BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ATTENTION, SUBJECT, AND WORLD

by

CAROLYN M. SUCHY-DICEY

M.A., University of St Andrews, 2004

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2011

Approved by

First Reader

Daniel O. Dahlstrom, Ph.D.
John R. Silber Professor of Philosophy
Boston University

Second Reader

Walter Hopp, Ph.D.
Associate Professor of Philosophy
Boston University

Third Reader

Ned Block, Ph.D.
Silver Professor of Philosophy, Psychology,
and Neural Science
New York University

Dedication

To my father, Timothy Richard Dicey, who used mathematics to train my attentional control. Without this training I would not have been able to sustain such a long and arduous study. And to my mother, Donna Pauline Suchy, who finds possibility everywhere. I aim to follow her example.

Acknowledgments

I would like to acknowledge the intellectual contribution of several philosophers and scientists. My committee is an obvious starting point. Daniel Dahlstrom took care to determine my intended message and helped me to communicate it. Walter Hopp likewise pushed me toward clarity and philosophic depth. Ned Block inspired and sustained my interest in the topic. John Campbell guided me toward more defensible versions of my theses. On these and countless other contributions, the foundation of the dissertation rests.

Others outside of the committee deserve considerable credit for shaping the work. Takeo Watanabe has spent several years pointing me to research in psychology and neuroscience that has proven vital. Christof Koch met with me on a few occasions to discuss my ideas, in one case leading to an important shift in my thinking. Ruth Millikan, Sean Kelly, Jesse Prinz, Christopher Hill, and Justin Junge gave me insightful comments on drafts of various chapters.

Also, the philosophical community at Boston University and its surrounding universities has influenced the work both directly, through discussion of ideas contained within it, and indirectly, through discussion on connected topics. Most notably, the influence of Susanna Siegel, Alex Byrne, Daniel Dennett, Daniel Pollen, Tian Yu Cao, David Roochnik, Rosa Cao, and David Jennings has been valuable.

Finally, I am grateful for the material support of Boston University over the past five years, as well as the financial contribution of the Philosophy Department, the Center for Philosophy and History of Science, and the Mind Science Foundation to some of my pet projects. Such contributions allowed me to bring in speakers on this and other topics, considerably improving my understanding of the issues.

ATTENTION, SUBJECT, AND WORLD

(Order No.)

CAROLYN M. SUCHY-DICEY

Boston University, Graduate School of Arts and Sciences, 2011

Major Advisor: Daniel O. Dahlstrom, John R. Silber Professor of Philosophy

Boston University

ABSTRACT

Everyone is familiar with attention. In *It's a Wonderful Life*, Zuzu attends to her flower. In *The Scarlet Letter*, Pearl attends to the "A" adorning her mother's clothing. In *Christina's World*, Christina attends to the house as she crawls toward it. Attention binds Zuzu, Pearl, and Christina to the objects of their respective interests.

Despite this familiarity, a satisfactory account of how attention works has proven stubbornly elusive. Specifically missing is an account of how an experiencing subject, like Zuzu, is able to attend to a physical object, like her flower. The experiencing subject and the objective world, thought by many to be irreconcilable, appear to be brought together through attention.

This dissertation supplies an account of attention and its role in the relationship between subject and world. After establishing certain markers of attention afforded by phenomenology, psychology, and neuroscience, it applies these markers to debates concerning the contribution of attention to other mental processes. This application supports the peripheral findings that attention is not necessary for sensory selection, action, or non-perceptual conscious experience. This application also makes room for the key finding that attention *is* necessary for conscious perception.

The dissertation presents this key finding in the form of a theory, dubbed the "Standard Theory." According to the Standard Theory, attention is necessary for

the structure of conscious perception. Attention is able to provide this structure, it is argued, by prioritizing subconscious sensations with respect to a subject-based standard, allowing for the differentiation of those sensations according to this unifying standard.

The Standard Theory has the singular advantage of starting from an account of attention that allows attention to operate both within and outside of conscious experience, thereby enabling the theory to escape difficulties besetting other theories that link conscious perception and attention. In this account, the subject *brings about* attention, but this bringing about need not be conscious so long as it is within the subject's control. It is accordingly argued that conscious perception is an achievement enacted by the subject through attention, but not necessarily a conscious achievement.

Contents

List of Figures

List of Acronyms

ERP Event-Related Potential

RT Reaction Time

IOR Inhibition of Return

SOA Stimulus Onset Asynchrony

FEF Frontal Eye Fields

LIP Lateral Intraparietal

DLPFC Dorsolateral Prefrontal Cortex

fMRI Functional Magnetic Resonance Imaging

FIT Feature-Integration Theory

IB Inattentional Blindness

SDT Stimulus Detection Task

Chapter 1

Introduction and Background

Once upon a time, two strange events occurred. The first occurred early one morning when my fiancé, David, turned to me and asked, "Don't you love the sound of the trumpet?" I listened for a moment, but heard only natural sounds – birds chirping, trees swaying. "What trumpet?" I responded. The expression on David's face was one of muted horror. "Oh no!" I thought, "Am I going deaf?" David regularly hears sounds that I don't, but the sound in question has never been one so striking as a trumpet. I listened a little harder, turning my attention fully to the soundscape. A faint trumpet song emerged, some distance away. I said, with relief, "Ah, I hear it now."

The second event occurred some time later. While cooking dinner, David remarked, "Dill is one of my favorite herbs." Having a terrible olfactory memory, I responded, "I don't know how dill smells." David cut off a sprig and offered it to me. I put the leaves quite near my face, sniffing. "I can't smell it!" David's horrified expression returned. I was perplexed. How is it that I could not smell anything for what I know to be a striking scent? I sniffed and sniffed, until the leaves were pressed against my nostrils, attending fully to the smellscape. Nothing. And then, suddenly, there it was: the clear and distinct smell of dill.

I thought to myself, "What did I hear before I recognized the sound of the trumpet?" and "What did I smell before I recognized the scent of dill?" It seems unlikely that I did not hear the sound of the trumpet before that revelatory moment. Yet, the sound came as an epiphany, suddenly, against what I took to be only natural

morning sounds. Likewise, the scent of dill came suddenly against what I previously took to be the absence of smell. What changed that allowed me to perceive what I was hitherto unable to detect?

While the rest of the dissertation focuses almost entirely on the case of vision, I here use examples of sound and smell because my visual sense is keener than my other senses. It is a rare occurrence that someone will point something out to me that I cannot see. Sound and smell remind me that the world is not always there for me to observe, that I must sometimes work to perceive it. This work is known to most of us as "attention" – whether we see, hear, smell, taste, or feel something seems to be partly determined by whether we attend to that thing. A central question of this dissertation is how much our ability to perceive something is determined by whether we attend to it. I find, ultimately, that conscious perception relies on the act of attention, such that we do not consciously perceive something without our first attending to that thing.

There are at least two reasons to be initially suspicious of this conclusion. First, it does not seem as though we attend to everything that we perceive. There is a background in visual perception, for example, that seems to exist before any act of attention. Second, it appears contradictory to say that we must attend to something in order to perceive it, since we might then have nothing on which to focus our attention. How can we attend to the sound of the trumpet, for example, before hearing it?

I find that these worries arise partly out of a misunderstanding of what attention is and how it is directed. One's understanding of attention could easily be misdirected by the prominence of the phrase, "pay attention." When asked to "pay attention," one is being asked to consciously alter one's state of attention. I remember, for instance, my ballet teacher asking me to "pay attention" while I was attending to the

interesting patterns created by a sandstorm outside. My teacher was asking me to stop attending to the sandstorm and to start attending to the movements I was supposed to learn. The phenomena picked out by the command to "pay attention" might include all the cases in which we consciously direct our attention, but it will leave out any attention that occurs prior to conscious direction. The attention, for example, employed *before* one is asked to "pay attention" is often a result of *subconscious* direction. In my memory of the event, I did not attend to the sandstorm because of a conscious directive, but because of a subconscious directive that was nonetheless mine. That is, I found myself attending to the sandstorm, but allowed myself to continue attending to it – my attention was not outside of my control. To accommodate the cases in which attention is directed subconsciously, our understanding of attention will have to include more than the type brought about after being asked to "pay attention."

If we allow for attention that is not consciously directed, but is nonetheless directed by the subject, we see that the above worries dissolve. With respect to the first worry, even if we do not *consciously* attend to everything that we perceive, we may yet *attend* to everything that we perceive (through a mix of conscious and subconscious attention). With respect to the second, even if it is a contradiction to say that we must *consciously* attend to something in order to consciously perceive it, it is not a contradiction to say that we must *attend* to something in order to consciously perceive it. This is because attention can be subconsciously directed, and subconscious direction is not limited by conscious perception. In other words, the subject can subconsciously direct attention to sensations that are not yet conscious.

The working definition of attention that I propose centers on the role of the subject, rather than conscious experience. If we see conscious experience as a relation between an experiencing subject and an experienced world, then neither the subject

nor the world will be fully contained within experience. One reason to believe that the subject and the world exist outside of and allow for experience is their continuation in the absence of conscious experience. This understanding of the subject – as outside of conscious experience – is already used in the commonsense understanding of thought and belief, where it is natural to assume the existence of both subconscious thoughts and subconscious beliefs that belong to the subject.[1] I presume in what follows that the reader can at least countenance this basic picture.

Presuming, then, a subject outside of conscious experience, I understand attention to be a process that emerges from interaction with the subject, or a process that occurs at the "subject level." Specifically, I understand attention to be the process of *subject-level prioritization*. In other words, attention is what accounts for the volitional[2] ranking of some activities and percepts over others according to goals held by the subject, where the subject is understood as that to which we attribute such capacities as consciously experiencing, knowing, thinking, planning, and perceiving.[3] I take this working definition of "attention" to reflect its ordinary usage, and it is from that definition that I attempt to find the markers of attention that help me to later reveal the relationship between attention and other mental processes.

The scholarly background of the dissertation began with an article by Christof Koch and Naotsugu Tsuchiya called "Attention and consciousness: two distinct brain processes." In this article the authors claim that the concepts of "attention" and "con-

[1]For example, it is natural to assume that you, the reader, might have subconsciously believed that the use of second-person writing is annoying, without your consciously believing it prior to reading this sentence. You might now be subconsciously thinking that second-person self-reflexive writing is more annoying than second-person non-reflexive writing, without committing this to your conscious thoughts.

[2]As will be seen in Chapters Two and Three, my sense of "volitional" is weaker than its normal sense of "volitionally directed," and can include actions that are only "volitionally accepted."

[3]I initially thought that I could leave out a specification of the subject, but was convinced by Dan Dahlstrom and Sean Kelly that I should provide a provisional sketch. Ned Block further pointed out to me that my early assumption that the subject of experience and the subject of action are united was unwarranted, leading me to a helpful division of these in Chapter Three.

sciousness" (where they seem to mean "conscious perception") are often conflated but that they are both conceptually and empirically distinguishable (Koch and Tsuchiya, 2007). Specifically, they claim that the function of attention is distinct from the function of consciousness and that there are instances of both attention without consciousness and of consciousness without attention. In support of the latter claim they cite evidence of conscious perception outside the "tunnel" of volitional attention, gist perception, single-object perception, dual-task paradigms, and so-called "zombie" behaviors, all of which will be discussed in the dissertation.

Koch and Tsuchiya's article fits into a wider landscape of articles on the relationship between attention and conscious perception. On the one hand, philosophers from Immanuel Kant to William James have discussed the (at least partial) reliance of conscious perception on attention. Kant writes in *The Critique of Pure Reason* that attention reveals the influence of the "understanding" on our "inner sense" (Kant, 1996, p. 195). James writes in *The Principles of Psychology* that we choose what we experience through attention (James, 1981, p. 424). More recently, this purported reliance of conscious perception on attention has been taken up in both philosophy and the cognitive sciences. In "Attention: the mechanisms of consciousness," Michael Posner contends that an understanding of attention is necessary for an understanding of consciousness (Posner, 1994, p. 7398). Philip Merickle and Steve Joordens stipulate in their paper "Parallels between perception without attention and perception without awareness" that "perception with and without awareness and perception with and without attention are equivalent ways of describing the same underlying process distinction" (Merickle and Joordens, 1997, p. 219). In their article "To see or not to see: the need for attention to perceive changes in scenes," Ron Rensink, J. K. O'Regan, and J. J. Clark find that the conscious perception of change relies on attention (Rensink et al., 1997). Arien Mack and Irwin Rock write in *Inattentional*

Blindness that there is no conscious perception without attention, basing this claim on their finding that the competition for attention can make a significant difference in whether participants detect a stimulus (Mack and Rock, 1998a, p. ix). Jesse Prinz presents a theory that defines conscious experience as attended "intermediate-level" representations, in "A neurofunctional theory of visual consciousness" and in later work (Prinz, 2000b,a). Stanislas Dehaene, et al. suggest a brain mechanism for the reliance of conscious perception on attention, according to which "bottom-up activations" require accompanying signals from "top-down areas" to allow for conscious perception, in "Conscious, preconscious, and subliminal processing: a testable taxonomy" (Dehaene et al., 2006).

Several philosophers and cognitive scientists have argued, on the other hand, that conscious perception does not rely on attention, adding their voices to those of Koch and Tsuchiya. In their study "Neural mechanisms of selective visual attention," Robert Desimone and John Duncan hypothesize that attention limits conscious perception only for higher-order representations, such as those we can report (Desimone and Duncan, 1995, p. 39). John Campbell writes in "Sense, reference, and selective attention" that attention provides for reference to objects rather than the conscious experience of those objects (Campbell, 1997), an argument that is continued in his later work (Campbell, 2002, 2004). In "Why visual attention and awareness are different," Victor Lamme writes that attention is only necessary for report and not for awareness (Lamme, 2003). Eric Schwitzgebel finds that participants in his experiment report having consciously perceived unattended stimuli; the details can be found in his work "Do you have constant tactile experience of your feet in your shoes?: Or is experience limited to what's in attention?" (Schwitzgebel, 2007, p.14). In "Attention and consciousness," Christopher Mole puts forward the hypothesis that attention is only necessary for "conceptually structured representations" (Mole, 2008, p. 98).

This sketch, despite its brevity, hopefully conveys a sense of the vivacity of the debate between the two sides. This dissertation presents mostly positive evidence for the first side. I initially thought to present negative arguments against the views on the second side but found that many of those views emerge from an understanding of attention that I consider insupportable (because they unduly limit attention, e.g., to focal attention), and so I have taken a different approach. My approach is to first present what I take to be a defensible understanding of attention and to then show what the consequences of that understanding must be.

Specifically, I start the dissertation in Chapter Two by presuming the working definition of attention that I supply above, "the process of subject-level prioritization," and examine its fit with various phenomenal, behavioral, and neural evidence. Because some of the available evidence concerns subject-level or volitional activity, in general, I organize the chapter around this more general type of process, applying my findings to attention along the way. In the course of searching for markers of subject-level activity in these domains, I look for a means of distinguishing attention from non-attention – one that allows for later discussions on the relationship between attention and other mental processes. The phenomenal evidence that I review calls for an extension of what is considered to be subject-level activity to include the mere receptivity or allowance of that activity. Thus, attention becomes any process of prioritization that is volitionally-directed or volitionally-accepted by the subject. This extension of what is considered to be subject-level activity makes it more difficult, from within experience, to distinguish from pre-subjective activity, and so I suggest that we study subject-level activity by supplementing the phenomenal evidence with both behavioral and neural evidence. Chapter Two culminates in the suggestion of particular neuropsychological markers of subject-level activity, where such markers depend on all three types of evidence for their definition (i.e. phenomenal, behav-

ioral, and neural evidence). Namely, I suggest that we distinguish attention through the neuropsychological markers of long-range recurrency between the planning and sensory areas and serial resource-pooling.

In Chapter Three I use the second of these neuropsychological markers to examine the relationship between attention and action. The chapter explores the relationship between attention and selection in general, since many researchers hold that attention is necessary for selection.[4] I present reasons to think that attention is required for neither sensory selection nor action selection. However, it is action selection that I find deserves the most discussion, and in that discussion I use the neuropsychological marker of "serial resource-pooling" to show that attention is not necessary for the selection of action. I argue that voluntary actions, in particular, or actions that could have been controlled by the subject but are not performed at the subject-level (i.e. they are neither volitionally-directed nor volitionally-accepted by the subject), do not require attention.

Chapter Four delivers the heart of the dissertation in the form of a theory, which I call the "Standard Theory." This theory holds that attention is necessary to organize sensations with respect to a subjective standard (a standard that is, in some sense, rooted in or held by the subject), thus allowing for the transformation of sensations into percepts. I find such organization necessary because of the structure of division and unity found within conscious perception, which could not otherwise be supported by the neural correlates of sensation. This is a higher level division and unity than that invoked by most other researchers, who for the most part look at the division and unity of features and objects. The structure of division and unity that I cite spans

[4]For the entirety of this dissertation, when I say that a is "necessary" for b, I mean that all known *types* of b can be shown to coincide with a, where the purported existence of new types of b must be justified by either conceptual or empirical evidence that seems relevant to the necessity relation in question. Haecceitistic properties, for example, are unlikely relevant to most, if not all, necessity relations. This is discussed at more length in Suchy-Dicey (2011).

the entirety of perceptual experience. The Standard Theory, in short, is the theory that attention is necessary for conscious perception because it is required to provide the division and unity found across all instances of conscious perception. Attention achieves this by transforming sensations into percepts by prioritizing those sensations with respect to a subjective standard.

Chapter Five looks at evidence of types of conscious experience that may not require attention. Although I say above that the dissertation presents mostly positive evidence for the claim that attention is necessary for conscious perception, there are exceptions to this strategy. For the arguments of two views potentially at odds with mine, namely, those of John Campbell and Ned Block[5], warrant explicit discussion. Both views present types of conscious experience that may exist outside the reach of attention, where these types of experience may or may not be perceptual in nature. I claim that there is insufficient evidence to infer on the basis of the phenomena put forward by these theorists that there is a second form of conscious experience that is separable from attention. I present, instead, evidence for a new type of conscious experience that likely exists outside the reach of attention. I call this type of conscious experience "immersion consciousness." Immersion consciousness does not present an obstacle to the Standard Theory, since it does not have the structure of conscious perception.

