be required for adequate behavior has initiated a significant rethinking of the inner workings of human beings. Also the study of perception has increasingly focused on dynamic patterns of engagement with the environment. Active exploration plays a crucial role in perception, and it has become questionable whether the assumption that the brain contains elaborate inner models is helpful to the study of perception. Indeed, some have argued that if the world is out there to explore, no inner model is required for perceptual experience at all.[1] In this thesis I have explored the consequences of this perspective for understanding the phenomenal character of experience. More particularly, I explicate, develop, and defend

[1] There are two lines of argumentation against reliance on inner models. First, as a conceptual matter, we should not presuppose that perception involves inner models – no model is self-interpreting and the possession of an inner model cannot be what perception consists in (Chapter 1). Second, there are reasons to believe that the most parsimonious empirical model of perception may not involve inner models (Chapter 1, 2).

a *sensorimotor approach*, as outlined in O'Regan and Noë (2001). I argue that this approach can help to understand the phenomenal character of perceptual experience by focusing on the perceiver's sensorimotor engagement with the environment.

As already recognized by Helmholtz (1876) and others, perceptual experience cannot generally be understood in terms of sensory stimulation alone. Perceptual experiences depend on the way in which sensory stimulation relates to actual or potential motor action; they depend on the implicit grasp of *sensorimotor dependencies*. For example, to see that an object stands before its background is to grasp the sensory consequences that can be expected from a bodily movement; it is to implicitly know how bodily movements allow you to look behind the object. Similarly, the visual experience of a movement in the environment involves changes of sensory stimulation in absence of eye movements, while it involves a lack of change of sensory stimulation coming from the moving object if one is tracking the object with one's eyes: in either case experience relies on sensorimotor dependencies. The key question is how to conceive of them. In particular, the question is whether we should conceive of sensorimotor dependencies as the basis for the construction of inner models of the world, or whether such models are not needed. The sensorimotor account proposes that we can do without inner models, and that we should think of experience as a mode of skillful perceptual engagement with the environment.

An explanatory account of the phenomenal character of experience, I have stressed, must do two things (Chapter 1). First, it must provide an accurate description of the phenomenal character of experience, at the level of the person. Second, it must be able to link this description to a description of the underlying processes, at the subpersonal level of description. If the first requirement were not fulfilled we would at best have a list of the processes that correlate with an experience: we may then claim that somehow these processes should explain the character of experience, but until we understand how these processes can be considered descriptive of a particular experience, such an explanation has not yet been given. If the second requirement is not fulfilled, we at best have a narrative of experiences: 'first I experienced this, and then I experienced that'. To make experience fit in with our view of the natural world, this narrative must be connected to non-phenomenal descriptions, and in particular with descriptions of the processes underlying the experience. A *dual currency* account does this by providing a description that fits both experience and its underlying processes.

I have argued that a sensorimotor account has several advantages over accounts relying on inner models. The following three considerations in particular count in favor of the sensorimotor account. *First*, the

sensorimotor account enables us to articulate what experiences are like for the person. For example, the differences between visual experience and tactile experience can be explicated in terms of differences in the characteristic sensorimotor patterns (Chapter 2). Furthermore, descriptions in terms of visual images fail to capture crucial aspects of visual experience, and we must appeal to sensorimotor patterns if we are to provide an accurate description of the phenomenal character of spatial vision (Chapter 6). *Second*, a sensorimotor account simplifies our view of the subpersonal processes involved in perceptual experience. A grasp of sensorimotor dependencies need not be complemented by an additional construction of inner models of the environment. As a result, a whole layer can be eliminated from the subpersonal account of perceptual experience as compared with inner model accounts (Chapter 2). We are then one step closer to understanding the phenomenal character of experience than inner model accounts suppose. *Third*, the sensorimotor account can help to link personal and subpersonal levels of description of perceptual experience. Such an account provides a description of subpersonal processes that remains close to the way we experience the world. It thereby helps to find a match between personal and subpersonal levels of description, and thus to attain the dual currency ideal (Chapter 1, 2, 3).

In Chapter 3 I have connected the sensorimotor account to *neural workspace accounts*. While sensorimotor approaches focus on patterns of interaction cutting across the brain-environment divide, workspace theories strongly focus on the brain. The approaches are often seen as competitors, but I argue that they are complementary. At the basis of the argument lies the distinction between two explanatory issues regarding phenomenal experience (Chalmers 1996; Hurley & Noë 2003). There are the *comparative gap* issues, concerning the specific phenomenal character of experience (e.g. what explains the difference between the experience of red and blue, or between visual experience and auditory experience). Furthermore, there is the question of the *absolute gap*, concerning the fact that we have conscious experience at all. I argue that sensorimotor and workspace approaches are best seen as addressing different explanatory gaps. Sensorimotor approaches are best suited to explain the specific phenomenal character of experiences, but they do not explain why there is experience in the first place (Noë 2004). A sensorimotor characterization of the requirements for conscious experience remains too descriptive, and needs to be fleshed out in subpersonal terms. Workspace approaches, in contrast, address the neural processes that are specific for conscious experience (e.g. Dehaene & Naccache 2001). These approaches are relevant to the absolute gap because they may explain the contrast between conscious experience and its absence for a conscious perceiver. But these

approaches are less suited for comparative gap concerns. I discuss three scenarios for combining the approaches (more on this in the concluding thoughts below), and I argue that on all three scenarios advantages are to be expected from the combination.

In Chapter 4 I point out a distinctive contribution of the sensorimotor account by arguing that its focus on the phenomenal character of experience offers a crucial addition to accounts focused on what perception may tell us about the world. Starting-point is the schism between 'qualophobic' and 'qualophilic' tendencies in the philosophy of mind. While qualophobes argue that conscious experience requires nothing over and above the processes that play a role in our capacities to act and to have access to the world (Dennett 1991), qualophiles insist that there is more to explain than cognitive functions or discriminatory abilities (Levine 1994; Block 1996). I propose to follow the qualophobes in rejecting special qualitative ingredients within our account of consciousness, while accepting the qualophilic idea that different questions can be asked regarding phenomenal character and cognitive access. These positions can be reconciled by building on two perspectives we can take towards conscious perceivers: the *intentional stance* and the *phenomenal stance*. When we apply Dennett's (1987) intentional stance to a perceiver, we take an interest in a perceiver's potential perceptual knowledge, focusing on states that can be evaluated as true or false. From the phenomenal stance, in contrast, what matters is what the experience is like for the perceiver (cf. Robbins & Jack 2006). As a result of its focus on epistemically evaluable content, the intentional stance abstracts away from the processes through which we acquire perceptual knowledge. I argue that from the phenomenal stance experience should be construed as a perceiver-centered notion, which concerns the way in which the perceiver is perceptually engaged with his or her environment. The sensorimotor approach offers a natural way to flesh out such a notion of experience.

Chapter 5 draws consequences of the skill-oriented sensorimotor account for the explanation of color vision. Traditional approaches have capitalized on a dichotomy between inner and outer factors in the explanation of experience. While some have focused on the environmental properties that are thought to be tracked by color vision, others have insisted that properties of our perceptual systems determine crucial aspects of experience. An aspect of color vision that is often brought in connection with neural factors is that some 'unique' colors appear as 'pure', or containing no trace of any other color (red, green, yellow and blue), while others can be considered as a mixture of these colors, or as 'binary colors'. According to a widespread assumption, this unique/binary structure of color experience is to be explained in terms of

neurophysiological structuring (e.g. by opponent processes) and it has no objective basis in the physical stimulus. In this chapter I challenge this assumption, and an argument building on it, namely the *argument from structure*, which concludes that color experiences are neural properties lacking a proper environmental basis (Hardin 1988). I point out that, according to present vision science, the activity of known neurophysiological opponent processes does not in fact correlate with the experience of unique colors. Moreover, a logical point must be stressed: even if a neural correlate of the unique/binary structure of color experience is to be found, this does not in itself support the claim that the unique/binary structure derives exclusively from neurophysiological factors. Indeed, a recent analysis of Philipona and O'Regan (2006) has revealed environment-involving patterns which may be at the basis of the unique/binary structure of color experience. In particular, it turns out that under changing lighting conditions, surfaces that appear as having unique colors result in more constrained patterns of retinal stimulation as compared to surfaces with binary colors. This suggests that if the structure of neurophysiology matches the unique/binary structure, this matching derives from the patterns of sensory stimulation that arise from the perceiver's interaction with the environment. I conclude that purely neurophysiological accounts have little explanatory force. Instead of focusing exclusively on neural or environmental correlates, explanations of the structure of experience should be sensitive to the patterns of perceptual engagement with the environment.

In Chapter 6, the sensorimotor approach is applied and vindicated in a phenomenological case-study. Drawing on my experience with wearing *left/right inverting glasses*, I show how a sensorimotor analysis helps to characterize visual experience, and I report findings on mental imagery. When first wearing inverting glasses, movement of the head leads to a breakdown of visual stability: the world appears to sweep in front of one's eyes. The metaphor of 'visual images' is inadequate to describe this experience, for the experience is importantly different from the experience of watching moving images. Indeed, the lack of visual stability is better described as a failure to track objects with one's gaze, an ability that is usually in place when watching moving images. I further point out that inverting glasses introduce a previously underappreciated conflict at the heart of spatial vision. Although the glasses change the relation between eye movements and the distal stimuli, they leave unaltered how head movements change the direction of view. As a result, I was often aware of the general direction of my view, despite the inversion within my visual field. Again, the resulting experience cannot be adequately characterized in terms of (inverted) visual images. This conclusion is reinforced by my

subsequent finding on perceptual adaptation. During adaptation to inverting glasses, visual stability is regained, and even the skill to see where objects are located can be re-acquired (as also reported for example by Stratton 1897; Taylor 1962; Kohler 1964). Crucially, I found that perceptual adaptation came with an experience comparable to Gestalt-switches (as in the different ways of experiencing the famous picture that can be seen either as a duck or as a rabbit): I could experience the environment in different ways, without experiencing a pictorial inversion. My experience with wearing inverting glasses thereby vindicates the distinction between an object-oriented stance and the phenomenal stance, for the same perceptual judgments (with and without inverting glasses) could come with markedly different experiences. Finally, I show that visual imagery depends on real or imagined bodily movements, supporting the idea that also imagery is grounded in sensorimotor engagement with the environment.

Further issues

In this thesis I have emphasized explanatory advantages of a sensorimotor approach for understanding the phenomenal character of experience. It must be noted that my explorations have mainly concerned the comparative advantages of the general framework, and no elaborate attempt has been made to provide detailed sensorimotor analyses of specific experiences (Chapter 6 is a modest exception here). Furthermore, as a matter of strategic choice, this thesis has been concerned with basic aspects of perceptual experience. By focusing on what is distinctive for a given experience, we have deliberately set aside issues concerning for example individual aesthetic appreciations, cognitive associations, or affective aspects of experiences. Clearly there is more to say about aspects of the phenomenal character of perceptual experience that have not been included in the analysis. I shall not attempt to sketch a more encompassing account here. The point I wish to stress, however, is that the sensorimotor account offers a relatively well-developed framework from which further issues may be approached. If a sensorimotor account is along the right lines for basic aspects of perceptual experience, our understanding of other aspects had better fit in with this account.

At the same time I think we should acknowledge that in important respects also the general framework of the sensorimotor approach requires further development. The main focus has been on sensorimotor patterns that are *distinctive* for specific perceptual experiences, such as patterns that differentiate between visual experiences and auditory experiences. This emphasis on differentiation has been at the expense of an exploration of the integration of the different perceptual modalities. For example, colors may

look cool or warm, and sounds may be described as round or sharp. It seems that the 'warmth' of a color is an aspect of color experience that requires explanation, and it seems plausible to bring this into relation with our sense of temperature. While the precise nature of this relation remains to be established, it is clear that there can be an affinity between experiences, as evidenced by the fact that it seems natural to describe certain colors as 'cool' or 'warm'.[2] From the perspective of a sensorimotor account, we could then speculate that different sensorimotor skills may intertwine and that this is reflected at the level of experience. Alternatively, there may at some level of description be a similarity between the sensorimotor dependencies characteristic of the different experiences. Be this as it may, a full account of the phenomenal character of experience should eventually be able both to explain the differences and the affinities between experiences.

So far we have been concerned with the question what it is that we reflect on, when we reflect on the phenomenal character of perceptual experience. This naturally raises the further question what it is to reflect. How should thought – including thought about experience – be understood? Since thought can have a phenomenal side this question is also relevant to the understanding of the phenomenal character of experience. Without being a native speaker of English, I sometimes think in English. It seems clear that this phenomenon should be understood in the light of actual engagement with an environment filled with language. Thus the question becomes whether or to what extent we can 'scale up' the sensorimotor approach to include the perception of language – not just sounds but meaningful sentences – and to capture our thought-infused engagement with the environment. Now of course the point here is not to anticipate in any detail what such an account would look like. The point is that further exploration of these matters could increase our understanding of human experience. Again, I would stress the relevance of the contrast between skillful engagement and inner models here. If our basic perceptual engagement can best be understood without proposing inner models, the same may be true for our linguistic engagement with the world.

---

[2] A striking case of affinity between experiences is the fact that sounds can be strongly associated with shapes, as in the 'bouba/kiki effect', in which subjects, when asked to apply these words to round or pointy shapes, tended to make the same association (Ramachandran & Hubbard 2001). For a discussion of the perhaps related phenomenon of synesthesia in relation to a sensorimotor perspective, see Hurley and Noë (2003).

Zooming out, zooming in

The sensorimotor account, it has been argued, gains crucial explanatory advantages from expanding our perspective to include interactions with the environment. Let me conclude by tentatively relating this claim to a recent discussion on the location of the processes that constitute conscious experience. There is certainly no lack of correlations between neural activity and conscious experience. One way to approach the explanatory challenge of conscious perceptual experience is by considering some neural activity that is correlating with experience, asking what it is that explains this correlation. Sensorimotor theorists have argued (as also quoted in Chapter 1 above):

> "To find explanations of the qualitative character of experience, our gaze should be extended outward, to the dynamic relations between brain, body, and world." (Hurley & Noë 2003, p 132)

A reason for adopting this extended gaze is this. The extended dynamics of sensorimotor engagement with the environment can be considered descriptive of the phenomenal character of experience. Arguably, it is by participating in this dynamics that the local neural activity contributes to the phenomenal character of experience. If, as I have argued, perceptual experience is a matter of characteristic sensorimotor dependencies, and given that these sensorimotor dependencies presumably can best be described and explained in environment-involving terms, it would follow that perceptual experience can best be described and explained in environment-involving terms.

In recent years, there has been a discussion in the philosophy of mind concerning the location of the processes we must study if we are to understand mental phenomena in general and conscious experience in particular. Some have argued that there is no principled reason to presuppose that the processes that constitute and explain experience are restricted to the organism or the brain, defending the view that not all the processes underlying consciousness are contained within the perceiver (e.g. Hurley 1998; 2010; Rowlands 2003). Others have defended the widespread view that the processes underlying consciousness do reside in the brain (e.g. Clark 2009).

Now if we were to think of conscious perceptual experience in terms of the possession of a model, this issue of localization would be of vital interest to the explanation of experience. For if this model were located partly outside the head no scrutinizing of inner processes could ever allow us to characterize experience. It would then be crucial to find the location of the physical processes that embody the model, or as it is often put, of the

'vehicles' of the 'content' of conscious experience. If we think of experience as a model of the world, we would like to know where this model is to be found.