The final chapter of the dissertation presents a case study for the Standard Theory. Where the Standard Theory holds that attention is necessary for conscious perception, Chapter Six involves a study of the phenomenon of gist perception, which is thought by some to be evidence of conscious perception outside the reach of attention. I find that gist perception is sometimes modulated by attention but is sometimes outside

[5]I did not mention Ned Block's view above because he does not explicitly discuss the relationship between attention and conscious perception. This point will be addressed in Chapter Five, where I discuss his work.

the reach of attention, according to the markers that I lay out in Chapters Two and Three. In the case where gist perception is outside the reach of attention, I find a lack of evidence that the gist percept is consciously experienced, and so suggest some tests that could verify whether it is consciously experienced. If such tests were able to determine that the gist percept is consciously experienced in these cases, this would be a serious blow to the Standard Theory.

The reader may note the absence, in the dissertation, of any discussion on three connected points. First, I do not discuss whether attention is necessary for perception, in general. The reason for this is that the evidence I use to argue for the Standard Theory comes from conscious experience. Thus, although it seems likely that subconscious perception (should it exist) has the same character of division and unity found in conscious perception, I do not have evidence to that effect.

Second, I do not discuss whether attention is sufficient for conscious perception. I do not discuss this matter for the simple reason that attention, on my account, need not apply exclusively to conscious perception. I claim above that attention is the process of subject-level prioritization, but this prioritization may be of processes that have little to do with perception, such as the prioritization of thought. Since attention need not be limited to the prioritization required for conscious perception, it can hardly be considered sufficient for conscious perception.

Third, I do not discuss in the dissertation whether attention is necessary for conscious sensory imagination and memory. I presume that these may have structures analogous to conscious perception that likewise require attention, but I am not sure whether attention is necessary in all cases, especially regarding memory.

Chapter 2

The Search for Attention

ABSTRACT: In debates on the relationship between attention and other mental processes, researchers use conflicting notions of attention, generating contradictions and misleading results. My aim in this chapter is to find a clear understanding of attention that complements the normal use of the term but also fits its many associated phenomena. I find that the normal use of "attention" fits the working definition of "subject-level prioritization," where the "subject" is what we take to be the source of conscious experience and thought and where "subject-level" applies to any process currently within the subject's conscious or subconscious control. I thus approach the problem by trying to specify the more general "subject-level activity." I contend that neither phenomenal nor behavioral evidence yield a fully successful specification of subject-level activity. I suggest in their place a neuropsychological account in terms of recurrency between the planning and sensory areas. Attention can then be understood as prioritization that occurs through such recurrency.

2.1 Introduction

When parents direct their children to pay attention, they are invoking the commonly held truth that paying attention helps one to see, hear, and generally experience the world. We might then ask ourselves the question: must we attend to something in order to consciously perceive it? An affirmative answer has old philosophical roots. Georg W. F. Hegel, for instance, appears to hold this view when he asserts that "apart from such attention there is nothing for the mind" (Hegel, 1971, p. 448). A

common worry against such an answer is that this reliance on attention might render conscious perception an act of will. Commenting on this passage, for example, Hegel scholar Willem de Vries comments that "this notion seems implausible, for attention (*Aufmerksamkeit*) is a fairly strong word, implying a high degree of conscious mental activity and willful self-control" (de Vries, 1988, p. 112).[1] Given that the world sometimes jumps in front of you without a corresponding sense of willfulness or volition, if attention is essentially tied to felt volition then attention must not be necessary for conscious perception. On the other hand, if attention is not tied to felt volition, then de Vries' worry (while properly applied to Hegel) should not concern us in our pursuit of a general answer to the question of whether attention is necessary for conscious perception.

The above considerations illustrate just why a common understanding of attention is required for debates on the relationship between attention and other mental processes, such as conscious perception. Participants in these debates, however, often express the concern that the very use of "attention" is nonstandard. As Jeremy Wolfe and Todd Horowitz complain: "even though William James famously declared that 'Everyone knows what attention is', there is no single, satisfying definition of attention" (Wolfe and Horowitz, 2004). Felipe de Brigard and Jesse Prinz contend:

> those who try to move beyond that suggestion that "everyone knows what attention is" often replace the folk concept with idiosyncratic definitions that settle crucial questions by fiat rather than facilitating the process of scientific investigation and discovery. (Brigard and Prinz, 2010)

Without a clear understanding of attention it is difficult to find points of agreement among participants in these debates and to make conceptual progress.

[1] In this dissertation, all emphases contained in quotes are original to the author(s) of the quote.

In this chapter I approach the problem of finding a clear understanding of attention by first presuming that attention is a process of prioritization performed by subjects.[2] I use the more particular "prioritization" over the more general "selection" because "selection" is ambiguous between processes that fully separate the selected from the deselected entities and processes that simply bias the selected over the deselected entities. Attention is a process of selection in the latter sense, which I (and others) call "prioritization."[3] As a provisional definition of "subject," the term is intended to pick out that to which we attribute such capacities as consciously experiencing, knowing, thinking, planning, and perceiving. As such, I am here claiming that the normal use of the term "attention" requires that neither rocks, trees, nor bodily limbs can attend. This feature of attention – that it is carried out exclusively by subjects – is central to the normal use of the term, and so I presume that it should be included in any more technical definition.[4] Unfortunately, a full set of criteria for what counts as a subject, whether in the physical or metaphysical domain, seems to be beyond our current capabilities. We can instead specify "subject-level activity," or find criteria that reliably separate processes currently within the subject's control (or under the subject's influence) from those processes outside of the subject's control, and this specification should help us to mark off the territory of attention as one such activity.

In order to specify subject-level activity I start with phenomenal evidence and then proceed to behavioral and neural evidence. I take this route in part to reflect the

[2]I find unconvincing Christopher Mole's adverbial view of attention, which takes attention to be a mode of other processes rather than its own process (as in Mole (2011a)). For the most part, Mole's arguments are negative ones. That is, attention is taken to be adverbial because he cannot find a process-first understanding of attention that can capture its many phenomena. I consider my understanding of attention to fit its many phenomena and, thus, to remove the motivation for Mole's position.

[3]For example: "Selective attention is the prioritization of processing of certain stimuli within the external and internal environment for further processing" (Mansell et al., 2008).

[4]This is not to say that normal use should fully dictate technical use but that the central meanings of normal use should be preserved if we are to be talking about the same phenomenon.

historical development of work on attention but also to address the three areas of study that I take to have foundational importance for any study of mental processes. I find that subject-level activity (of which attention is an example) is best specified through neuropsychology, where the neuropsychological division of the subject's influence from that of the world is triangulated by directing phenomenal and behavioral evidence at the available neural evidence. A neuropsychological division, I argue, is fundamental to research on attention and thus to any investigation of the contribution of attention to other processes of the mind, such as conscious perception.

2.2 The Phenomenal Search

A natural first attempt to specify subject-level activity might fixate on the phenomena immediately and directly available to the subject, since recognition of one's own subjectivity through such phenomena seems to be fundamental to the recognition of subjectivity in others.[5] The phenomena immediately and directly available to the subject include at least two types of features which may serve to specify subject-level activity, which I call "introspectible qualities" and "felt activity."[6] To illustrate these two types of features, pay attention to

this.

As you attend to the written letters, reflect on any changes in the quality of the percept. You will likely notice an increase in detail in the structure of the letters and in the contrast between the ink and paper. The subjective phenomena of attending, then, appear to involve changes in *what* is experienced. Now, attend to the written

[5]The following philosophical papers take this approach in discussing the relation of attention to consciousness while relying on a phenomenological account of attention: Dretske 2000 (p. 110), Mole 2008 (pp. 98-101), and Watzl 2010 (pp. 3 and 23).

[6]"Introspectible qualities" do not line up with the "qualities" referred to in the historical movement of phenomenology (e.g. by Husserl). The latter term refers to the phenomenal "feels" of different intentional acts.

letters again but reflect instead on your action in attending. You will likely notice an involvement of effort while you attempt to select the written letters over other competing visual stimuli. The subjective phenomena of attending thus also appear to involve a change in *how* one experiences, where this is not reducible to changes in what is experienced. I will assume in this chapter that these are the only types of features found through phenomenal reflection that are relevant to the specification of subject-level activity.

An examination of these two types of features reveals that neither successfully specifies subject-level activity. As I discuss below through the example of attention, the subject is not required for changes in introspectible quality: increased detail and contrast occur in cases where the features of a stimulus are exaggerated without the influence of the subject. The letters may appear more detailed, for example, if seen through the fovea rather than the perifovea, or if seen at a higher resolution. In fact, every introspectible quality seems to be jointly affected by the influence of the subject and by changes to the stimulus (whether through pre-subjective processing or through changes to the objective stimulus).

After dismissing introspectible qualities as candidates for specifying subject-level activity, I question the feasibility of felt activity as a candidate. I make this move, in part, because of the problem of captured attention. As an instance of captured attention, take the case of your attending to the laughing woman on the right. You might describe her image as having distracted you and as having drawn your attention away from the text. That is, you likely attended to her image with either no felt activity or with so little felt activity that it was difficult to distinguish from no felt activity. Ac-

Figure 2.1: The Laughing Woman

counting for the problem of captured attention renders the criteria of felt activity (at best) difficult to determine. Thus, although felt activity has the advantage of exclusively reflecting the contribution of the subject, it is not an attractive candidate for specifying subject-level activity.

In the paragraphs to follow I discuss these two attempts to specify subject-level activity through its subjective phenomena, focusing on the case of attention. Neither attempt is satisfactory, but the second attempt reveals an aspect of subject-level activity (i.e. felt activity) that will help us to develop our account through behavioral and neural evidence.

Introspectible Qualities We may sympathize with Edward Titchener when he says that "the analytical study of attention must center about the sensory attribute of clearness or vividness" (Titchener, 1910, p. 180), particularly after observing that the experience of attention usually involves clarity or vividness. When we attend to the written letters in the example above, they seem clearer and more vivid, while the surrounding de-selected text simultaneously becomes less clear and less vivid.[7] However, one worry that counts against this early approach is that the commonsense attribution of "attention" to a subject does not depend on the experienced qualities of an object but on the relation between the subject and (what we take to be) the received stimulus, of which the qualities are a mere guide. For example, we would not necessarily say of someone who is staring at a low-resolution image that he or she is not paying attention, even if his or her resulting experience is ill-defined. Similarly, we would not necessarily say of someone who glances at a high-definition image that he or she is attending to that image, even if that person has a well-defined percept.

[7]The difference here is subtle, and has been measured by researchers to be of the order of only a few percentage points. See, e.g., Carrasco et al. (2004).

Thus, our ordinary ascriptions of "attention" to a subject appear to depend on the relationship between the subject and the received stimulus.

William James likely shares this worry, though at first glance he may appear to agree with Titchener that attention is tied to the experience of clarity or "concentration."

> Every one knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. (James, 2010, p. 403)

Note that James here says that concentration is the "essence" of attention. Later in the passage he says that the "real opposite" of attention is "distraction." James next asserts that any degree of concentration constitutes a state of attention:

> The awakening [of attention] may come about either by reason of a stimulus from without, or in consequence of some unknown inner alteration; and the change it brings with it amounts to a concentration upon one single object with exclusion of aught besides, or to a condition anywhere between this and the completely dispersed state. (James, 2010, p. 404)

As Christopher Mole points out, this account appears self-contradictory since it holds both that attention is *essentially* concentration and that attention is anything between concentration and complete dispersal, or non-concentration (Mole, 2011a).[8] This apparent inconsistency can be resolved if we do not think of James' "attention" as a mere quality of experience but instead as a process that results in a range of qualities.

[8]Christopher Mole tries to solve this seeming self-contradiction by suggesting that we ignore James' willingness to extend attention to dispersed states and "take seriously" his initial description of attention as focused rather than dispersed (Mole, 2011a, p. 158).

That is, it is resolved if we understand James to define "attention" as the *act* of concentration which has the *resulting experience* of anything from concentration of the introspectible qualities to their complete dispersal.[9]

James might avoid tying attention to introspectible qualities, relying instead on the activity of concentration, because one cannot fully separate the contribution of the subject from that of the world through the introspectible qualities of experience. That is, changes in the subject can, in some cases, produce the same types of qualities that changes in the world do. For this reason, the subject is often unable to distinguish the source of changes in quality.

This effect is described in the work of Marisa Carrasco et al. who show that attention mimics pre-attentional changes to the stimulus "by enhancing its 'effective contrast' or salience" (Carrasco et al., 2004). Carrasco et al. discovered this effect by asking participants to compare the contrast of different gratings after the participants' attention was drawn to one of those gratings with a peripheral cue. Participants typically rated the contrast of an attended grating as having equal contrast to that of an unattended grating with lower contrast.

Keith Schneider disputes the claim that this effect is due to attention, rather than to pre-attentional processing. He replicates the Carrasco et al. effect for luminances close to the threshold of what participants are able to detect ("perithreshold stimuli"), but shows that the attentional effect drops off for higher luminances ("suprathreshold stimuli"). He concludes:

> The results of the four experiments suggest the presence of two distinct

[9]I am assuming that James means by "concentration" in the second quote ("...amounts to a concentration...") the resulting concentration of introspectible qualities rather than the concentration of attention's resources. If the latter, the conflict dissolves in a different way, since it would not be a problem for attention to be essentially concentration but to concentrate its resources in different ways. John Campbell pointed this out to me as a possible criticism of the interpretation I give above.

> mechanisms that can appear to alter perceived contrast: (1) an attentional mechanism independent of the sensory nature of the cues that operated only on perithreshold stimuli and (2) a sensory interaction mechanism that also operated on suprathreshold stimuli. (Schneider, 2006, p. 808)

He further concludes that attention may only alter the participant's detection threshold rather than the perceived contrast (Schneider, 2006, p. 809), which, if true, would seriously undermine Carrasco et al.'s findings.

Against Schneider's findings, Sam Ling and Marisa Carrasco highlight some problematic features of his experiments. First, they show that his task was different from that of Carrasco et al. in a crucial respect: Schneider's participants were asked to detect luminance rather than perceived contrast, but luminance and perceived contrast are thought to be processed differently in the brain and to operate independently in "natural images." Second, they show that the cue used in Schneider's task may have been responsible for the variability in his results, since this cue is known to operate as a "paracontrast" mask, a mask that interacts with perceived luminance. Ling and Carrasco report that the cue "in [their] original study was specifically chosen to elicit transient attention while avoiding any confounds from masking" (Ling and Carrasco, 2007, p. 1053). Finally, Ling and Carrasco test a prediction set forward in Schneider's paper, a prediction that Schneider claims will prove that any alteration in appearance is due to pre-attentional mechanisms. They show that the prediction is not upheld, even while remaining true to Schneider's experimental specifications.

Thus, Carrasco et al. seem to successfully show that attention can alter perceived contrast. However, as Schneider points out, perceived contrast can also be brought about by pre-attentional mechanisms (a point that is not disputed by Ling and Carrasco). In other words, the subject-level activity of attending can change

phenomenal experience by increasing perceptual contrast, but perceptual contrast can also be increased by changes in the stimulus, in the "hard-wired" physical features of a participant, or in other pre-attentional processing. Thus, an account of subject-level activity solely based on introspectible qualities will not be able to fully separate the contribution of the subject from other processes that bring about those same qualities, as demonstrated through the example of attention.[10]

Felt Activity Instead of relying on introspectible qualities we could follow James in turning to the other feature of subjective phenomena that may reflect the contribution of the subject – felt activity. Returning to the example of attention, when Edmund Husserl, for instance, writes about attention and how it relates to conscious experience, he describes attention in its active[11] sense: "In general, *attention* is a *tending of the ego toward an intentional object*" (Husserl, 1975, p. 80). For Husserl, attention is the activity of the "ego" or subject in favoring an intentional object (where this attentional activity occurs prior to the subject's corresponding motor response) (Husserl, 1975, p. 85).

Although attractive in fitting our working definition of "subject-level prioritization," an apparent difficulty for Husserl's understanding of attention is that something can capture or grab attention, as with the laughing woman above. That is, the example of the laughing woman shows that attention need not be felt as active but can

[10]See Block (2010) for further discussion of Carrasco's work. Unlike Block I find the effect described by Carrasco to be illusory because the participant takes the change in contrast to derive from the stimulus instead of from their own attention. For example, a 20 percent contrast stimulus and a 23 percent contrast stimulus might be illusorily perceived to have the same value of contrast with no contribution from the subject, rather than differential contrast with a contribution of approximately 3 percent added by the subject's attention. More on this below.

[11]Whereas my sense of "activity" in the phrase "subject-level activity" is that of "happenings" or "processes," Husserl's sense of "activity" is that of "volitional happenings" or "volitional processes." I use "activity" in this more general way so as not to confuse the two senses of the word ("happenings" and "volitional happenings"), preferring to use "subject-level" to dictate the type of happenings in question.

sometimes be experienced as relatively passive. Husserl was aware of such cases and says that where a stimulus seems to intrude upon consciousness the activity of the ego is in "receptivity" (Husserl, 1975, p. 79). Husserl thus broadens the concept of activity to include mere receptivity or acceptance of a stimulus (Husserl, 1975, p. 60). In the example above, Husserl might say that the reader actively allows the image of the laughing woman to distract the reader from the text.

Thus, Husserl's specification of subject-level activity includes even the mere acceptance of a stimulus. Because this pushes the boundary between active and passive into the realm of what we hitherto might have called "passive," Husserl says that we must accept the existence of a "deeper passivity":

> Thus *the distinction between active behavior and passive acceptance or suffering does not have the same meaning for naïve consciousness, turned directly toward pregiven objects, as it does for the reflective regard* which already finds in acceptance of the pregiven, in contemplative apprehension of it, an element of activity and, as a result, must obtain a *more radical conception of passivity* than that entertained by naïve consciousness. This [more radical] conception is that of *pure affective pregivenness*, of *passive belief in being*, in which there is nothing yet of cognitive achievement: the mere "stimulus" which proceeds from an existent in the environing world, as, e.g., the barking of a dog which "just breaks in on our ears," without our previously having given our attention to it, without our having turned toward it as a thematic object. Wherever it is a matter of attention, such an activity of the lowest level is already present. (Husserl, 1975, p. 60)

This passage is difficult, but I interpret Husserl to be claiming that the naïve observer separates active from passive by separating what the subject wills from what the

subject willfully accepts, whereas we should separate active from passive by separating what the subject willfully accepts from what is merely presented to the subject by his or her environment, which is the result of a truly passive process. That is, I take the final line to read: "Wherever it is a matter of attention, such an activity of [turning toward a stimulus as a thematic object] is already present." It seems that Husserl considers conscious experience outside of attention to involve a system of salience, such that prominent items can "pull" at the ego, but "apprehension" of those items requires that the "ego" turns to them.[12] Thus, in Husserl's picture, the problem of captured attention shows us that the activity of attention can include anything from willed activity to willful acceptance of an otherwise passively given stimulus.

Although Husserl's answer to the problem of captured attention has merit, the account leaves us with two significant hurdles in our attempt to specify subject-level activity. First, the extension of activity into the realm of what was hitherto considered passive makes activity (in Husserl's sense) more difficult to identify. We normally use the experience of willing as a guide to the presence of activity, but for willful acceptance the experience of willing is either absent or so weak as to be difficult to detect. The activity of acceptance is, at best, a subtle experience that I gather many will have a hard time distinguishing from the passive experience of givens (supposing the possibility that such passive experiences occur). Second, and more problematic to this tack, is that phenomenal experience is full of illusions of inherent saliency, or illusions that the source of saliency in a particular case is the stimulus rather than the subject. The phrase "love is blind," for instance, arguably points to the illusion that the prominence of a loved one in our mental life is due to his or her merit rather than to our loving attention. Similarly, relative perceptual contrast can be illusory

[12]Husserl prefers to describe this pull of the bare stimulus on the ego with the phrase "relief of salience" (Husserl, 2001, p. 215), but his translator calls it "a passive attention" (Husserl, 2001, p. xlix).

when we fail to track our attention, as revealed in the experiments by Carrasco et al. mentioned above. In these experiments, two gratings of differential contrast are subjectively perceived to have equal contrast by some number of participants with working perceptual systems. One explanation for this is that attention changes relative perceptual contrast by increasing the contrast of the attended stimulus with respect to the unattended stimulus (Carrasco et al., 2004; Treue, 2004). However, since we do not normally confuse changes due to subject-level activity with changes in the received stimulus, the difference between the normal case and the illusory case is that in the illusory case we perceive the perceptual changes as belonging to the stimulus. Whatever the explanation behind these failures to track subject-level activity, they reveal the need for information about the received stimulus apart from what we can gather through our own subjective phenomena. That is, Husserl's account will have to be supplemented if we wish to find a clear and consistent divide between subject-level activity and pre-subjective processes.[13]

2.3 The Behavioral Search

Although the example of attention reveals that phenomenal evidence will be insufficient for a clear and consistent specification of subject-level activity, we can supplement phenomenal with behavioral evidence. One reason we might take this approach is to better detect illusions of saliency that proved problematic for the phenomenal approach: an external perspective allows us to separate the objective stimulus from the participant's response to the stimulus.[14] In psychology this has led to research on the

[13]It is possible that these illusions of inherent saliency result from a failure to track processing outside the realm of the subject rather than from a failure to track attention. As I point out in the next section, behavior does not parse the difference between subject-level activity and other processing by the participant. If this alternative explanation is the right one, this "hurdle" collapses, leaving only the first.

[14]I purposefully change the language from "subject" to "participant" here because the latter, which is typically used by scientists to refer to the experimentee, need not indicate subject-level

difference between stimulus-driven, or "exogenous," behavior and internally-driven, or "endogenous," behavior. Such research successfully reveals behavioral markers that separate responses driven by changes in the stimulus from those driven by changes in the participant, which I review below. However, this research is generally assumed to represent not simply the difference between participant involvement and stimulus influence, but the difference between voluntary and reflexive processes, where voluntary processes are considered to be those within the subject's control.[15] Herrmann et al. (2010), for example, declare: "We found that attention, both exogenous (involuntary) and endogenous (voluntary), can affect performance by contrast or response gain changes, depending on the stimulus size and the relative size of the attention field." In the following paragraphs I show that endogenous processes reflect a combination of voluntary and non-voluntary influences. To do this, I review the best available behavioral evidence for specifying subject-level through endogeny and conclude that this evidence is not sufficient to divide the contribution of subject-level activity (what these researchers would call "voluntary activity") from the contribution of other aspects of the participant, such as the participant's "hard-wired" features.

ERP, or Two-b or Not Two-b? One way that researchers have found differences between endogenous and exogenous influences is through comparison of the electric potentials present on the surface of the scalp for particular tasks, or Event-Related Potentials (ERPs). By averaging the recordings from participants listening to "oddball" tones and then subtracting those of frequently repeated tones, one can create an ERP waveform that isolates the impact of the salient stimuli. That is, the disruptions in the waveform reveal the difference between the unusual and the usual

involvement.

[15] Although in the next chapter I define "voluntary" much more weakly, most researchers mean by "voluntary" either volitionally directed or volitionally allowed, which falls within my understanding of subject-level activity. I am using the stronger sense of "voluntary" here.

sounds, where no difference would result in a flat horizontal line. If one were to add a task, such as having the participant count the oddball tones, one could look at the differences between these endogenous ("count" condition) and exogenous ("read" condition) waveforms.

Using this method, C. Bottcher-Gandor and Peter Ullsperger find a pattern of activity *common* to both the count and read conditions, which is called the "N1-P1 complex." This is the first positive peak in the waveform (P1) and its negative "deflection" (N1) (see figure 2.2), where positive and negative are measured with respect to the recordings for the frequently repeated tones. Bottcher-Gandor and Ullsperger determine that this complex can be reduced, but not eliminated, when participants ignore the stimulus. This shows that the neural underpinnings for the N1-P1 complex can be modulated[16] by changes in subject-level activity (i.e. the peaks can be reduced by ignoring the stimulus), but are nonetheless partly independent of subject-level activity (i.e. the peaks cannot be eliminated by ignoring the stimulus). Consequently, Harold Pashler asserts that the N1-P1 complex results from *both* features of the stimulus and the influence of attention: "there seems to be little doubt that these changes reflect modulations in the magnitude of underlying exogenous components as a function of spatial attention" (Pashler, 1999, p. 75).

Perhaps more importantly, Bottcher-Gandor and Ullsperger find that the count condition involves a component *absent* from the read condition:

> In the count condition, where an N2b is almost certainly present and occurred together with a subsequent positive deflection, the peak latency of the negativity in the difference waveform was revealed to be significantly longer than in the read condition....This again suggests that in the read

[16] "Modulation" is the change of one wave's (or signal's) properties through the interference of another wave (or signal).

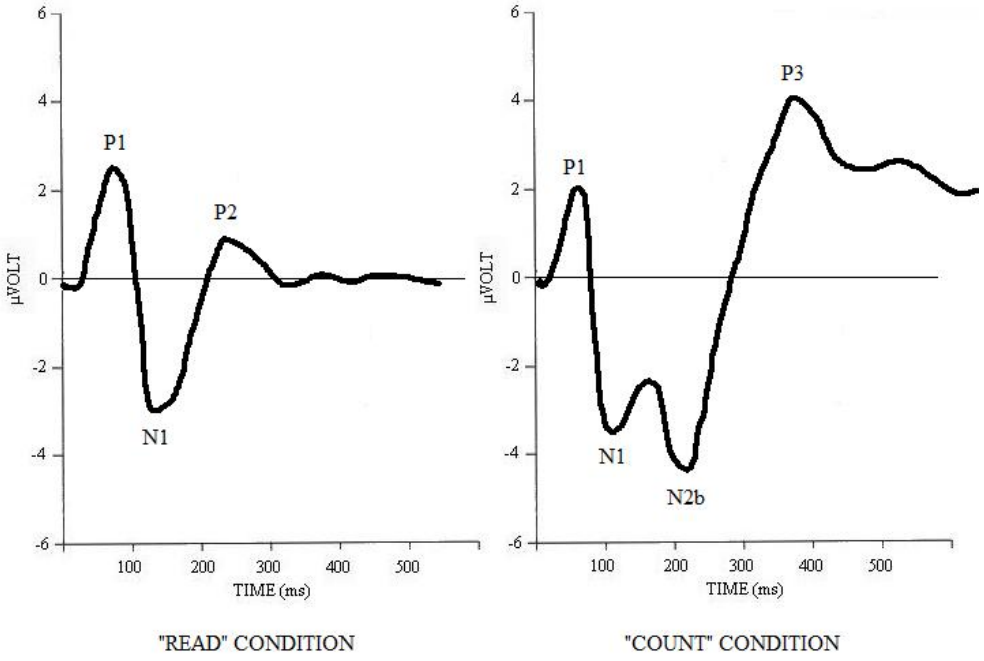

Figure 2.2: Approximate ERP Waveforms

condition, the N2b was not present. (Bottcher-Gandor and Ullsperger, 1992)

The N2b is the second negative peak in a waveform, where the "N2b-P3 complex," as it is sometimes called, is the combination of the N2b and the positive P3. Thus, the N2b-P3 complex is considered to be a marker of endogeny, or of the voluntary influence of the participant.

These studies show that when we add a task to passive listening, the waveforms of the passive listening condition and the task condition will have marked differences, such as the presence or absence of N2b, seen above. These studies do not show, however, that the difference between these conditions is one of subject-level activity, since it could be related to other internal processing required by the addition of a counting task. That is, the counting task requires more internal processing than the passive listening task does (because it is more difficult). This difference might be sufficient to explain the presence of N2b in the count condition, and its absence from

the read condition. Moreover, studies examining the relationship between ERP and musical sequences have found that ERP can differentiate unstructured sequences of notes from structured melodies, but not subjectively meaningful or pleasant melodies from randomly generated ones (Minati et al., 2008, 2010). This points to the possibility that ERP can reflect only processing power and time course but not anything so particular as subject-level activity. At least, the behavioral evidence appears unable to separate subject-level activity from other aspects of participant involvement, such as increased internal processing.

RT, or How Fast Is Too Fast? The reaction time to a target (RT) is likewise thought to reflect a difference between exogenous and endogenous influences (Yantis, 1998). For one set of experiments on RT, two different kinds of cue are used to indicate the location of a target while the gaze is fixed at a central point.[17] The "central cue" is presented at fixation, symbolically directing the participant to expect the target either to the left or to the right of the central point (e.g. with an arrow). The "peripheral cue," on the other hand, directs the participant to expect the target at the cue's location, which is either to the left or to the right of the central point. In figure 2.3, for example, the cues direct the participant to the right side of fixation. The difference between the RT following a peripheral versus a central cue is meant to tell us the difference between stimulus-driven (exogenous) expectation and participant-driven (endogenous) expectation.

The RT experiments use two different conditions to measure the impact of these cues on RT. In one condition (the "valid condition") the participant is cued to the correct target location. In the other condition (the "invalid condition") the participant is cued to the location on the opposite side of fixation from the target (see figure

[17]Gaze is verified in these experiments through either eye tracking or tasks that are shorter than the time it takes to shift gaze (around 100 ms).

2.3). If the participant expects the target at an invalidly-cued location, RT to the target should be longer than if the participant expects the target at a validly-cued location or does not expect the target at any location in particular. The RT difference between peripheral and central cues in the valid condition is meant to illustrate the difference in processing speed for exogenous and endogenous cues. The RT difference between peripheral and central cues in the invalid condition is meant to illustrate the difference in strength of influence by exogenous and endogenous cues.

In a 1981 study by John Jonides, the RT is shortest when the target is preceded by a valid peripheral cue, or when the target is correctly predicted by a cue in the same location as the target. The RT is longest when preceded by an invalid peripheral cue, or when the target is incorrectly predicted by a cue on the opposite side from the target. The central cue brings about a longer RT in the valid condition (i.e. longer than the peripheral cue) but a shorter RT in the invalid condition (i.e. shorter than the peripheral cue) (Jonides, 1981). Because the central cue does not hinder expectation much in the invalid condition, it appears to have less of an overall effect on expectation (it is less "potent"). Because the central cue results in longer RT's in the valid condition,

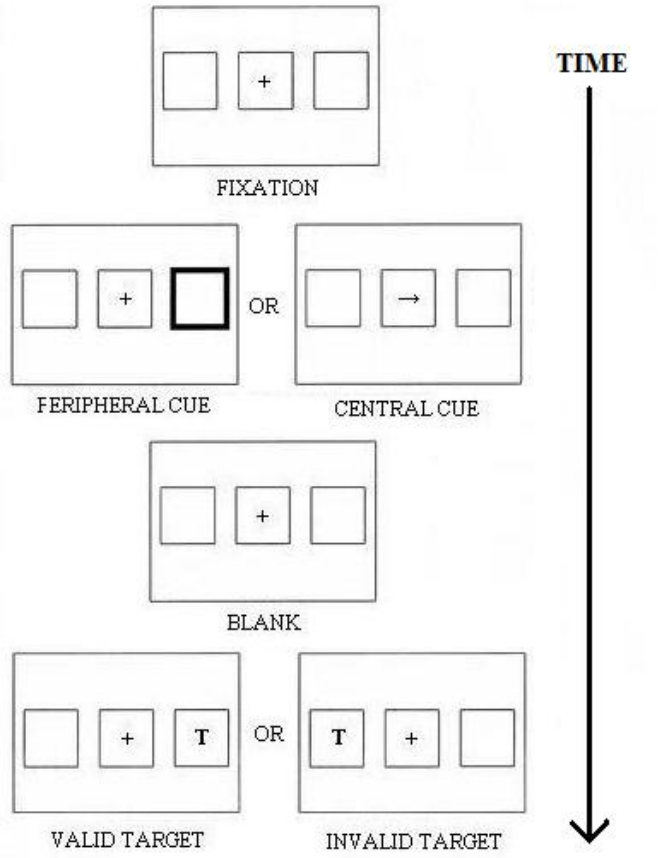

Figure 2.3: The RT Experiment

it appears to engage expectation more slowly. Michael Posner reviews more recent literature on these experiments and finds that the peripheral cue increases sensitivity at the expected target location at around 50 ms after cue onset, whereas the central cue does not bring about this sensitivity until much later.

Although these RT findings illustrate a difference between peripheral and central cues, and perhaps between more and less internal processing, it is not clear that the difference derives from subject-level activity. A possibility mentioned by Jonides is that the central cue has to be processed more deeply than the peripheral cue, where this processing could occur outside of the purview of the subject; perhaps the central cue is processed more slowly because it requires symbolic processing in brain areas some distance from the visual areas required to process the peripheral cue. As with the ERP studies, the RT experiments do not successfully separate the effect of subject-level activity from that of other internal processing, and so we cannot conclude from these findings that the slower RT times represent the slower and less potent involvement of the subject.

IOR, or How to Avoid Being Turned Into a Pillar of Salt Inhibition of return, or IOR, is another type of evidence thought to distinguish endogenous from exogenous influences on behavior. IOR studies use the same paradigm as RT studies but look instead at the *inhibition* of target detection by peripheral and central cues. Both RT and IOR experimental setups involve central gaze, central and peripheral cues, and valid and invalid conditions, but in IOR experiments the time between cue and target onset (the "Stimulus Onset Asynchrony," or SOA) is varied (see figure 2.4). As I discuss below, one of the important findings from this research is that peripheral, but not central, cues invoke inhibitory responses at the cued location after around 250 ms.

Michael Posner and Yoav Cohen published a paper in 1984 showing that peripheral and central cues yield different inhibitory effects. Namely, when the time between cue and target is increased beyond 250 ms, the initial advantage of peripheral cueing is reversed, whereas the advantage of central cueing is unchanged (Posner and Cohen, 1984). After this time, it appears as though the peripherally-cued location is inhibited, giving the advantage to the uncued location. That is, when there is more than 250 ms between a peripheral cue and its target, response to the target is faster for the uncued location than for the peripherally-cued location. As Posner and Cohen write: "These results show that facilitation can be obtained either from peripheral or from central cues. However, the inhibition effect does not occur if the cue is a central one" (Posner and Cohen, 1984, p. 541).

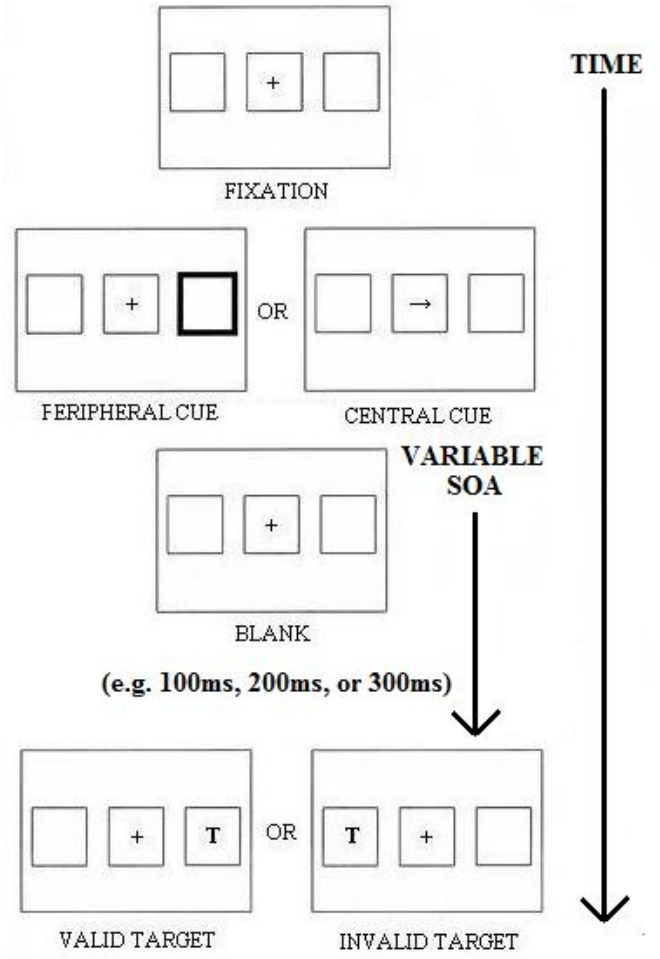

Figure 2.4: The IOR Experiment

Although it is tempting to attribute any inhibitory effect to subject-level activity, additional research seems to show that IOR represents the contribution of non-voluntary oculomotor planning, instead. A review by Bruce Millikan and Steve Tipper, for example, presents research showing that IOR *does* take place for central cues when the par-

ticipants are allowed to move their eyes and even when eye movements are planned and then cancelled by the participant. That is, if participants fixate and are cued at center, but plan to direct their eyes to the cued location, that location receives an inhibitory influence at around 250 ms, just as though it were peripherally cued. However, this inhibitory effect occurs even when the participant attempts to override it by "cancelling" the movement, showing that the effect is not voluntary. This indicates that IOR may be connected to "the automatic programming of an eye movement to its location" (Milliken and Tipper, 1998, p. 214).[18] Thus, IOR does not separate subject-level activity from the influence of the stimulus so much as the planning of eye movements from the absence of such planning, where the IOR effect of this oculomotor planning represents forces outside of the subject's control.