From a skill-oriented perspective, however, the import of the issue concerning the internal versus the external location of the processes underlying consciousness is not so evident. It may be a clear question which subpersonal processes would constitute a model, but it is certainly not so clear what we mean by asking which subpersonal processes constitute the exercise of a perceptual skill. Does the exercising of a perceptual skill consist in the extended process of causal interaction of brain, body and world? Or could we somehow conceive of perceptual engagement in terms of the internal activity of the organism or the brain? Alternatively, we could reject these views and insist that strictly speaking the skill only exists at the personal level, although of course enabled by the subpersonal processes. Chapter 3 above pointed towards three scenarios along these different lines for the further development of a sensorimotor account (in relation to workspace accounts of neural dynamics). But perhaps these questions about the constitution of experience have no more a real answer than similar constitution-questions about behavioral skills.

Consider the case of cycling. Given that we do not think of cycling as a matter of having local models or representations accompanying the activity, we would not ask whether or not cycling 'resides in the brain'. We have no doubt that it is possible to give a characterization of the subpersonal processes involved, and such a characterization can presumably be aligned with a personal level description of the exercise of cycling skills. But while a model or representation would be a strictly localized process, it seems that we lack a basis for claiming that exercising a skill resides in some definite set of subpersonal processes. In other words, from a skill-oriented perspective it is not so evident how we could distinguish between processes constituting the subpersonal basis of a high-level phenomenon and processes that are 'merely' causally involved. But this should not be a problem. After all, for explanatory purposes the key issue is not where to draw any causal/constitutive distinction. It is to provide accurate descriptions of behavior and experience and to relate these descriptions to subpersonal processes.

The fundamental contribution of the sensorimotor approach as defended in this thesis is certainly not that it would promote a shift from an inner model-framework to an extended model-framework. It does no such thing. The contribution is that it shifts away from the model-based framework altogether, to a framework oriented towards skillful perceptual engagement. I have shown in some detail how such a framework can be fleshed out in a sensorimotor account, to explain crucial aspects of the

phenomenal character of perceptual experience. If this account is right, we can increase our understanding of experience by zooming out from the local neural activity correlating with experience. To explain the phenomenal character of experience we should zoom in on the perceptual skills in which these neural processes participate.

References

Austin, John L. (1962) *Sense and Sensibilia* (Oxford: Oxford University Press).

Auvray, Malika; Hanneton, Sylvain; O'Regan, J. Kevin (2007) 'Learning to perceive with a visuo-auditory substitution system: localisation and object recognition with 'The vOICe'', *Perception*, 36, pp. 416-430.

Auvray, Malika; Myin, Erik (2009) 'Perception with compensatory devices: from sensory substitution to sensorimotor extension', *Cognitive Science*, 33 (6), pp. 1036-1058.

Baars, Bernard J. (1988) *A Cognitive Theory of Consciousness* (Cambridge: Cambridge University Press).

Baars, Bernard J. (2002) 'The conscious access hypotheses: origin and recent evidence', *Trends in Cognitive Sciences,* 6, pp. 47-52.

Bach-y-Rita, Paul (1984) 'The relationship between motor processes and cognition in tactile vision substitution', in: Prinz, W., Sanders, A.F. (eds.) *Cognition and Motor Processes,* pp. 149-160 (Berlin: Springer-Verlag).

Bach-y-Rita, Paul (2002) 'Sensory substitution and qualia', in: Noë, A., Thompson, E. (eds.) *Vision and Mind: Selected Readings in the Philosophy of Perception*, pp. 419-436 (Cambridge, MA: MIT Press).

Bennett, Maxwell R.; Hacker, Peter M.S. (2003) *Philosophical Foundations of Neuroscience* (Oxford: Blackwell Publishing).

Bentham, Jeremy (1823/1907) (second edn.) *An Introduction to the Principles of Morals and Legislation* (Oxford: Clarendon Press).

Berlin, B.; Kay, P. (1969) *Basic Color Terms: Their Universality and Evolution* (Berkeley: University of California Press).

Block, Ned (1996) 'Mental paint and mental latex' *Philosophical Issues*, 7, pp. 19-49.

Block, Ned (2007) 'Consciousness, accessibility, and the mesh between psychology and neuroscience', *Behavioral and Brain Sciences*, 30, pp. 481-548.

Bompas, Aline; O'Regan, J. Kevin (2006a) 'Evidence for a role of action in colour perception', *Perception*, 35, pp. 65-78.

Bompas, Aline; O'Regan, J. Kevin (2006b) 'More evidence for sensorimotor adaptation in color perception', *Journal of Vision*, 6, pp. 145-153.

Bradley, P.; Tye, M. (2001) 'Of Colors, kestrels, caterpillars, and leaves', *The Journal of Philosophy*, 98(9), pp. 469-487.

Bridgeman, Bruce (2007) 'Efference copy and its limitations', *Computers in Biology and Medicine*, 37, pp. 924-929.

Bridgeman, Bruce; van der Heijden, A.H.C; Velichkovsky, Boris M. (1994) 'A theory of visual stability across saccadic eye movements', *Behavioral and Brain Sciences*, 17 (2), pp. 247-292.

Broackes, Justin (2011) 'Where do the unique hues come from?', *Review of Philosophy and Psychology*, 2, pp. 601-628.

Brooks, Rodney A. (1991) 'Intelligence without representation', *Artificial Intelligence*, 47, pp. 139–159.

Brouwer, G.J.; Heeger, D.J. (2009) 'Decoding and reconstructing color from responses in human visual cortex', *The Journal of Neuroscience*, 29, pp. 13992-14003.

Buchsbaum, G.; Gottschalk, A. (1983) 'Trichromaticy, opponent colours coding and optimum colour information transmission in the retina', *Proceedings of the Royal Society of London, Series B*, 220, pp. 89-113.

Byrne, Alex (2001) 'Intentionalism defended', *The Philosophical Review*, 110, pp. 199-240.

Byrne, A.; Hilbert, D.R. (2003) 'Color realism and color science', *Behavioral & Brain Sciences*, 26(1), pp. 3-64.

Calvo, Paco; Gomila, Toni (eds.) (2008) *Handbook of Cognitive Science: An Embodied Approach* (San Diego: Elsevier).

Chalmers, David J. (1995) 'Facing up to the problem of consciousness', *Journal of Consciousness Studies,* 2, pp. 200-219.

Chalmers, David J. (1996) *The Conscious Mind* (Oxford: Oxford University Press).

Chemero, Anthony (2009) *Radical Embodied Cognitive Science* (Cambridge, MA: MIT Press).

Churchland, Paul (2005) 'Chimerical colors: some phenomenological predictions from cognitive neuroscience', *Philosophical Psychology*, 18(5), pp. 527-560.

Churchland, Paul (2007) 'On the reality (and diversity) of objective colors', *Philosophy of Science*, 74(2), pp. 119-149.

Clark, Andy (1997) *Being There: Putting Brain, Body, and World Together Again* (Cambridge, MA: MIT Press).

Clark, Andy (1999) 'An embodied cognitive science?', *Trends in Cognitive Sciences,* 3, pp. 345-351.

Clark, Andy (2008) *Supersizing the Mind* (Oxford: Oxford University Press).

Clark, Andy (2009) 'Spreading the joy? Why the machinery of consciousness is (probably) still in the head', *Mind*, 118, pp. 963-993.

Clark, Andy; Chalmers, David J. (1998) 'The extended mind', *Analysis,* 58, pp. 7-19.

Clark, Austen (1993) *Sensory Qualities* (Oxford: Oxford University Press).

Cohen, Michael A.; Dennett, Daniel C. (2011) 'Consciousness cannot be separated from function', *Trends in Cognitive Sciences*, 15 (8), pp. 358-364.

Cooke, Ed; Myin, Erik (2011) 'Is trilled smell possible? How the structure of olfaction determines the phenomenology of smell' *Journal of Consciousness Studies*, 18 (11-12), pp. 59-95.

Crapse, Trinity B., Sommer, Marc A. (2008) 'Corollary discharge across the animal kingdom', *Nature Reviews Neuroscience*, 9, pp. 587-600.

Damasio, Antonio (2010) *Self Comes to Mind: Constructing the Conscious Brain* (New York: Pantheon Books).

De Valois, R.L.; Abramov, I.; Jacobs, G.H. (1966) 'Analysis of response patterns of LGN cells', *Journal of the Optical Society of America*, 56 (7), pp. 966-977.

De Valois, R.L.; De Valois, K.K. (1993) 'A multi-stage color model', *Vision Research*, 33(8), pp. 1053-1065.

Degenaar, Jan (*forthcoming*) 'Perception from the phenomenal stance', *Logique et Analyse*.

Degenaar, Jan; Keijzer, Fred (2009) 'Workspace and sensorimotor theories: complementary approaches to experience', *Journal of Consciousness Studies*, 16 (9), pp. 77-102.

Degenaar, Jan; Myin, Erik (*submitted*) 'The structure of color experience and the existence of surface colors'.

Dehaene, Stanislas; Naccache, Lionel (2001) 'Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework', *Cognition,* 79, pp. 1-37.

Dennett, Daniel C. (1969) *Content and Consciousness* (London: Routledge).

Dennett, Daniel C. (1978) 'Mechanism and responsibility', in: *Brainstorms* (Sussex: Harvester Press).

Dennett, Daniel C. (1987) *The Intentional Stance* (Cambridge, MA: MIT Press).

Dennett, Daniel C. (1991) *Consciousness Explained* (New York: Little, Brown and Company).

Dennett, Daniel C. (1991b) 'Real patterns', *The Journal of Philosophy*, 88, pp. 27-51.

Dennett, Daniel C. (1995) 'Animal consciousness: what matters and why', *Social Research*, 62 (3), pp. 691-710.

Dennett, Daniel C. (2001) 'Are we explaining consciousness yet?', *Cognition,* 79, pp. 221-237.

Dennett, Daniel C. (2003) 'Who's on first? Heterophenomenology explained', *Journal of Consciousness Studies*, 10, pp. 19-30.

Dennett, Daniel C.; Kinsbourne, Marcel (1992) 'Time and the observer: the where and when of consciousness in the brain', *Behavioral and Brain Sciences*, 15 (2), pp. 183-247.

Derrington, A.M.; Krauskopf, J.; Lennie, P. (1984) 'Chromatic mechanisms in lateral geniculate nucleus of macaque', *Journal of Physiology*, 357, pp. 241-265.

Dolezal, Hubert (1982) *Living in a World Transformed: Perceptual and Performatory Adaptation to Visual Distortion* (New York: Academic Press).

Dretske, Fred (1995) *Naturalizing the Mind* (Cambridge, MA: MIT Press).

Edelman, Gerald M.; Tononi, Giulio (2000) *A Universe of Consciousness. How Matter Becomes Imagination* (New York: Basic Books).

Engel, Andreas K.; Fries, Pascal; Köning, Peter; Brecht, Michael; Singer, Wolf (1999) 'Temporal binding, binocular rivalry and consciousness', *Consciousness and Cognition,* 8, pp. 128-151.

Engel, Andreas K.; Singer, Wolf (2001) 'Temporal binding and the neural correlates of sensory awareness', *Trends in Cognitive Sciences,* 5, pp. 16-25.

Findlay, John Malcolm; Gilchrist, Iain D. (2003) *Active Vision: The Psychology of Looking and Seeing* (Oxford: Oxford University Press).

Gegenfurtner, K.R. (2003) 'Cortical mechanisms of colour vision', *Nature Reviews Neuroscience*, 4, pp. 563-572.

Gibson, James J. (1950) *The Perception of the Visual World* (Cambridge, MA: The Riverside Press).Gibson, James J. (1964) 'Introduction', *Psychological Issues*, 3 (4), pp. 5-13.

Gibson, James J. (1966) *The Senses Considered as Perceptual Systems* (Boston: Houghton Mifflin Company).

Gibson, James J. (1979) *The Ecological Approach to Visual Perception* (Boston: Houghton Mifflin Company).

Grush, Rick (1998) 'Skill and spatial content' *Electronic Journal of Analytic Philosophy*, 6 (6).

Grüsser, Otto-Joachim (1995) 'On the history of the ideas of efference copy and reafference', in: Debru, Claude (ed.) *Essays in the History of the Physiological Sciences*, pp. 35-55. (Amsterdam: Rodopi).

Guillery, R.W.; Sherman, S.M. (2011) 'Branched thalamic afferents: what are the messages that they relay to the cortex?' *Brain Research Reviews*, 66, pp. 205-219.

Hardin, C.L. (1988/1993) *Color for Philosophers: Unweaving the Rainbow*. Second edition (1993). (Indianapolis: Hackett Publishing Company).

Harman, Gilbert (1990) 'The intrinsic quality of experience', *Philosophical Perspectives*, 4, pp. 31-52.

Harris, Charles Samuel (1965) 'Perceptual adaptation to inverted, reversed, and displaced vision', *Psychological Review*, 72 (6), pp. 419-444.

Held, Richard; Hein, Alan (1963) 'Movement-produced stimulation in the development of visually guided behavior', *Journal of Comparative and Physiological Psychology*, 56 (5), pp. 872-876.

Helmholtz, Hermann von ([1876]/1924/1962) *Treatise on Physiological Optics*, Volume 3, Southall, J.P.C. (ed.) (New York: Dover Publications).

Helmholtz, Hermann von ([1878]/1977) 'The facts in perception (with notes and comments by Moritz Schlick)', in Cohen, Robert S.; Elkana, Yehuda. (eds.) *Hermann von Helmholtz: Epistemological Writings*, pp. 115-185. (Dordrecht: Reidel Publishing Company).

Hering, E. (1920/1964) *Outlines of a Theory of the Light Sense*. (Cambridge, MA: Harvard University Press).

Holst, Erich von, Mittelstaedt, Horst (1950) 'Das Reafferenzprinzip (Wechselwirkungen zwischen Zentralnervensystem und Peripherie)' *Naturwissenschaften*, 37, pp. 464–476.

Humphrey, Nicholas (2000) 'How to solve the mind-body problem', *Journal of Consciousness Studies*, 7, pp. 5–20.

Hurley, Susan L. (1998) *Consciousness in Action* (Cambridge, MA: Harvard University Press).

Hurley, Susan L. (2001) 'Perception and action: alternative views', *Synthese*, 129, pp. 3-40.

Hurley, Susan L. (2010) 'The varieties of externalism', in Menary, Richard (ed.) *The Extended Mind* (Cambidge, MA: MIT Press).

Hurley, Susan L.; Noë, Alva (2003) 'Neural plasticity and consciousness', *Biology and Philosophy,* 18, pp. 131-168.

Hurvich, L.M. (1981) *Color Vision*. (Sunderland: Sinauer Associates Inc).

Hutto, Daniel D. (2005) 'Knowing *what*? Radical versus conservative enactivism', *Phenomenology and the Cognitive Sciences*, 4, pp. 389-405.

Jacob, Pierre; Jeannerod, Marc (2003) *Ways of Seeing* (Oxford: Oxford University Press).

James, William (1890/1950) *The Principles of Psychology* (New York: Dover Publications).

Keijzer, Fred (2001) *Representation and Behavior* (Cambridge, MA: MIT Press).

Keijzer, Fred; Schouten, Maurice (2007) 'Embedded cognition and mental causation: setting empirical bounds on metaphysics', *Synthese,* 158, pp. 109-125.

Koch, Christof (2004) *The Quests for Consciousness: a Neurobiological Approach* (Englewood: Roberts and Company Publishers).

Kohler, Ivo (1964a) 'The formation and transformation of the perceptual world', *Psychological Issues*, 3 (4), pp. 19-133.

Kohler, Ivo (1964b) 'Rehabituation in perception', *Psychological Issues*, 3 (4), pp. 135-164

Krauskopf, J.; Williams, D.R.; Heeley, D.W. (1982) 'Cardinal directions of color space', *Vision Research*, 22, pp. 1123-1131.