All the examples given above illustrate differences between endogenous and exogenous influences on behavior.[19] Although the experimenters themselves may not have held this view, many interpret this difference between endogenous and exogenous influences to reveal a difference between voluntary and reflexive processing. I take it that while the differences between exogenous and endogenous influences are genuine differences that are interesting in their own right, they need not stem from a difference between the influence of subject-level activity and that of the stimulus. I

[18] A more recent review by Raymond Klein also cites oculomotor planning as the likely cause of IOR, though it notes a more complex relation between oculomotor planning and IOR than that suggested by Millikan and Tipper (Klein, 2000).

[19] I purposefully leave out fMRI (functional Magnetic Resonance Imaging) evidence because I find the relevant evidence inconclusive. Namely, while conducting the RT paradigm in an fMRI scanner, Allyson Rosen et al. found that both the exogenous and endogenous conditions activate a common neural network (echoing other studies (Nobre et al., 1997)) but that the endogenous condition also activates areas associated with planning and working memory (Rosen et al., 1999). However, in a later study, Marius Peelen et al. found no such difference between the endogenous and exogenous conditions. Peelen et al. suppose that one reason for the discrepancy is that the Rosen et al. study uses a "block" design, where the endogenous and exogenous conditions are run on separate occasions. This design is more open to error since the difference found in the separate runs may be due to "expectation, arousal, effort, and mnemonic, behavioral, and other demands" unconnected to subject-level activity (Peelen et al., 2004, p. 823).

have discussed in each case why the evidence offered does not indicate such a difference, but there is a general reason why endogeny does not line up with subject-level activity: subject-level activity is just one of the many "internal" drivers of behavior captured by endogeny, which also include emotion, arousal, and associative processing. Thus, separating endogenous from exogenous behavior will not be sufficient to separate subject-level activity from other processing by the participant.

This problem with specifying subject-level activity through endogenous behavior is one example of a more general problem that applies to any purely behavioral study. Namely, all behavior necessarily combines the influences of physical propensity (the hard-wired features of the participant), subject-level activity, and other internal processing. Even if one could hold pre-subjective internal processing constant, one would not be able to hold the hard-wired features constant across participants in every case because these depend on the history (i.e. the skills and capacities) of the participant. The orienting response to one's name, for example, is both reflexive and personal. Thus, not only is behavior sometimes a result of internal factors outside of subject-level activity, it is sometimes a result of hard-wired features, which vary across participants.

Thus, behavioral evidence, even when added to phenomenal evidence, will be insufficient for specifying subject-level activity. In the next section I attempt to resolve these difficulties by adding a third type of evidence: neural evidence.

2.4 The Neuropsychological Search

Recall that the purpose of turning to behavioral evidence was to supplement the phenomenal evidence for hard cases. That is, the phenomenal evidence allows for easy identification of subject-level activity when there is a high degree of felt activity but leaves us unsure at lower levels. However, the debates concerning the

relationship between attention and other mental processes often spin around what happens at these lower levels. Thus, to allow for discussion on these issues, we will have to supplement the phenomenal evidence. Supplementing the phenomenal with behavioral evidence, though, leaves us with a new difficulty. Namely, the behavioral evidence allows for easy identification of subject-level activity when pre-subjective processing is held constant, but leaves us unsure in all other cases.

Luckily, given a certain set of assumptions about the subject, we can supplement these two types of evidence with a further type to complete our specification of subject-level activity. Namely, if we assume that the subject can operate outside of conscious experience and that there are correlates of subject-level activity in the brain, then neural evidence can complete the picture left in between phenomenal and behavioral evidence. That is, where phenomenal evidence captures only subject-level activity as found in experience and where behavioral evidence captures only subject-level activity mixed in with other aspects of the participant, neural evidence can separate subject-level activity from other aspects of the participant. It does not necessarily do this by capturing the particular history of the subject but by separating the type of processing brought about pre-subjectively from the type of processing occuring at the subject-level. The neural evidence will rely on the phenomenal and behavioral "easy" cases to determine the kind of neural activity that is likely to correlate with subject-level activity, after which it can extend these findings to harder cases. Thus, where neuropsychology is the study of elements of psychology through neuroscience (psychology being informed by both reports of subjective phenomena and behavior), we might call any such specification of subject-level activity "neuropsychological," rather than "neural," because it relies on all three types of evidence (phenomenal, behavioral, and neural).

There are currently two promising avenues for separating subject-level activity

from pre-subjective activity using neuropsychological evidence. One neuropsychological divide is that between "global" recurrency (subject-level recurrency, or long-range recurrency involving the planning areas) and either "local" recurrency (recurrency that does not involve the areas associated with the subject, such as the planning areas) or feed-forward sweeps (see, e.g., Spratling and Johnson (2004)). "Recurrency" is a catch-all term for synchrony between processing in different neural areas (Lamme and Roelfsema, 2000). The concept of recurrency is similar to that of the relationship between full two-way radios that are in sync. Imagine, for example, that two friends, David and Jamie, are playing music together from their respective homes via two-way radio. If their two-way radios only allow one transmission at a time (like walkie talkies) then we might call the relationship between their radios "feed forward" or "feed back." David might play a tune on his guitar, for example, while Jamie waits to transmit back a tune on his banjo. Regardless of how quickly Jamie responds, David's tune will have ended. If David and Jamie have full two-way radios, on the other hand, or two-way radios that allow simultaneous transmission and retrieval, then the relationship between the music played over the radios can take on the structure of recurrency. That is, David's tune and Jamie's tune can share air space and, if all goes well, begin to harmonize or "sync." If David and Jamie record an entire song, we might say that the recurrency is "global" with respect to that song. If David and Jamie divide the song into parts, recording each part on an eight-track, we might say that the recurrency is "local" for each part of the song with respect to the song as a whole. One reason to think global recurrency an ideal candidate for specifying subject-level activity is the fact that it involves interaction between multiple brain areas and states, including those that are likely to reflect the goals of the subject, without necessarily involving conscious experience. Thus, if global recurrency properly captures subject-level activity then it will have the advantage of capturing both

conscious and subconscious subject-level activity.

Anatomically, recurrency can be identified by determining the areas of the brain that receive signals from and/or provide feedback to the entire cortex in their domain – such as the Frontal Eye Fields (FEF) or the Lateral Intraparietal (LIP) for the visual cortex – or by determining states that involve these areas (see figure 2.5). The "planning areas" (such as the dorsolateral prefrontal cortex, or DLPFC) are so-called because they have been correlated with executive planning, both in imagination and in action. Because a central role of the subject is to plan and execute activity, global recurrency, or recurrency that correlates with the subject level, can thus be identified by examining either regions that connect the signals from the sensory cortices to the planning areas or by examining states involving these areas. Local recurrency in the sensory areas alone, on the other hand, does not seem attributable to subject-level activity, since the subject is not usually involved in controlling visual norms outside of more general planning.

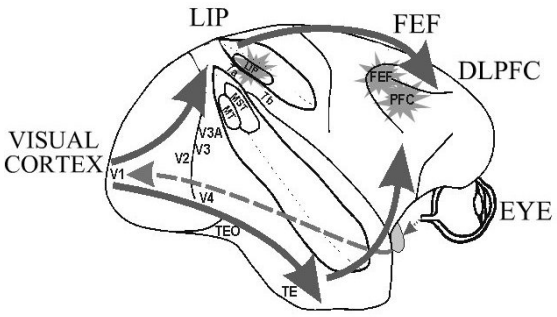

Figure 2.5: Neural Areas

As I will discuss further in Chapter Three through a discussion on mental selection, another neuropsychological candidate for specifying subject-level activity is that of resource-pooling.[20] That is, some processes seem to share neural resources, resulting in "serial" behavioral markers (markers of a single set of available resources at any one time), whereas other processes seems to use resources that are not in conflict,

[20]I suppose here that "resource" is an umbrella concept that includes both biological movers, such as sugar and oxygen, and more general restricted properties, such as space and time. The concept of capacity or resource-limitation inherent in resource-pooling can be found in the work of Daniel Kahneman, as in his Capacity Model of Attention (Kahneman, 1973, p. 13) and Nilli Lavie, as in her Perceptual Load Theory of Attention (Lavie, 1995).

resulting in "parallel" behavioral markers (markers of multiple simultaneous sets of available resources). This difference will be used in Chapter Three to further separate subject-level activity from pre-subjective processing.[21]

It seems likely that these two candidates for specification are related. Long-range (global) recurrency between the planning and sensory areas would result in more signal interference from other long-range processes than short-range (local) recurrency would receive from other short-range processes. That is, whereas local recurrency can occur in parallel across the brain without (much) interference, for this type of recurrency to occur at a level that involves the planning areas, the resources available at that level would be in competition, since this type of recurrency spans the entire cortical structure. Thus, the neural system might take on the character of serial processing, or processing that splits a common pool of resources, for long-range recurrency involving the planning areas.

2.5 Conclusion

In this chapter I have argued that a neuropsychological specification of subject-level activity has clear advantages over its competitors. Namely, neither the phenomenal nor behavioral specifications fully distinguish subject-level activity from other processes. By contrast, the neuropsychological approach provides a means of specifying that activity. As discussed above, subject-level activity mimics changes to the received stimulus, so conscious experience does not always capture the difference between changes in subject-level activity and stimulus variability. Similarly, since behavior does not always parse the difference between the influence of subject-level activity

[21] Importantly, this measure of behavior does not depend on response complexity. Just as subject-level activity need not result in phenomenal clarity and phenomenal clarity can be achieved without subject-level activity, subject-level activity need not result in response complexity and response complexity can be achieved without subject-level activity. Thus, any marker of subject-level activity in behavior should not depend on response complexity.

and that of either the participant's hard-wired features or other internal factors, it cannot reliably capture subject-level activity. By combining phenomenal, behavioral and neural research through neuropsychology, however, we can find a clear specification of subject-level activity, and can thus limit our research on attention to processes that fit into this spectrum. It is thus recommended that attention be understood as subject-level prioritization that takes place through relatively resource-limited recurrency between the planning and sensory areas, or what I will call "subject-level neural governance."

In the next few chapters I discuss the relationship between attention and other mental processes. Namely, in Chapter Three I discuss the relationship between attention and mental selection, in Chapter Four I discuss the relationship between attention and conscious perception, and in Chapter Five I discuss the relationship between attention and conscious experience in general. Finally, in Chapter Six I examine a case study for the claims made in Chapter Four.

Chapter 3

Attention and Mental Selection

ABSTRACT: What is the relationship between attention and mental selection? This question is rarely asked in the philosophical and cognitive science literature, and in some cases the two are simply assumed to be equivalent. In this chapter, I investigate the contribution of attention to mental selection by looking at two forms of such selection: sensory selection and what is known as "selection-for-action." I find that attention is not necessary for selection in either case, but that it can provide for more optimal selection than non-attentional processing. This discussion will serve as a contrast to what I take to be the more direct contribution of attention, that of providing for conscious perception, which I discuss in the next chapter.

3.1 Introduction

The concept of "selection" is more general than that of "attention." Selection is the separation of some entities from other entities based on particular properties. Attention is the subject-level prioritization of some entities over other entities based on particular properties, which is a form of selection. That is, "attention" always involves a two-part relation between the subject and the attended entities, but "selection" can involve the selected entities by themselves, selected simply against the background of the laws of nature. Thus, attention cannot be necessary for selection in general unless there is no selection outside of the realm of the subject.

However, attention may be necessary for all selection that occurs in the life of the mind, or for all *mental* selection, depending on whether all mental processing

involves the subject.[1] That is, if some mental processing is accomplished through either neural anarchy or local governance, then the subject-level neural governance of attention introduced in the last chapter is not likely necessary for all mental selection. To determine whether or not attention is necessary for all mental selection I look at two types of mental selection that are commonly linked to attention, sensory selection and selection for action.

First, at the start of many research papers on attention, one finds a brief statement about how attention enables sensory selection, where selection is there understood as the reduction of what is otherwise a supposedly unmanageable amount of available sensory "information." The following examples illustrate the mantra-like quality of many of these statements:

> At any given moment, our visual system is confronted with far more information than it can process effectively...Visual attention serves as a mediating mechanism... (Carrasco et al., 2004)

> At any given moment, our visual system (and indeed every sensory system) takes in far more information than can be fully processed. Selective attention allows an individual to choose certain subsets of that information to receive additional processing. (O'Craven, 2005)

> At any given moment, our visual system is confronted with more information than it can process. Thus, attention is needed to select behaviorally

[1] By "mind" I mean the entirety of internal drivers of experience and behavior, where the mind includes both subject-level and pre-subjective drivers. An example of a pre-subjective mental driver is one's distaste for certain foods. Oddly enough, I cannot stand to drink milk in the morning if a window is open because the milk then tastes disgusting to me. I take it that this is related to an event in my personal history, but it is neither a physical feature of my body nor something within my (direct) control, and so seems to count as a *pre-subjective* driver.

> relevant information in a visual scene for further processing. (Bichot and Desimone, 2006)

Many other papers start the same way (Kim and Cave, 2001; Chua and Chun, 2003; Busse et al., 2008; Brisson et al., 2009, etc.), all testifying that attention plays an essential role in the reduction of sensory processing. In philosophy, attention has long been linked with selection based upon limited processing capacities, as can be gleaned from Augustine's contrasting description of divine perception in *City of God* (426 CE): "He sees in some other manner, utterly remote from anything we experience or could imagine. He does not see things by turning his attention from one thing to another. He sees all..." (Augustine, 2003, p. 452). According to Augustine, human experience differs from divine experience in requiring selection like that described in the quotes above, which I will call "selection-from-limitation."

Henri Bergson writes of the connection between attention and another form of selection. Bergson claims that "attention to life" keeps the conscious mind focused on a particular activity, selecting both the sensory input and the motor response relevant to that activity (Bergson, 2007, p. 226).[2] Without attention-to-life, Bergson claims that the conscious mind would be disconnected from the living body, likening the conscious mind to the Cartesian soul (Bergson, 2007, p. xiii). Thus, for Bergson, attention-to-life is essential for connecting the unlimited consciousness with the limited body, allowing for conscious life: "To live is to be inserted into things by means of a mechanism which draws from consciousness all that is utilizable in action...and darkens the greater part of the rest" (Bergson, 1920, p. 71). Since Bergson's attention-to-life is required to allow the conscious mind to act through the limited body (Bergson, 1920, p. 59), it is a selection for the sake of action, or what

[2]This is a low-level form of attention according to Lapoujade (2005).

has come to be called "selection-for-action."[3]

In this chapter, I investigate the contribution of attention to selection by looking at both selection-from-limitation and selection-for-action. I find that attention is not necessary for selection in either case, but that it sometimes allows for more optimal selection than non-attentional processing does. Specifically, after discussing the defining features of attention that I argued for in the last chapter, I dismantle the idea that attention is necessary for selection-from-limitation since this type of selection can be provided by pre-attentional processes. I then present and reject the arguments given in favor of the view that attention is necessary for selection-for-action by giving evidence of action that occurs without attention. Finally, I present arguments in favor of the view that attention can provide optimality in both forms of selection as it allows for more flexibility in selection.

3.2 Selection-from-Limitation

As I argued in the last chapter, my understanding of attention first proceeds from the assumption that attention is a subject-level process[4] and that this is an essential part of our use of the concept. We do not say, for instance, that coffee filters attend to coffee or that digestive fluids attend to nutrients; although these filtering processes "select" their respective goods, they are not subject-involving and so do not count as instances of attention. Thus, attention will have to involve more than pre-subjective filtering. At the other extreme, I set aside felt volition as a reliable marker because there are many instances of attention that are not rooted in felt

[3]Bergson says that attention ("attention-to-life") is necessary for connecting the conscious mind to the living body, but not that attention is necessary for all selection-for-action. Thus, my criticisms later in the chapter against the view that attention is necessary for selection-for-action do not apply to Bergson.

[4]Where "subject" stands for that to which we ascribe conscious experience and thought, "subject-level" stands for that which is currently within the control of the subject, or is under the influence of the subject.

volition, such as the turning of attention to an appealing stimulus against the better part of one's will. In between these limits I found a natural division of subject-level from pre-subjective processes in the neural contrast between long-range (global) recurrency and either local recurrency or feed-forward sweeps. I thus recommended that attention be understood as subject-level prioritization that takes place through recurrency between the planning and sensory areas. I call this "subject-level neural governance." Importantly, pre-subjective neural governance, or prioritization based on local recurrency, does not count as attention in my understanding.