Kuehni, R.G. (2005) 'Focal color variability and unique hue stimulus variability', *Journal of Cognition and Culture*, 5, pp. 409-426.

Ladyman, J.; Ross, D. (2007) *Every Thing Must Go: Metaphysics Naturalized*. (Oxford: Oxford University Press).

Levin, Daniel T., Momen, Nausheen, Drivdahl, Sarah B., Simons, Daniel J. (2000) 'Change blindness blindness: the metacognitive error of overestimating change-detection ability' *Visual Cognition*, 7, pp. 397-412.

Levine, Joseph (1983) 'Materialism and qualia: the explanatory gap', *Pacific Philosophical Quarterly,* 64, pp. 354-361.

Levine, Joseph (1993) 'On leaving out what it's like', in: Davies, M. Humphreys, G.W. (eds.) *Consciousness* (Oxford: Basil Blackwell).

Levine, Joseph (1994) 'Out of the closet: a qualophile confronts qualophobia', *Philosophical Topics*, 22, pp. 107-126.

MacKay, Donald M. (1962a) 'The use of behavioural language to refer to mechanical processes', *The British Journal for the Philosophy of Science*, 13, pp. 89-103.

MacKay, Donald M. (1962b) 'Theoretical models of space perception', in: Muses, C.A. (ed.), *Aspects of the Theory of Artificial Intelligence*, pp. 83-103. (New York: Plenum Press).

MacKay, Donald M. (1967) 'Ways of looking at perception', in: Wathen-Dunn, W. (ed.), *Models for the Perception of Speech and Visual Form*, pp. 25-43. (Cambridge, MA: MIT Press).

Madary, Michael (2008) 'Specular highlights as a guide to perceptual content', *Philosophical Psychology*, 21 (5), pp. 629-639.

Marr, David (1982) *Vision: A Computational Investigation into the human Representation and Processing of Visual Information* (New York: Freeman and Company).

McGurk, Harry; MacDonald, John (1976) 'Hearing lips and seeing voices', *Nature*, 264, pp. 746-748.

Merikle, Philip M.; Daneman, Meredyth (1999) 'Conscious vs. unconscious perception', in: Gazzaniga, M.S. (ed.). *The New Cognitive Neurosciences, Second Edition*, pp. 1295-1303 (Cambridge, MA: MIT Press).

Metzinger, Thomas (ed.) (2000) *Neural Correlates of Consciousness: Empirical and Conceptual Questions* (Cambridge, MA: MIT Press).

Metzinger, Thomas (2003) *Being No One: The Self-Model Theory of Subjectivity* (Cambridge, MA: MIT Press).

Miyahara, E. (2003) 'Focal colors and unique hues', *Perceptual and Motor Skills*, 97, pp. 1038-1042.

Mollon, J.D. (2006) 'Monge: The Verriest lecture, Lyon, July 2005', *Visual Neuroscience*, 23, pp. 297-309.

Mollon, J.D.; Jordan, G. (1997) 'On the nature of unique hues', in: Dickinson, C., Murray, I., Carden, D. (eds.) *John Dalton's Colour Vision Legacy*, pp. 381-392. (London: Taylor & Francis).

Mossio, Matteo; Taraborelli, Dario (2008) 'Action-dependent perceptual invariants: from ecological to sensorimotor approaches', *Consciousness and Cognition,* 17, pp. 1324-1340.

Myin, Erik (2003) 'An account of color without a subject?', *Behavioral and Brain Sciences*, 26, pp. 42-43.

Myin, Erik; O'Regan, J. Kevin (2002) 'Perceptual consciousness, access to modality and skill theories: a way to naturalize phenomenology?', *Journal of Consciousness Studies*, 9 (1), pp. 27-46.

Nagel, Thomas (1974) 'What it is like to be a bat?', *The Philosophical Review*, 83, pp. 435-450.

Noë, Alva (2002) 'On what we see', *Pacific Philosophical Quarterly*, 83, pp. 57-80.

Noë, Alva (2004) *Action in Perception* (Cambridge, MA: MIT Press).

Noë, Alva; Pessoa, Luiz; Thompson, Evan (2000) 'Beyond the grand illusion: what change blindness really teaches us about vision' *Visual Cognition*, 7, pp. 93-106.

Noë, Alva; Thompson, Evan (2004) 'Are there neural correlates of consciousness?', *Journal of Consciousness Studies,* 11, pp. 3-28.

O'Regan, J. Kevin (1992) 'Solving the "real" mysteries of visual perception: the world as an outside memory', *Canadian Journal of Psychology*, 46(3), pp. 461-488.

O'Regan, J. Kevin (2010) 'Explaining what people say about sensory qualia', in: Gangopadhyay, Nivedita; Madary, Michael; Spicer, Finn (Eds.) *Perception, Action, and Consciousness: Sensorimotor Dynamics and Two Visual Systems*, pp. 31-50. (Oxford: Oxford University Press).

O'Regan, J. Kevin (2011) *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness* (New York: Oxford University Press).

O'Regan, J. Kevin; Myin, Erik; Noë, Alva (2004) 'Towards an analytic phenomenology: the concepts of "bodiliness" and "grabbiness"', in Carsetti, Arturo (ed.) *Seeing, Thinking and Knowing: Meaning and Self-Organization in Visual Cognition and Thought*, pp. 103-114. (Dordrecht: Kluwer Academic Publishers).

O'Regan, J. Kevin; Myin, Erik; Noë, Alva (2005) 'Sensory consciousness explained (better) in terms of 'corporality' and 'alerting capacity'', *Phenomenology and the Cognitive Sciences,* 44, pp. 369-387.

O'Regan, J. Kevin; Noë, Alva (2001) 'A sensorimotor account of vision and visual consciousness', *Behavioral and Brain Sciences*, 24, pp. 939-973.

O'Regan, J. Kevin; Noë, Alva (2001b) 'Acting out our sensory experience', *Behavioral and Brain Sciences* 24, pp. 1011-1031.

O'Regan, J. Kevin; Noë, Alva (2001c) 'What it is like to see: A sensorimotor theory of perceptual experience', *Synthese*, 129, pp. 79-103.

Palmer, Stephen E. (1999) *Vision Science: Photons to Phenomenology*. (Cambridge, MA: MIT Press).

Parkes, L.M.; Marsman, J.B.C.; Oxley, D.C.; Goulermas, J.Y.; Wuerger, S.M. (2009) 'Multivoxel fMRI analysis of color tuning in human primary visual cortex', *Journal of Vision*, 9 (1), pp. 1-13.

Peacocke, Christopher (1983) *Sense and Content* (Oxford: Clarendon Press).

Philipona, David L.; O'Regan, J. Kevin (2006) 'Color naming, unique hues, and hue cancellation predicted from singularities in reflection properties', *Visual Neuroscience*, 23, pp. 331-339.

Pridmore, R.W. (2008) 'Chromatic induction', *Color Research and Application*, 1, pp. 77-80.

Prinz, Jesse (2008) 'Is consciousness embodied?', in: Robbins, Philip; Aydede, Murat (eds.) *Cambridge Handbook of Situated Cognition* (Cambridge: Cambridge University Press).

Putnam, Hilary (1999) *The Threefold Cord: Mind, Body, and World* (New York: Columbia University Press).

Quine, W.V.O. (1960) *Word and Object*. (Cambridge, MA: MIT Press).

Ramachandran, V.S.; Hubbard, E.M. (2001) 'Synesthesia – a window into perception, thought and language', *Journal of Consciousness Studies*, 8 (12), pp. 3-34.

Ramsey, William M. (2007) *Representation Reconsidered* (Cambridge: Cambridge University Press).

Regier, T.; Kay, P.; Cook, R.S. (2005) 'Focal colors are universal after all', *Proceedings of the National Academy of Sciences of the United States of America*, 23, pp. 8386-8391.

Rensink, Ronald A.; O'Regan, J. Kevin; Clark, James J. (1997) 'To see or not to see: the need for attention to perceive changes in scenes' *Psychological Science*, 8 (5), pp. 368-373.