Some readers may be concerned that this way of understanding attention conflicts with the usage of the term in some scientific papers. Attention is, for example, in some places attributed to lower-level selection processes, which I call non-attentional. This conflict, unfortunately, precedes my investigation. The use of "attention" in the philosophical and scientific literature is not consistent, and a more determinate understanding is required.[5] Thus, to object to my understanding of attention one must not simply find conflicting usage but show that such usage is superior. The use of "attention" that I have proposed fits a natural division between serial and parallel resource-pooling, where both serial and attentional processes are known to be subject-involving. Thus, even if some other neural marker (besides long-range recurrency between the planning and sensory areas) is later found to better specify attention, the parallel processing contrasted with attention in the experiments to follow will likely still be considered outside the scope of attention because it does not appear to be subject-involving.

In any case, the understanding of attention that I provide makes it clear why

[5]Another approach to this problem is suggested by Alan Allport: "There can be no simple theory of attention, any more than there can be a simple theory of thought. A humbler but also a more ambitious task for the next twenty-five years will be to characterize, in cognitive neurobiological terms, as much as possible of this great diversity of attentional functions" (Allport, 1992, p. 206).

selection-from-limitation need not involve attention. "Selection-from-limitation" is the selection of some sensory processing at the expense of other processing that results from limited processing resources, a concept stated univocally at the start of this chapter. When we restrict attention to subject-level neural governance, then we see that there is no reason to suppose that attention is necessary to distribute limited resources. In the so-called "natural selection" of evolutionary theory, for example, there is no overarching governance of the resources available for survival, but direct, anarchical distribution based on fitness. Even selection under tribal governance, or the local distribution of resources by tribal leaders, lacks the overarching governance across tribes that would be analogous to subject-level neural governance or attention. Thus, non-attentional selection can be based either on anarchical neural distribution or on pre-subjective neural governance via local recurrency. Either type of distribution may correspond to the selection of particular sensory features, such as location and color. That is, the prioritization of the center of one's visual field over the periphery or of the color red over other colors may occur through either anarchic distribution or local recurrency.

To illustrate the most basic form of non-attentional selection think of the sensory system as a system of doors: a door "selects" the size of objects that are able to enter a room with respect to a spatial limit without need of attention. Thus, adult humans are (usually) unable to enter dog doors, and large couches are (usually) unable to enter human doors (even allowing for creative angling and Herculean efforts). Similarly, the eyes non-attentionally select certain properties of objects in a region of space at a particular acuity that they pass up the sensory chain. The visual cortex uses an analogous type of non-attentional selection when groups of neurons respond preferentially to certain types of visual stimuli, whether by responding exclusively to

those stimuli or by biasing them over other available stimuli.[6] In short, attention is not required for selection-from-limitation, since this can occur without subject-level neural governance, or prioritization by the subject.

3.3 Selection-for-Action

Odmar Neumann and Alan Allport introduced selection-for-action as an alternative concept to selection-from-limitation (Neumann, 1987; Allport, 1987). According to Neumann, the concept of what I call "selection-from-limitation" is based on the purported existence of function-general limitations on neural processing, or processing "Capacities," where the concept of selection-for-action is based instead on the idea of function-specific limitations: "independent of all Capacity considerations, selection is evidently needed for the control of action" (Neumann, 1987, p. 374). In looking at their concept of selection-for-action, I leave out from this discussion all actions that *necessarily* involve the subject (such as the action of attending) and focus only on actions that are *not* necessarily subject-involving, such as behavioral responses to perceptual stimuli. In this way I hope to avoid confusing the agent (that which is responsible for action in selection-for-action) with the subject.[7]

Neumann claims that the function-specific limitations for such actions (behavioral responses to perceptual stimuli) derive from competition for *effector recruitment* and from underdetermined *parameter specifications*, where these can be related through an example. Imagine that you arrive home from work after a long commute and want to change into more comfortable clothing, get a glass of water, put your belongings away, return some phone calls, and relax on the couch. As the effectors required for these actions are limited – the hands, for instance, cannot both pour a glass of

[6]Wayne Wu makes a similar point by comparing non-attentional selection to the sorting of a gumball machine (Wu, 2011a, p. 159).

[7]As stated in the Introduction, I owe the division of these to Ned Block.

water and pull off an uncomfortable sweater without likely disaster – the actions must be prioritized, whether consciously or subconsciously. Thus, due to limitations on effector recruitment, you will have to select one action to perform first. Once the actions have been prioritized, you must select parameter specifications, such as specific movements of the hands. That is, the parameters of motion will have to be specified (either consciously or subconsciously) for the hands to perform an action. For Neumann, these two types of selection must be accomplished in order for you to act, where action is understood as the matching of a goal with particular behaviors in order to achieve that goal.

Allport adds to Neumann's requirements of effector recruitment and parameter specification the requirement of perceptual-motor coupling:

> The need for such a mechanism (of selective coupling and decoupling of perceptual and motor processes) arises directly from the many-to-many possible mappings between domains of sensory input and of motor output within the very highly parallel, distributed organization of the nervous system. (Allport, 1987, p. 397)

That is, the fact that neural processing in the sensory and motor areas takes place in parallel (i.e. without interference or communication) means that some mechanism of coupling between these areas is required to achieve a motor response that is based on the perceptual stimulus. In other words, even if the agent[8] limits itself to one action per effector at a time and sets the parameters for action, there still has to be a determination of which perceptual stimulus to act upon (from the perceptual

[8]Again, by "agent" I mean the source of action at a given time, which may or may not be the subject. That is, I am leaving it open here whether selection-for-action needs to be solved by the laws of nature, neural anarchy, local neural governance, or subject-level neural governance. I am leaving open, for instance, the possibility that the subject normally coincides with the agent but that action is nonetheless sometimes controlled by lower-level mechanisms.

"inputs" represented in the sensory areas) and then this stimulus must be coupled to the selected motor response (from the motor "outputs" represented in the motor areas). This is not merely to say that *successful* implementation of an action requires perceptually selecting a stimulus and coupling it to the action; it is not the idea that, for example, the chosen stimulus needs to be the source of parameter specifications so that, in the case of reaching for a glass, you do not accidentally direct your grip to the wrong location and knock the glass to the floor. Allport's idea is that the failure to act *at all* rests on the failure to connect percept to response, since action is based in neural processing and these are not processed together in the brain. In other words, Allport claims that action of the type in question definitionally requires a one-to-one relation between percept and response, but not necessarily a perfect coupling between the stimulus and the resulting behavior.

Although it seems clear that selection-from-limitation can take place without attention, the story with selection-for-action is more complicated. The reason for this complication is that selection-for-action is required for the goal-based coupling of percept to response, rather than a simple reduction in processing, and it is not obvious that this can be provided by a system of pre-subjective filters. To discover the relation between attention and selection-for-action I will first look at reasons offered by Alan Allport and Wayne Wu to think selection-for-action is linked to attention before reviewing evidence that they can be separated.

Linking Attention to Selection-for-Action Allport localizes both attention and selection-for-action at agentive control:

> The observable criterion for successful "attention" to or awareness of an environmental event depends on the subject's ability to *act* voluntarily

> (i.e. arbitrarily, selectively) in response to that event...In effect, then, this criterion is the same as for selection-for-action. (Allport, 1987, p. 414)

Note, first, that Allport's sense of "subject" is that of "experimental subject," which can be equivalently expressed by "participant." If we thus call the acting participant the "agent," what Allport is claiming here is that attention and selection-for-action both rely on the agent's ability to control the goal-based coupling of percept to motor response; if this coupling is fully pre-determined then it is not controllable by the agent and is therefore not a candidate for attention or selection-for-action. Allport is thus committed to the requirement of voluntariness[9], or the *possibility* of control through a minimal degree of freedom, for both attention and selection-for-action.

This condition for the possibility of attention and selection-for-action does not entail that attention is necessary for selection-for-action. Even if both attention and selection-for-action require voluntariness, this common requirement does not entail that selection-for-action requires attention. In the picture of attention I have given so far, attention requires the *volition* of the subject. The degree of volition can be minimal, as in Husserl's mere acceptance of a stimulus, but is nonetheless distinct from the absence of volition. Unless selection-for-action, or the goal-driven coupling of percept to response, likewise requires the volition of the subject, it may not be coincident with or rely upon attention.

Wayne Wu, however, makes the stronger claim that attention is necessary for selection-for-action in "Confronting Many-Many Problems: attention and agentive control." Wu's "Many-Many Problem" is that of reducing the many available targets and many-available responses to a single action, which corresponds to Neumann and Allport's concept of selection-for-action. As Wu describes the problem,

[9]Note that this sense of "voluntary" is weaker than that of Chapter Two.

> Action requires that the Many-Many Problem be solved by reducing the many-many set of options to a specific mapping between target and response. Throughout the execution of action, the agent must continue to perceptually select, and hence attend to, relevant information so as to guide the execution of specific movements. Since perceptual attention is a necessary part of solving the Many-Many Problem, it is a necessary part of action. (Wu, 2011b, p. 50)

In order to justify his claim that action requires a solution to the Many-Many Problem Wu uses a thought experiment: he imagines that creatures without the Many-Many Problem would have pre-set responses to stimuli and would thus respond reflexively. Wu claims that reflexive behavior does not fulfill the minimum requirements of agency, and so these creatures do not act.

> To see the necessity of the Many-Many Problem for agency, consider a world whose creatures do not face the Problem. The presentation of possibilities is denied them. To the extent that they exhibit bodily behaviors in response to the environment, this must be driven by preset one-one mappings between stimulus and response....If these creatures are in possession of a variety of preset stimulus-response mappings, they may exhibit a certain complexity in behavior over time. Nevertheless, their behavior does not count as action for they are driven by what are essentially a set of reflexes, and these, I take it, never exemplify agency. (Wu, 2011b, p. 54)

Wu thus closes the possibility left open by Allport above; for Wu, all behavior is either fully determined reflex or involves selection-for-action, and therefore also attention.

It may seem as though this last bit of reasoning moves too quickly from "selection" to "attention." However, according to Wu, perceptual selection just *is* attention, and so solving the Many-Many Problem with perceptual selection means solving it with attention: "Since visual, and broadly perceptual, selection of an object and property among others just is a form of *attention*, attention is required to identify a target" (Wu, 2011b, p. 53). This claim is only true, in my understanding of attention, if the subject necessarily coincides with the agent of bodily action. If the agent of bodily action sometimes selects a perceptual target without the involvement of the subject, then attention need not be involved in solving the Many-Many Problem. Either Wu understands attention differently than I do or he is assuming that the subject and agent are inseparable (at least in the case of bodily action). In my understanding of the term, "attention" should not be applied to selection outside the scope of the subject. If this is correct and if the agent is separable from the subject, then some operations of the agent, such as selection-for-action, can occur without attention.

I think that it is by confusing the agent and subject that Wu goes wrong in thinking that there is no behavior between reflex and attentive action, and, thus, no action outside of attentive action. Wu's position in a later paper makes this assumption more striking: "The relevant conception of action, in the first instance, is *intentional* action" (Wu, 2011a, p. 160), where intention is defined in a footnote as "attunement of subjects towards certain inputs and outputs given their goals" (Wu, 2011a, p. 186). Note that Wu also uses the term "subject" in terms of "experimental subject," which can be meaningfully replaced with "participant" or "organism," instead of the more narrow "that which consciously experiences, perceives, and thinks." With that note in mind, this definition of "intention" undercuts Wu's previous claims since attunement does not require the involvement of the subject (in my sense) or attention. In fact, this statement seems to fit a much weaker account of action – as goal-based

coupling of percept to motor response directed by pre-subjective mechanisms that are neither necessarily conscious, attentive, nor within the subject's control. Imagine, for instance, a hungry sleepwalker who makes his or her way to the kitchen to prepare a snack. This behavior, although it results from attunement between the sleepwalker's intention to get a snack and snack availability, is not available to the subject's control (or is at least not within the subject's current control). Wu claims in the earlier paper that such "habitual" actions are either entirely non-intentional (in the sense of "intention" sketched out above) or are reflexive at a rough-grained level that will still require attention for fine-grained coupling (Wu, 2011b, p. 67). As I will discuss below, this move seems ill-supported by the behavioral evidence; behavior like sleepwalking shows that we need not be forced into the position of accepting the reflex/attention dichotomy (that all behavior is either reflex or involves attention) as put forward by Wu.

As a means of fleshing out behavior that lies between attentive action and reflex, consider the following example. Rosa is sailing with her friend Justin, engrossed in conversation, when she suddenly notices that she is on a collision course with another boat. Rosa is sailing upwind on a port tack (where "tack" is defined with respect to wind direction – see figure 3.1) and so knows she should give way to the other boat, which she can see is on a starboard tack. Rosa decides to come about, such that she will sail in parallel to the other boat. That is, she will move her boat in the counterclockwise direction indicated by the darkly-shaded area in the image. ("Coming about" on the starboard tack would involve moving in the clockwise direction indicated by the lightly-shaded area.) To come about, Rosa quickly pushes the tiller toward the

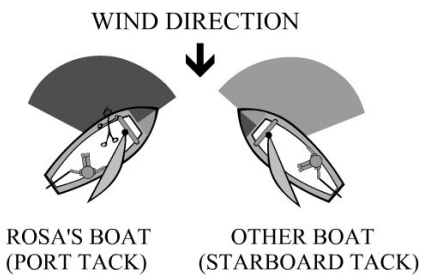

Figure 3.1: Tack Motion

sail until the boom swings over the centerboard, and then switches to the other side of the boat, sailing away into safety. I use this example, in part, because Rosa's action is enough complicated that most people would need to use conscious guidance and attention to achieve it. However, by orienting themselves to the direction of the wind, experienced sailors can accomplish even these difficult and complex behaviors without conscious guidance and, arguably, without attention.

Rosa's action of changing the direction of the boat could be any one of four (or more) types. If Rosa uses conscious guidance and attention to change the direction of the boat we might call her behavior "consciously attentive action." It is conceivable that she could instead unconsciously but attentively decide to come about, in which case we might call her behavior "unconsciously attentive action." Either of these should be considered within the purview of what I have so far called "attentive action," but there are two other options. Imagine that Rosa's conversation with Justin is so engaging that she unconsciously and inattentively decides to come about. Perhaps Rosa is adept enough at sailing that this behavior does not rely on attention. In this case we can stipulate that Rosa's behavior is still within her voluntary control, such that if the wind shifts or if the other boat moves in an unusual way she will seamlessly take over the attentive direction of her behavior (either by volitionally-directing or volitionally-accepting it). Thus, the third type of behavior is goal-based coupling of percept to motor response as directed by pre-subjective mechanisms that are neither necessarily conscious nor attentive but are within the subject's control, which we might call "voluntary action."[10] In the fourth case, when the wind shifts or the other

[10]As previously mentioned, this sense of "voluntary" is much weaker than what is used by most action researchers. It means that there is a possibility for control over one's behavior in the sense that the behavior is not pre-determined and that some parameters must still be set either by the subject (in which case the behavior is "volitional") or by a pre-subjective agent. It does not mean that the behavior is *within* the control of the subject or that it is under the influence of the subject, which is what defines the subject-level.

boat moves in an unusual way, Rosa remains on auto-pilot, and the matching of her behavioral response to the percept is not within her control. Nonetheless, she is capable of acknowledging the goal to come about as her own, even if the coupling of this goal to her behavior was achieved by some pre-subjective mechanism outside of her control. In retrospect, she may even take responsibility for the movement of the boat as something she would have deliberately done, had she been paying attention. Thus, the fourth and final type of behavior is goal-based coupling of percept to motor response as directed by pre-subjective mechanisms that are neither necessarily conscious nor attentive nor within the subject's control, but which are accepted by the subject as his or hers. We might call this sort of behavior "proprietary action."

It may seem counterintuitive to assign the label of "action" to all four types of behavior, but this weakened notion can be found in the work of many action theorists, such as that of Patrick Haggard (Haggard, 2001), and fits some popular accounts of moral responsibility, such as that of Harry Frankfurt (Frankfurt, 1969). Non-action, in the weakest case, would include any behavior that is not based on a goal of the subject or is not accepted by the subject. At the very least, voluntary behavior seems a reasonable candidate for action, as Allport would presumably agree (based on his reasoning above).

Aside from this intuition-based or conceptual evidence, there is empirical evidence that attention is not necessary for selection-for-action, or the goal-based coupling of percept to motor response. Namely, experimental participants can be trained to do goal-directed tasks without diminishing the behavioral resources available to other tasks. In the next section I will present reasons to think that attention is not required for these tasks, but that they are examples of non-attentional action.

Selection-for-Action without Attention The claim that selection-for-action relies on attention is likely weaker than the claim that it relies on conscious guidance, since some instances of attention do not appear to involve conscious experience. Thus, before we look at evidence in favor of rejecting the weaker claim that selection-for-action depends on attention, it may help to understand the reasons behind rejecting the stronger claim that conscious guidance is necessary for selection-for-action, where these reasons can be illustrated through the phenomenon of blindsight.

Recall, first, that selection-for-action is necessary to translate the multitudes of perceptual targets and motor responses to a one-to-one mapping with respect to a particular goal. The phenomenon of blindsight, or the functional sightedness of an experimental participant who reports (or otherwise indicates) the absence of conscious visual perception, reveals that we can achieve this sort of selection without consciously coupling target to response. Two experimental participants are of particular interest here: Helen, a rhesus monkey who had her striate cortex removed, and TN, a male participant with bilateral striate lesions. Whereas many of the human participants who demonstrate blindsight appear unwilling to initiate action, Helen was trained to navigate mazes in search of currants and chocolate pieces (Humphrey, 1974). TN was the first human participant to navigate a maze on his own, though this maze was much less complex than Helen's and the researcher stood directly behind him as he progressed (de Gelder et al., 2008). Helen and TN are said to exhibit "sophisticated visuo-spatial skills in the absence of perceptual awareness," (de Gelder et al., 2008, p. 1128). Helen and TN are relevant to this section of the paper because they both seem able to match their behavior to particular goals without consciously coupling visual percept to response.

Of course, there is a sense in which Helen and TN act consciously: although they cannot consciously perceive the maze, and so cannot consciously couple percept to

motor response, they consciously *intend* to navigate it. One can experience something similar to this blind intention by closing one's eyes and reaching into a bag of black and white marbles and intending to pick out a black marble. If, by some chance, you are able to pick out a black marble each time you reach into the bag, you will have some idea of what it is like to be a blindsighter.[11] Even in this case, though, the conscious intention to pick out a black marble is neither helping you pick out blackness nor connecting this percept to the behavioral response of grabbing the marble. This suggests that you could pick out a black marble without conscious intention if the goal of picking it out was generated subconsciously. Similarly, the coupling mechanism of the blindsight participant occurs outside of conscious experience, indicating that conscious intention just helps to set in motion a particular goal that the subconscious mechanism aims to fit. Thus, conscious guidance is not likely necessary for the coupling of selection-for-action.

Even if conscious guidance is not necessary for selection-for-action, there is a further question of whether attention is necessary.[12] Dual-task experiments, I argue, suggest otherwise. In a few classic dual-task experiments, William Hirst and his colleagues trained two participants, Mary and Arlene, to copy words from dictation while at the same time reading sentences for comprehension. After training, each participant "read as quickly and comprehended as well on the experimental as on the control trials" (Hirst et al., 1980, p. 111), where the control trials did not include the copying task. Furthermore, "the results of the testing stage provide clear evidence that Arlene and Mary were sensitive to the structure and meanings of the sentences

[11] A similar example was given by Sean Kelly in 2010 during a lecture at Harvard University on the topic.

[12] One might think that the blindsight case also counts against the view that attention is necessary for selection-for-action if one thinks that attention requires consciousness (as for Mole (2008)), but since the most established view of blindsight involves attention (Kentridge et al., 1999) I use other evidence to argue that attention is not necessary for selection-for-action.

they copied" (Hirst et al., 1980, p. 114), demonstrating that they were not merely copying the sentences from dictation, a relatively simple behavior, but grasping their content, a more complex behavior. Hirst et al. reason that Arlene and Mary's copying skills are not automatic[13] because "all versions of the concept seem to share a common assumption: that automatic activities must be relatively simple and routinized" (Hirst et al., 1980, p. 106). They reason that, because of its complexity, the skill must therefore require attention (and hence title their paper "Dividing attention without alternation or automaticity").

Unlike Hirst et al., I find a better explanation of Arlene and Mary's behavior to be that the behavior is non-attentional. First of all, as I discussed in Chapter Two, the divide between subject-level activity (such as attention) and pre-subjective processing (such as automatic processing) does not stand or fall on purely descriptive behavioral measures (such as simplicity). These measures mask the skills and capacities of the participant and, as such, do not indicate the level of subjective involvement. As I say there, even though reading is a complex behavior, it can be achieved pre-subjectively and even automatically for one's own name – perhaps the reading comprehension of Arlene and Mary is a higher-level version of the same. Given this possibility, we should not just assume that Arlene and Mary's behavior is attentional because of its complexity.

Second, there is a more useful behavioral divide for detecting the difference between non-attentional and attentional processes that indicates non-attentional processing in this case: resource-pooling. "Resource-pooling," in my understanding, is a task-relative measure and does not entail the existence of anything like function-general processing limitations of the sort criticized by Neumann and Allport. That

[13] "Automatic" is normally understood to mean that the behavior in question is non-voluntary, or that it cannot be controlled. It fits the notion of pre-determined behavior or reflex (Pashler, 1999, p. 357).

is, my distinction between serial and parallel resource-pooling is not based on the distinction between absolute resource-limitation and unlimited processing, as introduced by Richard Shiffrin and Walter Schneider in 1977 to explain the difference between attention and non-attention. Both of our distinctions are motivated by the empirical finding that there are two distinct types of behavior: one slow and "serial," or limited in the number of tasks that can be undertaken at the same time, the other fast and "parallel," or relatively unlimited in the number of tasks that can be simultaneously performed (Crick and Koch, 1990). However, in contrast to Shiffrin and Schneider's early description, these behavioral findings do not show that there is anything like unlimited resources but that certain types of tasks do not interfere with one another. Similarly, these findings do not show that there is anything like absolute resource-limitation, but that there is interference between certain types of tasks. The limited "resources" that are thought to cause behavioral interference have not yet been identified, much less quantified (as is pointed out, e.g., by Neumann (1987, p. 367) and Allport (1987, p. 395)). Thus, I suggest that we stick with the more conservative concept of resource-pooling, which refers to the sharing of resources between tasks, rather than undertake a description that quantifies the available resources. In my understanding, if two tasks share resources then we can say that they "pool" resources, where the terms "serial" and "parallel" are applied with respect to a single participant. That is, if a participant can perform two tasks without pooling, then the two tasks are performed in parallel. In a resource-pooling model of attention, any tasks performed without interference, or in true parallel, will be non-attentional.

I suggest that this understanding of attention and non-attention is superior to that of Hirst et al. because it leaves room to take into account the skills and capacities of the participant. That is, the measure of serial resource-pooling is one of interference between tasks, where at least one task is known to be attention-involving

for the participant. The measure could here descend into vicious regress if we are not careful: the determination of an attention-involving task cannot itself be dependent on these behavioral measures because we would then never find a first task for comparison. Thus, we have to use evidence from neuropsychology to find the first attention-involving task, after which we can set into motion the behavioral measure. Namely, we can then add a second task and if it interferes with the first task then the second task involves attention (for that participant). If adding the second task does not interfere with the first task then the second task is performed in parallel to the first task and does not require attention (for that participant).

To bring this discussion into the realm of the everyday, I will return to the sailing example. When Rosa uses attentive action to direct the movements of the boat she finds herself less able to keep up her conversation with Justin. When she slips into voluntary action, she finds that her conversation with Justin is uninterrupted. What allows her to have uninterrupted conversation with Justin is that she is no longer actively monitoring the movements of the boat, but achieving those movements through "auto-pilot." On the other hand, when she allows her conversation with Justin to slip into "auto-converse" she is more able to undertake attentive monitoring of the boat's movements.

I take it to be a (non-coincidental) advantage of the concept of resource-pooling that it can account for such everyday phenomena. This and other reasoning was behind my suggestion in the last chapter that we understand attention in terms of subject-level neural governance, which is marked by this kind of behavioral resource-limitation. Thus, although Hirst et al. call the behavior in their study "attentional," there is good reason to think the behavior non-attentional due to the parallel nature of the two activities (i.e. their lack of interference with one another).[14]

[14]A possibility that has not been mentioned which would prove problematic for my account of

Dual-task studies now largely assume this serial resource-pooling understanding of attention, rather than an understanding based on behavioral complexity. For example, the articles "Face-gender discrimination is possible in the near absence of attention" (Reddy et al., 2004) and "Rapid natural scene categorization in the near absence of attention" (Li et al., 2002) both claim to provide evidence of complex, goal-directed behavior without attention because of the absence of serial resource-pooling for those behaviors. That is, both of these studies use a dual-task paradigm similar to that of Hirst et al. that shows minimal performance loss from adding a second task to a first. They differ from Hirst et al. in assuming that the behavior is non-attentional due to its non-interference with the attention-involving task.

Thus, when we set aside an understanding of attention based on complexity we are faced with evidence of what appears to be selection-for-action without attention, revealing that attention is not necessary for selection-for-action. In the next section I will discuss what I consider to be the true role of attention in selection-for-action: optimization.

3.4 Optimality in Action and Perception

To explain what he takes to be the relationship between selection-for-action and attentional control, Neumann uses an analogy with trains:

> The problem of recruiting effectors in such a way that no mutually incompatible actions are attempted is similar to the problem of how to prevent train crashes on a railroad network where many trains use the same tracks. There are essentially two ways to solve this problem. One is to devise a

the Hirst et al. study is that the dual task is performed through task-switching ("alternation"). In other words, if the participants are very quickly switching from one task to another then the evidence against serial resource-pooling comes into question. This possibility is rejected by Hirst et al., but I think it is an option worth investigating further.

> schedule such that two trains will never be close to each other on the same track...The other method is to divide the network up into sections and to let a train entering a section block this section for other trains...The first allows a better usage of the limited resource...The second...needs less communication and coordination, and is much less error prone. (Neumann, 1987, p. 377)

Neumann uses the concept of centralized control for trains to illustrate the agentive governance of behavior: just as one might control a system of train tracks by matching the trains to the tracks through a schedule, governance allows the agent to match actions to effectors by "scheduling" upcoming actions through prioritization. Automatic control is evoked through a system of limited access, where only one train can occupy any given section of the track at a time, just as certain stimuli can sometimes unilaterally control effectors. That is, instead of prioritizing behavior through governance, a neural system can automatize behavior by allowing certain stimuli complete control over particular effectors in certain conditions, such as when the eye reflexively blinks upon seeing an encroaching object, like dirt, despite other competing visual tasks.

Using Neumann's train analogy, we can get a clue as to the agent's role in selection-for-action. As Neumann puts it, centralized control (like governance by the agent) "allows a better usage of the limited resource" (Neumann, 1987, p. 377). Centralized control is better for trains because it allows them to use every section of the track, instead of shutting off occupied sections. The analogy to effector control could be that more actions are accomplished through agentive governance because actions can be scheduled instead of waiting for the effectors to be out of use. That is, so long as a stimulus unilaterally controls an effector, that effector will be out of use for the

period of time that it responds to the stimulus, such that any other stimuli "fall on deaf ears" for that period. Centralized control allows new stimuli to interrupt the use of effectors, making effectors more responsive to new stimuli. Thus, the temporal grain of stimulus response for automatic control will be wider than that of governed control – as wide as the tasks are, on average, long.[15]

A benefit for attentional control that transcends not only the automatic control discussed in the analogy with trains but non-attentional governance is the ability to prioritize according to the subject's current goals. Whereas non-attentional processing may incorporate the long-term goals of the subject, attentional processing also involves the current or short-term goals by bringing to bear the planning areas. Thus, attentional action is not only optimal in the sense that it can allow for finer-grained behavioral response (like all governed behavior) but also in the sense that it allows for action that is better suited to the full range of the subject's goals. Donald Norman and Tim Shallice reason that attention is especially useful for actions that involve "planning," "trouble shooting," "novel sequences of actions," "overcoming a strong habitual response," or are "dangerous or technically difficult" (Norman and Shallice, 1986, p. 2). Thus, although not necessary for selection-for-action, attention likely optimizes it in some cases.

In order to show that not only selection-for-action, but selection-from-limitation, can be made optimal through attention, it is helpful to consider a paper by Jonathan Harel and Christof Koch. Remember that selection-from-limitation is the reduction of sensory processing, which I have argued can take place non-attentionally through filtering and local recurrency. To argue that attention sometimes optimizes selection-from-limitation I would need to show that attentional selection of processing is some-

[15]Though, of course, even agentive governance will not be perfectly efficient since, unlike trains, the scheduling of actions will have to occur in real time, as was pointed out to me by Dan Dahlstrom.

times more optimal than non-attentional selection. Harel and Koch demonstrate that processing reduction, when guided by the current goals of the participant, is superior to even unlimited computational power when it comes to stimulus recognition, hinting at an optimization role for attention (Harel and Koch, 2009, p. 1). This finding follows from the fact that stimulus recognition is always prone to false positives, and thus the more a participant scans a scene within a particular window of time to determine whether it contains a particular stimulus, the more false positives the participant will report. If we note that there are likely to be only a few of the desired stimuli per scene at any given time, then we see that even granting unlimited resources we should limit scene-scanning to keep the false positives from running down precision, or the number of correct identifications per false positives.

Imagine, for example, that a participant is given photos of groups of people and is asked to identify familiar faces. If the participant has enough time to scan each face, he or she will have more opportunities for false recognition. If there are only a few familiar faces in the group, even if the participant accurately identifies all of the truly familiar faces, the false recognitions will significantly reduce the value of precision. Thus, if we can reduce scene scanning without losing accuracy, this will optimize task performance. The problem with using non-attentional selection is that non-attentional selection diminishes accuracy, or the ability to detect proportionally high numbers of total targets, especially when targets move or change. By using attention to prioritize elements of the scene according to the participant's current goals, the participant can limit scene-scanning with only a small reduction in accuracy.

Attention's contribution to optimality in selection-from-limitation and selection-for-action can be understood more intuitively by thinking about how selection works. Non-attentional selection is at best based on the hard-wired features of the participant and his or her long-term goals. Attentional selection, on the other hand,

integrates both the long-term and current goals of the subject into a common framework (through the global state of recurrency) in order to determine which processing should be prioritized. Thus, attentional selection is more flexible than non-attentional selection. This flexibility makes attention more optimally suited to our many goals, although it is resource-expensive (because it uses serial instead of parallel resource-pooling) when compared to non-attentional selection.

3.5 Conclusion

Along the way I have indicated that a subject's involvement in action can take place at four different levels. First, the subject can consciously guide behavior, which is most common with unfamiliar or difficult activities. Second, the subject can attentively behave without conscious guidance, as when performing multiple tasks at once. Third, a subject can voluntarily behave without either conscious or attentive guidance, as when acting on auto-pilot. Fourth, a subject's behavior may be attuned to his or her goals (given certain perceptual inputs) such that the subject might later accept and take responsibility for the behavior even though it was outside of his or her control, as in sleepwalking. In this chapter I have shown that selection-for-action only requires attunement or voluntariness, but not necessarily the volitional involvement of the subject through attention or conscious guidance. Such volitional involvement can, on the other hand, optimize selection.

One application of this research is in separating attentional actions from complex behavior. That is, whereas attentional action involves the activity of the subject, complex behaviors can represent anything from practiced reflex to full conscious involvement by the subject. It is important that we make these divisions in order to understand the domain of attention and thus of one type of involvement by the subject. This could have implications for moral and legal responsibility, as when legal

responsibility depends on attentive involvement.[16]

In the next chapter I examine the relationship between attention and conscious perception. While in this chapter I found that attention is neither necessary to reduce sensory information nor for action, in the next I find that attention *is* necessary for conscious perception, which I argue through what I call the "Standard Theory" of conscious perception. Thus, the most direct contribution of attention to the mind is that of supplying conscious perception, which allows the subject to optimize sensory selection and action.

[16]Thanks to Doug Husak for pointing out the connection here.

Chapter 4

The Standard Theory of Conscious Perception

ABSTRACT: As I claimed in Chapter Two, attention prioritizes the neural processing associated with select percepts, corresponding with a prioritization of both the conscious experience of those percepts and the behavioral response to those percepts. In the last chapter I argued that the prioritization of select sensory processing and behavioral response need not be governed by attention, because this can be achieved pre-subjectively. In this chapter I argue that the subject-level prioritization brought about by attention is constitutive of and essential to perceptual consciousness. That is, where "conscious perception" refers to the experience of sensory content arranged within a space-time framework, attention is at the heart of any explanation of such content. Thus, attention is necessary for conscious perception as we know it. I present this argument through a theory, which I call the "Standard Theory."

4.1 Introduction

In the previous two chapters I have argued that attention, understood through neuropsychology, prioritizes neural processing in a way that can optimize both the selection of particular percepts and of behavioral responses to those percepts. Up to this point attention may have seemed a curious fascination for philosophers, disconnected from deeper philosophical questions on, for example, conscious perception. In this chapter I argue, to the contrary, that such prioritization is constitutive of and essential to the structure of conscious perception.[1] I present this argument through a theory,

[1]Sebastian Watzl has recently published a paper claiming that attention *is* this structuring of conscious experience: "*every* form of attention consists in a structuring of one's conscious point of

which I call the "Standard Theory." According to this theory, attention provides for the division and unity found across all perceptual features by prioritizing elements of those features against a global, subjective standard, using a spatio-temporal framework common to all types of features to do so.[2] The division and unity assumed in the Standard Theory is akin to that which is rejected as being a critical feature of conscious expereince by Semir Zeki (2003) and Jesse Prinz (2011).

The Standard Theory is original in its details but not in its vision. William James connects attention to the division and unity of conscious perception in his *Principles of Psychology.*

> Millions of items of the outward order are present to my senses which never properly enter into my experience. Why? Because they have no *interest* for me. *My experience is what I agree to attend to.* Only those items which I *notice* shape my mind – without selective interest, experience is an utter chaos. Interest alone gives accent and emphasis, light and shade, background and foreground – intelligible perspective, in a word. It varies in every creature, but without it the consciousness of every creature would be a gray chaotic indiscriminateness, impossible for us even to conceive. (James, 1981, p. 403)

James claims here that it is interest on the part of the subject that determines the content of perception, both in quantity and in quality. His claim rests on evidence of sensory selection ("Millions of items...never properly enter into my experience") to-

view" (Watzl, 2011, p. 2). I disagree that the concept of attention is necessarily tied to conscious experience and have thus argued that attention is the more general process of subject-level prioritization. Despite this disagreement, the description that Watzl gives of attention is very similar to what I argue is attention's key contribution to conscious experience.

[2]I call it the Standard Theory to highlight the role of this global standard in attention; the ambiguity between this and the other meaning of "standard" (i.e. generally accepted) is unfortunate but yet, I think, amusing.

gether with the idea that such selection must be governed by interests. As I discussed in the last chapter, quantitative selection does not need to occur at the subject-level, but can take place through filtering and local recurrency. Thus, the fact that we do not perceive all of the "items" surrounding us does not by itself show us that the interest of attention is the gateway to conscious perception. Nonetheless, James' intuition that the interest of attention provides for the qualitative content of conscious perception finds support in the Standard Theory.

Maurice Merleau-Ponty, on the other hand, entertains the notion that attention is the key to the common spatio-temporal framework of conscious perception. In the Introduction to *Phenomenology of Perception*, he makes the following observation.

> Attention first of all presupposes a transformation of the mental field, a new way for consciousness to be present to its objects. Take the act of attention whereby I locate a point on my body which is being touched....A *vaguely located spot*, this contradictory phenomenon reveals a pre-objective space where there is indeed extension, since several points on the body touched together are not confused by the subject, but as yet no univocal position, because no spatial framework persists from one perception to another. The first operation of attention is, then, to create for itself a *field*, either perceptual or mental, which can be 'surveyed' (Merleau-Ponty, 1962, pp. 33-34, translation amended)

Here, Merleau-Ponty claims that although experience outside of conscious perception is not itself unorganized (two points on the body are easily distinguished), attention allows for a new type of organization by creating a "field" of objective space and time for itself. Attention transforms conscious experience from pre-objective to objective

space-time by invoking a common spatio-temporal framework, according to Merleau-Ponty.