Revonsuo, Antti (2006) *Inner Presence: Consciousness as a Biological Phenomenon* (Cambridge, MA: MIT Press).

Robbins, Philip; Aydede, Murat (eds.) (2008) *Cambridge Handbook of Situated Cognition* (Cambridge: Cambridge University Press).

Robbins, Philip; Jack, Anthony I. (2006) 'The phenomenal stance', *Philosophical Studies*, 127, pp. 59-85.

Rock, Irvin (1977) 'In defence of unconscious inference', in Epstein, William (ed.) *Stability and Constancy in Visual Perception: Mechanisms and Processes*, pp. 321-373 (New York: John Wiley & Sons).

Rowlands, Mark (2003) *Externalism: Putting Mind and World Back Together Again* (Chesham: Acumen Publishing).

Ruppertsberg, A.I.; Bloj, M. (2007) 'Reflecting on a room of one reflectance', *Journal of Vision*, 7(13): 12, pp. 1-13.

Ryle, Gilbert (1949) *The Concept of Mind* (edn. 2000, London: Penguin Classics).

Ryle, Gilbert (1954) *Dilemmas* (Cambridge: Cambridge University Press).

Ryle, Gilbert (1979) *On Thinking* (Totowa: Roman and Littlefield).

Sandell, J.H.; Gross, C.G.; Bornstein, M.H. (1979) 'Color categories in macaques', *Journal of Comparative Physiological Psychology*, 93(4), pp. 626-635.

Saunders, B.A.C.; Van Brakel, J. (1997) 'Are there nontrivial constraints on colour categorization?', *Behavioral and Brain Sciences*, 20, pp. 167-228.

Schellenberg, Susanna (2007) 'Action and self-location in perception', *Mind*, 116, pp. 603-631.

Schwarz, J.H. (2000) in: Kandell, E. R.; Schwartz, J.H.; Jessell, T.H. (eds.) *Principles of Neural Science* (fourth edn.), (New York: McGraw-Hill Companies).

Shanahan, Murray (2006) 'A cognitive architecture that combines internal simulation with a global workspace', *Consciousness and Cognition,* 15, pp. 433-449.

Sheppard, R.N. (1992) 'The perceptual organization of colors: an adaptation to regularities of the terrestrial world?', in: Barkow, J.H.; Cosmides, L.; Tooby, J. (eds.) *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, pp. 495-532. (New York: Oxford University Press).

Sherrington, C.S. (1918) 'Observations of the sensual role of the proprioceptive nerve supply of the extrinsic eye muscles', *Brain*, 41, pp. 332–343.

Simons, Daniel J. (2000) 'Current approaches to changeblindness', *Visual Cognition*, 7, pp. 1-15.

Simons, Daniel J.; Levin, Daniel T. (1997) 'Change blindness' *Trends in Cognitive Sciences*, 1 (7), pp. 261-267.

Simons, Daniel J.; Levin, Daniel T. (1998) 'Failure to detect changes to people during a real-world interaction' *Psychonomic Bulletin & Review*, 5 (4), pp. 644-649.

Sperry, R.W. (1950) 'Neural basis if the spontaneous optokinetic response produced by visual inversion', *Journal of Comparative and Physiological Psychology*, 43, pp. 482–489.

Stratton, George M. (1896) 'Some preliminary experiments on vision without inversion of the retinal image', *The Psychological Review*, 3, pp. 611-617.

Stratton, George M. (1897) 'Vision without inversion of the retinal image', *The Psychological Review*, 4, pp. 341-360; 463-481.

Taylor, James G. (1962) *The Behavioral Basis of Perception* (New Haven: Yale University Press.)

Taylor, James G. (1964) 'What is consciousness?' *The British Journal of Statistical Psychology*, 17 (1), pp. 71-76.

Thomas, Nigel J.T. (1999) 'Are theories of imagery theories of imagination? An active perception approach to conscious mental content', *Cognitive Science*, 23 (2), pp. 207-245.

Thompson, Evan (1995) *Color Vision: A Study in Cognitive Science and the Philosophy of Perception* (London: Routledge).

Thompson, Evan (2007) *Mind in Life: Biology, Phenomenology, and the Sciences of Mind* (Cambridge, MA: Harvard University Press).

Thompson, Evan; Varela, Francisco J. (2001) 'Radical embodiment: neural dynamics and consciousness', *Trends in Cognitive Sciences,* 5, pp. 418-425.

Tononi, Giulio; Edelman, Gerald M. (1998), 'Consciousness and complexity', *Science*, 282, pp. 1846-1851.

Tononi, Giulio (2004) 'An information integration theory of consciousness', *BMC Neuroscience*, 5, 42.

Torrance, S. (2002) 'The skill of seeing: beyond the sensorimotor account?', *Trends in Cognitive Sciences*, 6 (12), pp. 495-496.

Tye, Michael (1995) *Ten Problems of Consciousness* (Cambridge, MA: MIT Press).

Uexküll, Jacob von (1920/1928) *Theoretische Biologie* (second edn.), (Berlin: Springer).

Uexküll, Jacob von (1934) 'A stroll through the world of animals and men', in: Schiller, Claire H. (1957) *Instinctive Behavior: The Development of a Modern Concept* (New York: International Universities Press).

Valberg, A. (2001) 'Unique hues: an old problem for a new generation', *Vision Research*, 13, pp. 1645-1658.

Velmans, Max (1990) 'Consciousness, brain and the physical world', *Philosophical Psychology*, 3 (1), pp. 77-99.

Velmans, Max (2007) 'Where experiences are: dualist, physicalist, enactive and reflexive accounts of phenomenal consciousness', *Phenomenology and the Cognitive Sciences,* 6, pp. 547-563.

Varela, Francisco J.; Lachaux, Jean-Philippe; Rodriguez, Eugenio; Martinerie, Jacques (2001) 'The brainweb: phase synchronization and large-scale integration', *Nature Reviews Neuroscience,* 2, pp. 229-239.

Vasquez-Corral, J.; O'Regan, J.K.; Vanrell, M.; Finlayson, G.D. (*forthcoming*). A new spectrally sharpened sensor basis to predict color naming, unique hues, and hue cancellation.

Ward, Dave; Roberts, Tom; Clark, Andy (2011) 'Knowing what we can do: actions, intentions, and the construction of phenomenal experience' *Synthese*, 181, pp. 375-394.

Wilson, M.H.; Brocklebank, R.W. (1955) 'Complementary hues of after-images', *Journal of the Optical Society of America* 45 (4), pp. 293-299.

Wittgenstein, Ludwig (1953) *Philosophical Investigations*, Anscombe, G.E.M. (trans.) 1972, (Oxford: Basil Blackwell).

Wuerger, S.M.; Atkinson, P.; Cropper, S. (2005) 'The cone inputs to the unique-hue mechanisms', *Vision Research*, 45, pp. 3210-3223.

Zeki, S. (1983) 'Colour coding in the cerebral cortex: the reaction of cells in monkey visual cortex to wavelengths and colours', *Neuroscience*, 9, pp. 741-765.