Merleau-Ponty argues for this view of attention (i.e. the view of attention as transformative) by looking at the phenomenon of directed search. Directed search implies both some determinacy and some indeterminacy on the part of the subject (Merleau-Ponty, 1962, p. 33). When I search for a pen on the surface of my desk, for example, I must have both enough determinacy to guide my search (e.g. that the pen is somewhere on my desk) and enough indeterminacy to drive my search (e.g. but I don't know where). Our ability to have some determinacy in our search entails that the resulting sensory knowledge will not be entirely new or "pure" (contrary to the "empiricist" view, as Merleau-Ponty construes it). On the other hand, if the result of our search is completely circumscribed by the determinacy guiding it (as the "rationalist" contends, according to Merleau-Ponty), then we do not have the requisite indeterminacy to account for the fact of searching. Thus, the search-mechanism will have to bring new sensory knowledge into the space of the subject. According to Merleau-Ponty, as the mechanism of this type of search, attention should be understood to play the role of connecting the realm of structured thought to sensory experience through the subject's transformative activity, making it possible for the unknown to enter into the realm of the knowable.

The Standard Theory gives a working account of these phenomenological intuitions. Namely, the Standard Theory shows how attention transforms the processing of sensations into the processing of percepts by prioritizing the former with respect to a subjective standard, where a subjective standard is a standard either rooted in or held by the subject. According to the Standard Theory, this transformation brings about both the content of perception, as flagged by James, and the determination of a perceptual field, as flagged by Merleau-Ponty. Thus, the Standard Theory is an ac-

count of how the subject brings about the structure of conscious perception through the activity of attention.

In what follows I first introduce the twin problems of division and unity. I then discuss Anne Treisman's answer to a subset of these problems. Finally, I show that the Standard Theory can fill in the gaps left by Treisman's work, providing a complete solution to these problems. In the next chapter I discuss conscious experience outside of attention before, in the final chapter, discussing the test case of gist perception.

4.2 Twin Problems

To start, it may be helpful to review my working definitions of attention and conscious perception. As I asserted in Chapter Two, "attention" should be understood as the process of prioritization by a "subject," where the subject is that which we infer from our own experience is having that experience and that which we infer from the behavior of certain organisms is normally governing their behavior. (Although, as I say in Chapter Three, their behavior can also be governed by a pre-subjective agent). I further argued that this process of prioritization is best specified through subject-level neural governance, or the prioritization of neural processing according to the goals, desires, and needs of the subject. Pre-subjective neural governance, or the prioritization of neural processing through local recurrent interactions, can be contrasted with neural anarchy, or the passing on of neural processing without feedback or recurrency. The attentional selection that occurs through subject-level neural governance, or prioritization that results from recurrency between the planning and sensory areas, on the other hand, can be contrasted with non-attentional selection – the selection of neural processing that occurs with either pre-subjective neural governance or neural anarchy, and thus without reference to the subject's current goals. As I argued in both Chapters Two and Three, attention also corresponds with the

behavioral markers of serial resource-pooling, or performance interaction between simultaneous attention-involving tasks. These behavioral markers play a key role in the empirical tests of the Standard Theory that I discuss in the final chapter, Chapter Six.

I take the primary understanding of attention, as subject-level prioritization, to be in keeping with phenomenal experience: attention is felt as an increase in phenomenal *priority* for a *subject*. I also take it to fit with behavioral evidence: attention improves the likelihood of stimulus response, or behavioral *priority*, by a *subject*. As I claimed in Chapter Two, the advantage of neuropsychological markers over purely phenomenal ones is that the former serve to indicate when and how changes in the stimulus are separate from changes in pre-attentional and attentional processing. Similarly, unlike purely behavioral markers, neuropsychological markers can serve to indicate when and how changes in the physical makeup of the subject and to other pre-attentional processing are separate from changes to his or her attentional processing. These points may be illustrated through the following example. If I notice a star in the sky getting brighter, then it might be because I am now attending more to that star, because of a change in my internal processing, or because that star really is now brighter (putting aside issues of perceptual judgment). If, on the other hand, my stargazing friends report experiencing more stars than I do, it might be because they are attending more to the night sky, because they have more relative expertise, or because they have better eyes (again, assuming uniform perceptual judgment). These differences in stargazing experience and behavior have neuropsychological markers that can more easily distinguish the disjuncts than evidence from experience or behavior alone can.

"Conscious perception," on the other hand, normally refers to the experience of sensory content within a space-time framework (or the process leading to that ex-

perience, where the process need not be conscious).[3] Although sensations may have spatio-temporal locations, sensation is thought to differ from conscious perception in lacking content. That is, unlike sensation, conscious perception presents the world as being a certain way to the subject. This "being a certain way" involves both general or shared attributes between the many items of content and particular instantiations of those attributes held by each item. The stargazer, for example, experiences the stars as each having relative brightness: they share the perceptual quality of brightness, but each of them instantiates a particular degree of brightness. As Tyler Burge puts it: "a perception–a representational perceptual state instance, or the content of a perceptual state instance–must always involve the context-dependent singular application of (general) perceptual attributives" (Burge, 2010, p. 381). Perhaps we could provisionally characterize conscious perception as the experience of sensory elements bound to a spatio-temporal structure, where "elements" signifies the particularities or instantiations of generalities, and where this spatio-temporal structure may include localized events, patterns, or simply spatio-temporal depth.[4]

A full characterization of conscious content has been undertaken by Giulio Tononi in his Information Integration Theory of Consciousness, where he describes such content in terms of information. As Tononi puts it,

> every time we experience a particular conscious state out of such a huge repertoire of possible conscious states, we gain access to a correspondingly large amount of information. This conclusion is in line with the classical definition of information as reduction of uncertainty among a number of alternatives (Shannon and Weaver, 1949)...the information generated by

[3] My use of the term "content" is purposefully ambiguous between external (as in direct realism), internal (as in sense-data theory), and relational content.

[4] Although I take it that these remarks hold for all varieties of conscious perception, I will focus in this chapter on visual perception.

> the occurrence of a particular conscious state lies in the large number of different conscious states that *could potentially* have been experienced but were not. (Tononi, 2005, p. 111)

An example that Tononi frequently usees to motivate this idea (the idea that conscious content is inherently informational) is the visual experience of a completely dark room: such an experience only has visual content if "content" measures what is present against what is possible. That is, a room devoid of light has visual content only when compared with other potential visual experiences one could have. Normal visual sensation, in contrast, relies on the presence of light.[5] The fact that we experience total darkness, then, shows us that conscious content is inherently informational. In Tononi's view, even the conscious perception of a homogenous plane of light (such as the absence of light, or total darkness) has a lot of information for us, since its content is generated in contrast to other potential experiences.[6]

Tononi further notes that the information of conscious content is integrated in a way that is not captured by "Shannon Information."[7] As phrased in the paragraph above, we might say of the content of perceptual experience that it is informational "for us." Tononi suggests that

> to measure information integration, it is essential to know whether a system of elements constitute a causally integrated system, or they can be broken down into a number of independent or quasi-independent subsets

[5]In making this claim I am assuming that the experience of "subjective gray" or *eigengrau* in total darkness is an illusory perception and not a sensation, which is debatable.

[6]A possible threat to the idea that a homogenous plane of light has *as much* information as a heterogenous, changing one is that such homogenous experiences are unstable and hallucination-inducing (Knau and Spillman, 1997; Wackermann et al., 2008).

[7]Although "Shannon Information" was put forward as data without meaning, and thus only part of a full definition of information (Shannon, 1948; Floridi, 2009), since it is now common to call such bare data "information," I follow the contemporary usage.

among which no information can be integrated. (Tononi, 2005, p. 113)[8]

Tononi claims that to get anything like conscious content we require not only information but integrated information, which we can measure by looking at causal integration. But integration with respect to what? It sometimes seems that Tononi wants to measure integration with respect to individual subjects. At other times Tononi describes integration more weakly, as when he says that integration is implicated by the fact that we cannot experience shape without color (Tononi, 2005, p. 112). The inseparability of shape and color will not tell us whether or not the information is integrated with respect to an individual subject, since these may be integrated with respect to a "quasi-independent subset" of the subject. For integration with respect to an individual subject we would need integration of shape and color with respect to some further thing, outside the processing of any set of elements. In contradiction to the darkness example, Tononi also says that integration is implicated when a "conscious state is experienced as an integrated whole" (Tononi, 2005, p. 112). However, that a particular experience is integrated does not tell us whether it is integrated with respect to an individual subject. For this, we would have to know whether the potential experiences of the subject together with the current experience are integrated. Without this level of integration the conscious experience of a completely dark room would not be able to yield information through contrast with other potential experiences. I thus suggest that the more plausible claim is the stronger claim that conscious content is integrated information to the extent that it is information-for-a-subject.[9] Importantly, I am not claiming here that conscious experience is integrated information (as Tononi claims), but only that conscious (per-

[8]Note that Tononi's use of the term "element" does not necessarily include participation in a complex as a defining feature.

[9]These thoughts on integration benefitted from a discussion with Justin Junge and Rosa Cao after attending a talk by Giulio Tononi, Christof Koch, and Ned Block at MIT in May 2011.

ceptual) content can be defined this way.

This understanding of conscious perception (as the experience of integrated information within a space-time framework) raises two connected problems. Specifically, assuming conscious perception is rooted in neural processing and that this requires a structural isomorphism between the two, neural processing must involve a part-whole structure that can support the experience of a unified complex of sensory elements. The problems arise when we note that the processing of sensations does not appear to contain this requisite structure. In order to do so, the processing of sensations would need to be divided into the processing of sensory elements and the processing of sensory elements would have to be held together as part of a unified complex. These are known as the problems of division and unity, where the problems apply both within and across sets of elements. Importantly, although these problems assume the possibility of a natural explanation of conscious perception, they need not confuse the content of conscious perception with its vehicles. That is, one can agree with Susan Hurley that "the properties of subpersonal processes, of vehicles of content, cannot simply be projected into personal-level mental content, or vice versa" (Hurley, 2002, p. 3) and still find particular vehicles wanting. As stated, it is the presumption of structural isomorphism, and not a simple projection of properties, that drives the claim that content-vehicles will have to instantiate a part-whole relationship that corresponds to the division and unity of content. I assume structural isomorphism because I take this to be the minimal constraint on the "is rooted in" relation.[10] Burge, mentioned above, claims that the development of perceptual content occurs through perceptual constancies (Burge, 2010, pp. 407-413), but I look further back to the source of these constancies and find a role for the subject through attention.

[10]To the extent that I manage to keep the personal and neural levels separate, I am indebted to the keen eye of Dan Dahlstrom.

As I will argue below, a solution to the twin problems of division and unity will necessarily involve a subjective standard, which can only be applied through attention. The solution to these problems lies, in other words, in the adoption of the Standard Theory. Before discussing the Standard Theory, though, I look at the work of Anne Treisman, who gives a partial solution to these problems.

4.3 Treisman and Her Critics

The recent surge of philosophical interest in attention was arguably inspired by John Campbell's *Reference and Consciousness*, a book in turn partly inspired by the work of Anne Treisman. Campbell's book can be seen as an account of how attention grounds thought in perception, answering Merleau-Ponty's charge that we find the middle way between rationalism and empiricism. Specifically, Campbell argues that attention is necessary for reference: "If you are asked the question, 'What colour is that block?', you must, to understand the question, consciously attend to the block" (Campbell, 2002, p. 12). In order to consciously attend to the block, Campbell argues, the block needs to be available as a bound object, which is where Treisman's Feature-Integration Theory (FIT) comes in. That is, Treisman's work is centered around the unity of objects, where objects involve multiple sensory features. The unity of objects is only one part of the problem of unity, but it is a worthy starting point. I thus present and defend Treisman's FIT in this section before building on its conceptual foundation with the Standard Theory in the next.

The Feature-Integration Theory, or FIT In 1980, Treisman published "A feature-integration theory of attention" with Garry Gelade, changing the landscape of writings on attention. Previous milestones in attention research looked at the role of attention in filtering sensory processing (primarily for auditory processing) (Broad-

bent, 1958; Deutsch and Deutsch, 1963). Treisman and Gelade's work shows that attention also transforms the character of sensory processing (primarily for visual processing), such that it is now largely taken for granted that there are two types of sensory processes – parallel and serial.[11] As Treisman and Gelade put it:

> In our model, which we call the feature-integration theory of attention, features are registered early, automatically, and in parallel across the visual field, while objects are identified separately and only at a later stage, which requires focused attention...focal attention provides the 'glue' which integrates the initially separable features into unitary objects. (Treisman and Gelade, 1980, p. 98)

In other words, according to Treisman and Gelade attention transforms the parallel processing of sensory features into the serial processing of objects, thus filtering out the features unrelated to those objects.

The evidence for FIT comes largely from visual search tasks: Treisman and Gelade found a distinction between two types of search tasks that later became the paradigm examples of parallel and serial processing. Namely, if a subject is given only one object in a visual display, whether they are asked to report on the presence of a single feature (e.g. something black) or a feature conjunction (e.g. a black "X"), the response time from stimulus onset to button press is the same – around 400 ms. However, if the number of objects in the display is increased, such that the subject has to *search* for the feature or feature conjunction, a difference emerges: single-feature

[11] "Parallel" and "serial," in this usage, refer to the processing of the percept rather than of the response. That is, where the "serial resource-pooling" of Chapter Three represents the sharing of processing resources between tasks as evinced by behavioral interaction between the tasks, Treisman and Gelade's "serial processing" represents the sharing of resources between the neural processing of percepts, as evinced by the relative ease the participant has in detecting those percepts and in the signal to noise ratio for neurons representing those percepts.

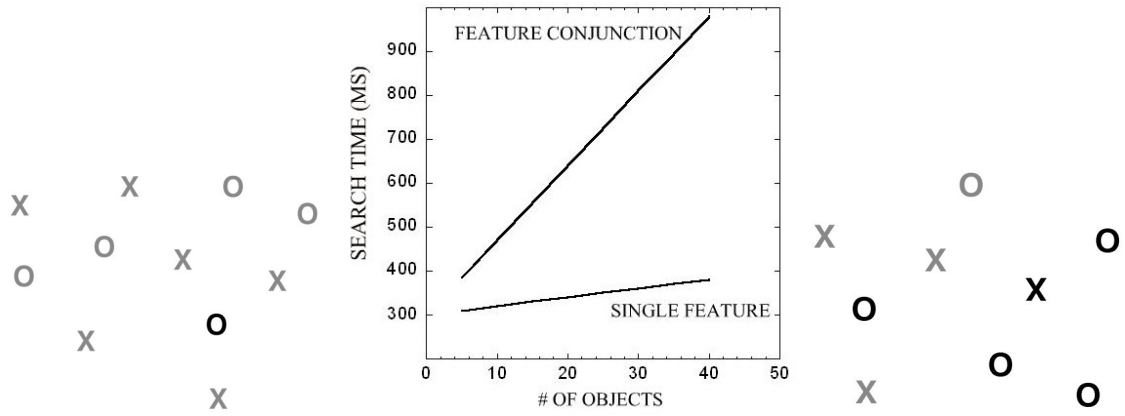

Figure 4.1: Single Feature and Feature-Conjunction Search

search takes around 400 ms *no matter how many objects are added to the display*, but response times in feature-conjunction search *increase linearly for each object added to the display*. For example, as shown in figure 4.1, a black "O" in a field of gray O's and X's takes no more time to detect than a black "O" on its own, but a black "X" in a field of gray O's, gray X's, and black O's takes more time to detect for each distractor that is added (i.e. each non black "X"). Treisman and Gelade reason from this that when search is based on single features the required processing takes place in parallel (i.e. there is no interference between the processing of the distractors and target). They further reason that the serial processing of feature-conjunctions derives from the need for attention to bind these features together. They postulate that in the case of feature conjunction the location of the features is essential to separating targets from distractors and that attention must therefore use feature locations to bind the features together.

The above is the standard description of FIT. However, Treisman and Gelade recognize that the story is a little more complicated than this since there is neural evidence of parallel feature binding and, moreover, we have experiences of bound objects beyond the focus of our attention. To account for the first fact they reason that parallel binding can occur for "integral" features (that act as a single feature) but

not for "separable" features (Treisman and Gelade, 1980, p. 101). For example, hue and saturation could be considered integral features, whereas color and orientation are separable. To account for the second, Treisman and Gelade suggest that inaccurate conjunctions of separable features (i.e. conjunctions of features that belong to separate objects) make up experience outside the focus of attention: "Our claim is that attention is necessary for the *correct* perception of conjunctions, although unattended features are also conjoined prior to conscious perception" (Treisman and Gelade, 1980, p. 98). Further cementing the connection between conjunction accuracy and attention, Treisman later describes attention as being either narrowly or widely distributed, where accuracy is the payoff (Treisman, 1988, p. 203). That is, attention can be spread widely for less conjunction accuracy or spread narrowly for more conjunction accuracy. Add to this the fact that, as Treisman and Gelade allow, familiar conjunctions can be formed with some reliability outside of attention, and it appears as though narrowly focused attention is neither necessary for conjunction nor for accuracy in conjunction, even for separable features. On a familiar walk one might, for instance, focally attend to one's cellphone screen while accurately conjoining features of the surrounding scene through a combination of widespread attention and the reliability of conjunction for familiar stimuli. That is, widespread attention and the reliability of conjunction for familiar stimuli may be enough for accurate conjunction in many instances. Thus, according to Treisman's later account, *some* amount of attention is necessary for new conjunctions of separable features and for accuracy in familiar conjunctions, but it is not necessary for conjunction in general.