Index

Samenvatting

Hoe kunnen we bewuste ervaring begrijpen als een deel van de natuurlijke wereld? Wat verklaart het specifieke *fenomenale karakter* van de ervaring, dat wil zeggen, wat verklaart hoe iemand zijn of haar omgeving beleeft? Pogingen om bewuste ervaring te begrijpen hebben zich in de afgelopen decennia vaak uitsluitend gericht op de hersenen. In het geval van waarneming wordt ervaring dan meestal opgevat als een intern model van de omgeving in het hoofd van de waarnemer. Tegelijkertijd is er in de gedrags- en cognitiewetenschappen een toenemende gevoeligheid voor een breder scala van processen, die dwars door de grenzen tussen organisme en omgeving gaan. De manipulatie van de omgeving blijkt bijvoorbeeld bij te kunnen dragen aan het vinden van een oplossing voor problemen waarvoor we ons gesteld zien, zoals wanneer iemand stukjes van een puzzel beweegt om te zien waar ze zouden kunnen passen. Het feit dat uitgebreide interne operaties niet noodzakelijk hoeven te zijn voor adequaat gedrag is er aanleiding voor om op zoek te gaan naar een ander beeld van de interne werking van de mens. Ook de studie van de waarneming is steeds meer gericht op dynamische patronen van betrokkenheid op de omgeving. Actieve verkenning speelt een cruciale rol in perceptie, en het is te betwijfelen of onze waarneming berust op uitgebreide interne modellen in de hersenen. Sommige hebben zelfs betoogd dat als de wereld voorhanden is, er helemaal geen interne modellen nodig zijn voor perceptuele ervaring. In dit proefschrift onderzoek ik de gevolgen van dit perspectief voor het begrijpen van het fenomenale karakter van de ervaring. Meer in het bijzonder, ik expliciteer, ontwikkel en verdedig een *sensomotorische benadering*, zoals beschreven in O'Regan en Noë (2001). Ik betoog dat deze aanpak helpt om het fenomenale karakter van onze waarneming te begrijpen in termen van onze sensomotorische betrokkenheid op de omgeving.

Zoals ook Helmholtz (1876) en anderen al hebben opgemerkt, kunnen perceptuele ervaringen niet altijd worden begrepen in termen van sensorische stimulatie alleen. Hoe wij de wereld ervaren is afhankelijk van de manier waarop de stimulatie van onze zintuigen samenhangt met werkelijke of potentiële motorische activiteit; het is afhankelijk van de impliciete erkenning van sensorisch-motorische interrelaties. Bijvoorbeeld, zien dat een object voor een achtergrond staat gaat om het impliciet begrijpen van de sensorische consequenties die kunnen worden verwacht van een lichamelijke beweging; het draait om het weten hoe beweging het mogelijk maakt om achter het object te kijken. We spreken hier van

impliciete kennis omdat je natuurlijk niet expliciet je hoeft te realiseren dat zowel sensorische als motorische factoren een rol spelen; het gaat hier eerder om een intuïtief weten, zoals je ook weet hoe je je evenwicht bewaard tijdens het fietsen zonder te kunnen zeggen hoe je dit doet. Zo is ook bijvoorbeeld de visuele ervaring van een beweging in de omgeving een sensomotorische kwestie. Je kunt de beweging waarnemen door veranderingen van sensorische stimulatie in afwezigheid van oogbewegingen, of juist door een gebrek aan verandering van sensorische stimulatie door het bewegende object in het geval dat men het object met de ogen volgt: in beide gevallen is de ervaring gebaseerd op de manier waarop sensorische stimulatie afhangt van motorische activiteit, ofwel de sensomotorische afhankelijkheden. De vraag is hoe we deze sensomotorische afhankelijkheden moeten begrijpen. Moeten we ze zien als de basis voor de constructie van interne modellen van de wereld, of zijn dergelijke modellen niet nodig? De sensomotorische benadering stelt dat we zonder interne modellen kunnen, en dat we ervaring moeten zien als het uitoefenen van een vaardigheid in perceptuele betrokkenheid op de omgeving.

Er zijn twee vereisten voor een verklaring van het fenomenale karakter van ervaringen (Hoofdstuk 1). Ten eerste is er een nauwkeurige omschrijving nodig van het karakter van de ervaring op het niveau van de persoon. Vervolgens dient deze beschrijving te worden gekoppeld aan een beschrijving van de onderliggende processen op een subpersoonlijk niveau van beschrijving. Als niet voldaan is aan de eerste voorwaarde zouden we in het beste geval een opsomming hebben van de processen die samenhangen met een ervaring, maar een verklaring van het karakter van ervaring is daarmee niet gegeven. Als aan de tweede voorwaarde niet is voldaan hebben we in het beste geval een verhalende opsomming van ervaringen: 'eerst heb ik dit meegemaakt, en toen ervoer ik dat'. Om ervaring in te passen in ons beeld van de natuurlijke wereld moet dit verhaal aansluiten bij niet-fenomenale beschrijvingen, en in het bijzonder bij beschrijvingen van de processen die aan de basis liggen van de ervaring. Een zogenaamde 'dual currency'-verklaring doet dit door een beschrijving te geven die zowel van toepassing is op de ervaring als op de onderliggende processen.

Ik betoog dat een sensomotorische benadering een aantal positieve eigenschappen heeft in vergelijking met benaderingen die uit gaan van interne modellen. Ten eerste helpt de sensomotorische benadering ons om te beschrijven hoe ervaringen zijn voor de persoon. Zo kunnen verschillen tussen visuele en tactiele ervaringen worden geëxpliciteerd in termen van verschillen in de karakteristieke sensomotorische patronen (Hoofdstuk 2). Een focus op sensomotorische patronen helpt ook bij het beschrijven van

het fenomenale karakter van ruimtelijke visuele waarneming (Hoofdstuk 6). Ten tweede, een sensomotorische benadering vereenvoudigt onze kijk op de subpersoonlijke processen die betrokken zijn bij perceptuele ervaring. Een impliciete kennis van sensomotorische afhankelijkheden hoeft niet te worden aangevuld met een additionele opbouw van interne modellen van de omgeving. Hierdoor kan een hele laag worden verwijderd uit de subpersoonlijke verklaring van ervaring, vergeleken met verklaringen gebaseerd op interne modellen (Hoofdstuk 2). We zijn dan een stap dichter bij het begrijpen van het fenomenale karakter van de ervaring dan benaderingen in termen van modellen in het hoofd doen vermoeden. Ten derde kan de sensomotorische benadering helpen bij het verbinden van persoonlijke en subpersoonlijke beschrijvingen van perceptuele ervaringen. De benadering geeft een beschrijving van subpersoonlijke processen die dicht blijft bij de manier waarop we de wereld ervaren, en helpt daarmee bij het vinden van een beschrijving die geldt voor zowel persoonlijke en subpersoonlijke niveaus van beschrijving (Hoofdstuk 1, 2, 3).

In Hoofdstuk 3 verbind ik de sensomotorische benadering met *neurale werkruimte*-benaderingen. Terwijl de sensomotorische benadering zich richt op patronen van interactie met de omgeving, zijn werkruimte theorieën sterk gericht op de hersenen: zij stellen dat een subset van de processen in het brein een 'werkruimte' vormen waarvan de activiteit correleert met onze ervaring. De benaderingen worden vaak gezien als concurrenten, maar ik stel dat ze elkaar juist aanvullen. Het hoofdstuk gaat uit van het onderscheid tussen twee kwesties met betrekking tot fenomenale ervaring (Chalmers 1996; Hurley & Noë 2003). Er zijn de *vergelijkende* kwesties, die betrekking hebben op het specifieke fenomenale karakter van ervaringen (bijvoorbeeld wat verklaart het verschil tussen de ervaring van rood en blauw, of tussen visuele ervaring en auditieve ervaring). Verder is er de *absolute* kwestie, betreffende het feit dat we überhaupt bewuste ervaring hebben. Ik betoog dat sensomotorische en werkruimte-benaderingen het best kunnen worden gezien als gericht op verschillende kwesties. Sensomotorische benaderingen zijn het best geschikt voor verklaren van het specifieke fenomenale karakter van ervaringen, maar ze verklaren niet waarom er ervaring is (Noë 2004). Een sensomotorische karakterisering van de vereisten voor bewuste ervaring blijft te beschrijvend, en moet verder worden uitgewerkt in subpersoonlijke termen. Werkruimte-benaderingen gaan juist wel in op de neurale processen die specifiek zijn voor bewuste ervaring (bijv. Dehaene & Naccache 2001). Deze benaderingen zijn relevant voor de absolute kwestie omdat ze het contrast uitleggen tussen bewuste ervaring en de afwezigheid van bewuste ervaring voor een bewuste waarnemer. Maar deze

benaderingen zijn minder geschikt voor de vergelijkende kwesties binnen de ervaring. Ik bespreek drie scenario's voor het combineren van de benaderingen, die verschillen in ontologische veronderstellingen, maar die alle drie de positieve kanten van de combinatie ondersteunen.

In Hoofdstuk 4 wijs ik op een specifieke bijdrage van de sensomotorische benadering, door te beargumenteren dat haar focus op het fenomenale karakter van ervaringen een cruciale aanvulling vormt op benaderingen die zich richten op wat de waarneming ons kan vertellen over de wereld. Uitgangspunt is de tegenstelling tussen 'qualofobe' en 'qualofiele' stromingen in de filosofie van de geest. Terwijl qualofoben stellen dat bewuste ervaring niets vereist dan de processen die een rol spelen in onze capaciteiten om te handelen en om toegang te hebben tot de wereld (Dennett 1991), dringen qualofielen erop aan dat er meer uit te leggen is dan onze cognitieve functies of onderscheidingsvermogens (Levine 1994; Blok 1996). Ik stel voor om de qualofoben te volgen in het verwerpen van speciale kwalitatieve ingrediënten in onze verklaring van het bewustzijn, terwijl ik tegelijkertijd voorstel om het qualofiele idee te aanvaarden dat verschillende vragen kunnen worden gesteld met betrekking tot fenomenale ervaring en cognitieve toegang. Deze posities kunnen worden verzoend door voort te bouwen op twee invalshoeken van waaruit we bewuste waarnemers kunnen beschouwen: de *intentionele houding* en de *fenomenale houding*. Wanneer we Dennetts (1987) intentionele houding toepassen op een waarnemer, stellen we belang in de potentiële kennis die de waarnemer opdoet van de omgeving, en zijn we daarmee gericht op mentale toestanden die kunnen worden geëvalueerd als waar of onwaar. Bij de fenomenale houding, daarentegen, gaat het er om hoe de ervaring is voor de waarnemer (zie Robbins & Jack 2006). De intentionele houding abstraheert van de processen waarmee we perceptuele kennis verwerven door zich uitsluitend te richten op epistemologisch evalueerbare inhoud; als gevolg hiervan verdwijnt onze belichaamde ervaring uit het zicht. Ik betoog dat, vanuit de fenomenale houding, ervaring moet worden geïnterpreteerd als een waarnemer-gecentreerd begrip, dat de manier betreft waarop de waarnemer perceptueel betrokken is op zijn of haar omgeving. De sensomotorische aanpak biedt een natuurlijke manier om een dergelijke notie van ervaring uit te werken.

Hoofdstuk 5 verbindt conclusies aan de sensomotorische benadering voor de verklaring van kleurwaarneming. Traditionele benaderingen van ervaring gaan vaak uit van een strikte tweedeling tussen interne en externe factoren. Terwijl sommigen zich in de studie van kleurwaarneming hebben gericht op de eigenschappen van de omgeving waarvoor we gevoelig zouden zijn, hebben anderen benadrukt dat cruciale eigenschappen van de

ervaring afhangen van de eigenschappen van onze perceptuele systemen. Een aspect van kleurervaring dat vaak in verband wordt gebracht met neurale factoren is dat een aantal 'unieke' kleuren geen spoor lijken te bevatten van een andere kleur (rood, groen, geel en blauw), terwijl andere kleuren kunnen worden beschouwd als een mengsel van deze kleuren, of als 'binaire kleuren'. Volgens een wijdverbreide aanname moet deze uniek/binaire structuur van kleurervaring worden verklaard in termen van neurofysiologische structureren (bijv. fysiologische 'opponente processen') en heeft deze structuur geen objectieve basis in de prikkel. In dit hoofdstuk bekritiseer ik deze aanname, en ik bekritiseer een argument dat op basis van deze aanname concludeert dat kleurervaring een eigenschap is van de hersenen die onafhankelijk is van de omgeving (Hardin 1988). Ik wijs erop dat, volgens huidige inzichten, de activiteit van bekende neurofysiologische opponente processen niet specifiek correleren met de ervaring van de unieke kleuren. Een belangrijk logisch punt is dat zelfs als er een neuraal correlaat van de uniek/binaire structuur van kleur ervaring te vinden is, dit geen ondersteuning biedt voor de bewering dat de uniek/binaire structuur uitsluitend het gevolg is van neurofysiologische factoren. Een recente analyse van Philipona en O'Regan (2006) heeft aangetoond dat patronen van interactie met de omgeving aan de basis kunnen liggen van de uniek/binaire structuur van kleurervaring. Het blijkt dat onder wisselende lichtomstandigheden oppervlakken met unieke kleuren resulteren in eenvoudigere patronen van stimulatie van de retina dan oppervlakken met binaire kleuren. Dit suggereert dat, mocht een neurofysiologische structuur overeenkomen met de uniek/binaire structuur, dit een gevolg is van de patronen van sensorische stimulatie die voortkomen uit de interactie van de waarnemer met de omgeving. Ik concludeer dat een zuiver neurofysiologische blik tot weinig verklaring in staat is. In plaats van zich uitsluitend te richten op neurale correlaten, of op de omgeving, moeten verklaringen van de structuur van ervaring rekening houden met de patronen die ontstaan in de perceptuele betrokkenheid van waarnemers op de omgeving.

In Hoofdstuk 6 wordt de sensomotorische benadering toegepast in een fenomenologische gevalsbeschrijving. Op basis van mijn ervaring met het dragen van een links/rechts omkeerbril laat ik zien hoe een sensomotorische analyse helpt bij het karakteriseren van visuele ervaring, en ik doe verslag van bevindingen over mentale voorstellingen. Als je voor het eerst een omkeerbril opzet leidt beweging van je hoofd tot een afbraak van de visuele stabiliteit: de wereld lijkt voor je ogen langs te trekken. De metafoor van 'visuele beelden' is onvoldoende om deze ervaring te beschrijven, want de ervaring is belangrijk anders dan de ervaring van het kijken naar bewegende beelden. Het gebrek aan visuele stabiliteit kan beter

omschreven worden als een gebrek aan het vermogen om objecten te volgen met de blik, een vermogen dat is meestal goed functioneert bij het bekijken van bewegende beelden. Ik wijs er verder op dat een omkeerbril een tot nu toe onderbelicht conflict introduceert in het hart van ruimtelijke visuele ervaring. Hoewel de glazen de relatie tussen oogbewegingen en de stimuli van objecten in de omgeving veranderen, laten ze ongewijzigd hoe hoofdbewegingen de kijkrichting te bepalen. Als gevolg hiervan was ik me vaak bewust van de algemene richting van mijn zicht, ondanks de omkering binnen mijn gezichtsveld. Ook wat dit betreft kan de resulterende ervaring niet adequaat worden weergegeven in termen van (omgekeerde) visuele beelden. Mijn latere bevinding met betrekking tot visuele adaptatie versterkt deze conclusie. Na verloop van tijd hervond ik visuele stabiliteit, en zelfs de vaardigheid om te zien waar voorwerpen zich bevinden was hervonden (zoals ook gerapporteerd in bijvoorbeeld Stratton 1897; Taylor 1962; Kohler 1964). Cruciaal is dat perceptuele adaptatie gepaard ging met een ervaring van Gestalt-switches (zoals in de verschillende manieren van ervaren van het beroemde plaatje dat kan worden gezien als eend of als haas). Ik kon de omgeving op verschillende manieren ervaren, zonder hierbij een omkering van een visueel beeld te ervaren. Mijn ervaring met het dragen van een omkeerbril illustreert daarmee het onderscheid tussen een object-georiënteerde houding en de fenomenale houding: dezelfde perceptuele oordelen (met en zonder omkeerbril) kunnen gepaard gaan met sterk verschillende ervaringen. Tot slot doe ik verslag van bevindingen waaruit blijkt dat ook visuele voorstelling afhangt van echte of voorgestelde lichamelijke bewegingen, wat ondersteuning biedt aan het idee dat ook voorstelling gegrond is in onze sensomotorische betrokkenheid op de omgeving.