The Guided Search Model Jeremy Wolfe, Kyle Cave, and Susan Franzel expose a glitch in this early account of FIT. First they show that even for standard conjunctions, *ruling out the presence* of a target is more linear (more dependent on the

number of distractors) than *detecting* a target, which is more efficient (Wolfe et al., 1989, p. 422). In other words, the reaction times of a participant asked whether a black "X" appears on the screen are less sensitive to distractors (i.e. less serial) than those of a participant asked whether the black "X" is absent from the screen. This means that, first, a process can be more or less serial and, second, feature binding is not the only factor separating serial and parallel processing. Wolfe, Cave, and Franzel also show that increasing feature contrast makes standard conjunction search nearly as efficient as single-feature search (Wolfe et al., 1989, p. 426). Larger letters, for example, might make the search for the black "X" nearly immune to distractors (i.e. more parallel). This shows that, first, a process can be more or less parallel and, second, feature binding is not a sufficient explanation of the separation between serial and parallel search. These two findings indicate that the distinction between serial and parallel search is not sharp, leading Wolfe, Cave, and Franzel to hypothesize that the relatively parallel processing of features is not wholly separate from but guides the relatively serial processing of objects (Wolfe et al., 1989, p. 428). In other words, the more visual information about the available features, the more parallel the process, which explains both the inefficiency of search for absent features in the first case and the efficiency of a search for salient features in the second. They call their hypothesis the "Guided Search Model."

Wolfe, Cave, and Franzel test their hypothesis with triple-conjunction search. FIT presumably predicts that searching for three features will lead to response times that are slower or equally slow to those for two features. The Guided Search Model, on the other hand, predicts that the extra information about the target yielded by the third feature will allow for faster search. They find that "overall, triple conjunction responses are faster than standard conjunction responses...and practically independent of the number of distractors," supporting their hypothesis (Wolfe et al., 1989,

p. 430).

Treisman's later work takes account of these and other experiments to suggest a mechanism of guidance between the parallel processing of feature analysis and the serial processing of accurate object binding. She theorizes that feature maps, or implicit placeholders for organized feature data and their saliencies, send location values to a master location map, on the basis of which attention serially binds features of the attended object to their spatial location, guided by the saliency data in the feature maps (see figure 4.2) (Treisman, 1988, p. 203). To illustrate the idea, Treisman uses the imagery of a needle sewing through layers of fabric, where attention is the needle that sews together features from different layers of fabric, using the location codes of those features to pick out a complete object in each stitch.

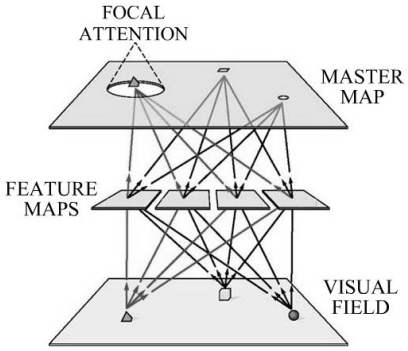

Figure 4.2: The Maps Model

The Synchrony Hypothesis Another criticism of Treisman's account emerges from a competing theory of feature binding that attributes binding to neural synchrony. Francis Crick and Christof Koch stipulate in 1990 that attention "helps set off the relevant neurons to fire in a coherent semi-oscillatory way," where the synchrony of these oscillations is thought to bind the features represented by the oscillating neurons (Crick and Koch, 1990, p. 263). In 1995, Wolf Singer and Charles Gray more radically hypothesize that binding *wholly* depends on neural synchrony, rejecting theories such as Treisman's that impute binding to location codes. Singer and Gray claim that location codes are overly resource-expensive in that they require neurons to keep track of not only the feature that they represent but where in space

(and, perhaps, time) that feature is. They reason that dropping this requirement allows for faster shifts and more flexible integration between different brain areas (Singer and Gray, 1995, p. 558). After reviewing an impressive range of neural evidence, Singer and Gray conclude in their favor:

> The experimental results reviewed in this chapter are compatible with the hypothesis that synchronization of neuronal activity on a millisecond time scale may be exploited to link featural information that is represented in different parts of the cortex. (Singer and Gray, 1995, p. 579)

This "Synchrony Hypothesis," as I call it, fits well with Robert Desimone and John Duncan's influential "Biased Competition Theory" of attention, where attentional selection emerges from the competition for synchrony between neurons instead of driving such synchrony through location codes (Desimone and Duncan, 1995, p. 214).[12] Thus, instead of holding that attentional selection causes binding through temporal synchrony, in Singer and Gray's account synchrony causes both binding and attentional selection.

Lynn Robertson defends FIT against the Synchrony Hypothesis in a 2003 *Nature* review (Robertson, 2003). Robertson offers evidence that the parietal cortex provides the spatial location codes necessary for accuracy in feature binding and that damage to these brain areas prevents accurate conjunction (Robertson, 2003, p. 96). Specifically, patients with unilateral visual neglect and Balint's syndrome (patients with damage to the parietal areas) exhibit increased rates of binding errors. As Robertson puts it:

[12]The hypothesis that attention emerges from competition and neural synchrony is further developed in the work of Alan Allport (Allport, 2011), Andreas Bartels and Semir Zeki (Bartels and Zeki, 1998), and in a paper by Andreas Engel, Pascal Fries, and Wolf Singer (Engel et al., 2001).

> Direct links between these neurons are not sufficient to explain surface feature binding, because damage to the parietal lobe disrupts such binding. An external source that synchronizes neural responses between cortical feature maps seems to be necessary, and this source seems to be spatial attention mediated by the parietal lobe. (Robertson, 2003, p. 97)

In other words, an account of binding that merely allows for competition and synchrony between neurons will not be able to explain the key role of spatial coding that is exhibited by pathological study. From this it seems likely that even if neural synchrony provides the mechanism of binding, it is not a sufficient source of binding, which is best characterized as the neural integration of features via a centralized spatial map, according to Robertson.

Binding in Unilateral Neglect Another criticism of Treisman's theory emerges from a couple of papers in a recently published collection on attention. Sebastian Watzl writes in a footnote: "I argue that [Treisman's] view is, among other things, undermined by recent findings of feature binding in the absence of attention (e.g. in Ro and Rafal (1996))" (Watzl, 2011). The paper by Tony Ro and Robert Rafal that Watzl refers to describes a patient with unilateral neglect who perceives the Judd Illusion on the neglected side. Unilateral neglect is a syndrome resulting from parietal damage that manifests itself as a failure to perceive and react to stimuli on the "neglected" side. The Judd Illusion, on the other hand, is the illusion that the midpoint on a line between two unidirectional arrowheads is further away from the open arrowhead, as shown in figure 4.3. As the Judd Illusion relies on a particular distribution of arrowheads to generate the illusory

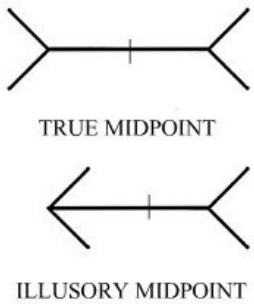

Figure 4.3: The Judd Illusion

midpoint, Ro and Rafal conclude that the patient perceives feature binding (Ro and Rafal, 1996, p. 977). Moreover, because the patient suffers from unilateral neglect, Ro and Rafal conclude that this binding occurs without attention. Christopher Mole concurs: "neglect patients are unable to *attend* to the neglected side of space. It is this that explains their neglect. The feature binding processes that occur in neglected patients must, therefore, be feature-binding processes that occur without constituting attention" (Mole, 2011b, p. 121).

I find two reasons to set Ro and Rafal's criticism aside. First, since Treisman is not committed to the view that attention supplies all binding but only that is supplies accuracy in the binding of separable features, a counterexample to her theory would have to show that the *accurate* binding of separable features can take place without attention. The purported successful binding of Ro and Rafal's single patient must be weighed against the body of evidence reviewed by Robertson that shows unilateral neglect patients suffering from higher than average binding errors for separable features. Second, an influential and oft-cited account of unilateral neglect by Marcel Kinsbourne in 1987 differs from Mole's in counting attention present in neglect:

> It is no explanation...of unilateral neglect of space that it has damaged a center for "attention"...Unilateral neglect results not from attentional deficit but from an attentional bias: imbalance in an opponent system that controls lateral orientation and action. (Kinsbourne, 1987, p. 69)

This stance by Kinsbourne and others – that neglect results from attentional bias rather than from attentional absence – would help to explain how it is that unilateral neglect patients can sometimes repair attentional pathologies on the neglected side with training (Robertson and Marshall, 1993). That is, biases, but not absences, can be re-set. Thus, feature-binding in unilateral neglect is neither compelling evidence

against Treisman's theory nor a clear instance of feature-binding in the absence of attention.

Illusory Conjunctions in Experience A final problem for FIT is its counter-intuitive prediction of regular illusory conjunctions to account for the richness of experience. Treisman and Gelade ask: "Can we reconcile our theory with the apparent speed and richness of information processing that we constantly experience? Perhaps this richness at the level of objects or scenes is largely an informed hallucination" (Treisman, 1988, p. 133). In their understanding, because attention is limited and because features outside the scope of attention are bound with some combination of habit and happenstance, much of our experience is constructed on the grounds of expectation.

Contrast this account with that of Alva Noë and J. Kevin O'Regan, who say that rich experience is nothing more than the amodal completion of what we are sensing according to our perceptual sensitivities. That is, when we experience a whole bottle while touching only part of it or see an entire room while looking at only part of it, experience is filled in according to what we are *able to* experience (the bottle or the room). This is not to say that our experience outside of attention is hallucinatory, they say, because we experience what we are sensing *as* a bottle or room without sensing the whole bottle or room. Noë and O'Regan ask, rhetorically:

> Does it seem to you as if you see all the environmental detail in uniformly sharp focus, all at once, now? Clearly not....While it does seem to us as if the detail is all there now, in the world, it is not the case that it seems to us as if all that detail is represented all at once in consciousness. (Noë and O'Regan, 2000)

Treisman and Gelade suggest, on the other hand, that the detail outside of attention *is* represented in consciousness, but that some of it is hallucinatory.

The evidence for illusory conjunction, Treisman and Gelade say, derives from behavioral studies that show subjects making more conjunction errors when attention is compromised (as in the evidence reviewed by Robertson, above). This is not evidence, however, that these subjects *experience* conjunction errors. Treisman admits in a later paper that experiences of illusory conjunction "are few and far between" in everyday life, but maintains that they occur and that they can be found through careful observation (Treisman, 1988, p. 215).

From FIT to the Standard Theory Despite the above criticisms, the later versions of FIT seem well-equipped to explain the binding of features into objects: sensory input is "automatically" divided into feature maps with priority values or saliencies and attention serially binds features from these maps into objects using spatial location codes stored in the parietal cortex, making this sort of binding more accurate than binding that does not make use of location codes. This theory gives us no reason, however, to think that attention is necessary for conscious perception; even according to Treisman herself, attention is not necessary for but aids perceptual processing, just as I argued that attention aids sensory filtering and the coordination of behavior in Chapter Three. I find that FIT misses the link to conscious perception because it stops short of fully answering the problems of division and unity. I will argue in the next section that attention is necessary for conscious perception because it is required to fully solve the problems of division and unity, of which FIT is only a start.

4.4 The Standard Theory

What we have so far is an account of how objects are bound from the features of separable feature sets, which answers only part of the problems of division and unity. Recall that the problems of division and unity concern the division and unity found throughout the perceptual field. If we want to fill in Treisman's account and fully answer the problems of division and unity, we need an account of both how feature-sets are created (of the division and unity in a single feature set) and how entire feature sets are held together (of the division and unity across feature-sets).[13] The Standard Theory provides these missing accounts. The basic claim of the Standard Theory is that attention provides for conscious perception by transforming unconscious sensations into conscious percepts. Attention performs this feat by prioritizing (and thus differentiating) the sensations according to a subjective standard, through which the sensations-turned-percepts are "unified" or integrated.

To start, let's return to the claim that conscious perception has division and unity. When I look across my living room, I see many shades of green – in plants, in books, in photographs, and in the house across the street. These parts of my experience share the quality of green-ness. Nonetheless, the instantiations of green are divided in my experience. The specific green of the plant is not confused with the specific green of the book. Thus, there is both a division and a unity of green-ness in my visual experience. Similarly, the sounds of two birds in a nearby tree both have pitch, but the sounds are not confused with one another. Thus, there is both a division and a unity of pitch in my aural experience.

How are these divisions and unities achieved? One might note the common role of space and time. The green of the plant and book are not blended into one an-

[13] I take it that Burge's perceptual constancies explain the unity within perceptual feature-sets without also explaining the existence of those constancies.

other, in part, because these objects are separated in space and time. The sounds of the birds have apparent source locations, which separate them in space and time. But can space and time account for the division and unity found in all perceptual experiences? It seems that I can perceive two pains without spatial and temporal location. That is, it seems as though the pains may be experienced as "here" and "now," as occurring in space-time, without being experienced as having a *location* in the space-time framework. My ability to separate the pains may be based on their different feels and intensities. Perhaps one pain is connected to emotional heartache and the other to the dehydration following alcohol consumption. One can imagine, I think, awaking to these pains without yet localizing them in space and time. If this is correct, then the space-time framework aids but is not necessary for division and unity. What can play this role? A clear common factor in all perceptual experience is the experiencing subject. That is, the parts of my perceptual experience are divided and unified for *me*. When we examine perceptual experience we discover a common pointing back, where the focus of that pointing is the experiencing subject, "me." This was brought out in in the discussion on Tononi, where "integrated information" was interpreted as "information-for-a-subject." Another way of putting this is that perceptual experience presents itself as something inseparable from the experiencing subject (without necessarily presenting other knowledge about the subject and what its status is with respect to other subjects). If we want to explain division and unity we should start with this most common factor of "for me"-ness, putting aside issues about the space-time framework and the division and unity at other layers, such as color and pitch.[14]

Whereas perceptual experience is marked by a pointing back to the subject, or a

[14]This move constitutes a departure from FIT, where attention is strongly linked to the common spatial framework and where binding is only described with respect to single objects.

"for me"-ness, the problems of division and unity concern processing outside of perceptual experience that cannot support this pointing back. Namely, the processing of sensations does not include interaction between its parts. The processing of color, luminance, and orientation, for example, can occur with little to no interaction. The processing of aural and visual sensations is further separated, despite these having interaction in experience. So how do we get from the parallel processing of these parts of experience to the interaction required for division and unity? Recurrency is a natural answer. Recurrency is a state of interaction between processing in different areas. It can take place through feedback from higher areas to lower areas, for example, where the processing thereby gains synchronicity. If we have recurrency between the neural areas that process visual and aural sensations, we have the potential for interaction between these areas, interaction of an analogous sort to that which we find between visual and aural percepts in experience. But recurrency at this level will not suffice for the division and unity of conscious experience, because the latter involves not just interaction between the parts but interaction with respect to the subject. So, recurrency with the areas associated with the subject, as an indication of subject-level activity, will likely be further required for conscious perception. In Chapter Two I claimed that such long-range recurrency is necessary for attention because attention is definitionally performed by subjects, a conceptual and semantic point. Here I claim that this long-range recurrency is necessary for conscious perception because conscious perception involves a pointing back to the subject, which will minimally require that the neural processing underlying conscious perception include interaction between the areas associated with the processing of the stimulus and areas associated with the processing of subject-level activity, such as the planning areas.

The integration of sensory processing into a state of recurrency with the planning areas is relevantly similar to the incorporation of nutrients into the body that occurs

in the process of digestion. When I digest peanut butter, the peanut butter is (eventually) transformed into parts of my body. How is this accomplished? The process of digestion shares some similarities with the process of turning light into plant cells for a plant: just as the light is taken in by the plant, the peanut butter is taken in by my body. But unlike plants, bodies have considerable (intra-species) variation in what they will incorporate. My friend Colin has an immune system that rejects peanut butter. The rejection by Colin's immune system reveals a self-organization that is not present in plants.[15] That is, unlike genetic rejection (e.g. from Celiac disease) and personal rejection (e.g. from disgust), rejection by the immune system reveals a self-governed determination of what will be incorporated by the body.

Just as the immune system (partially) determines what will be incorporated into the body, the subject (partially) determines what will be integrated from the mind into the subjective space as subject-level activity. This can be seen by looking at the variation between perceivers. For two perceivers watching someone bike past, one may perceive the color of the bike but not the clothing of the cyclist while the other perceives the color of the cyclist's clothing but not that of the bike. This shows that, just as the boundary between the body and the non-body is internally regulated, the boundary between the subjective and the pre-subjective is internally regulated. This boundary is the very complement to unity: the unity shared by perceptual elements implies a disunity between those elements and all other existing things. The consumption model that I have detailed so far is necessary because of this internal/external boundary that limits the subjective space of the mind together with the observation that the boundary is internally maintained.[16]

[15]That is, as far as I know it is not present in plants. I make this claim only to illustrate what is true about digestion, which does not hang on the truth or falsity of self-organization being absent in plants.

[16]I take it that non-consumption models of content, such as Dretske's magnetosomes (Dretske, 1993), do not have the integration required for their "information" to count as content.

If these observations are correct, the integration of sensation to conscious percept requires the volition of the subject. However, integration also has features that we do not find in incorporation or consumption. Unlike the peanut butter that becomes a part of my body by simply being allowed to stay within the physical confines of my body (or by being "incorporated"), sensation is *integrated* as *information* for the subject. That is, the consumption model gives us an account of how something might be incorporated by a self-organizing power, like the subject, but not an account of how the incorporated items become differentiated with respect to the subject, or become information-for-a-subject. Thus, anything in the subjective space must not only be incorporated but also differentiated with respect to the subject. This, I claim, is the role of attention. Importantly, "integration" and "differentiation" (like "division" and "unity") are twin concepts. That is, for parts to be integrated with respect to some standard is for the parts to be differentiated, and for the parts to be differentiated with respect to some standard is for them to be integrated. Thus, the role of attention in differentiating and the role of the subject in integrating are linked: the subject is the standard of integration with respect to which attention differentiates.

To see the link between informational differentiation and integration, consider that for something to count as information it must be differentiated with respect to some standard. The number "59" is not information, even though it can be differentiated from other numbers, unless it is set against some standard. The phrases "My father is aged 59 years" or "I saved 59 dollars on my wedding dress," do count as informational because they supply a context and an implied standard. For something to be information-for-a-subject the subject must either hold the standard or the standard must be rooted in the subject. That is, the differentiated parts will have to point back to the subject according to the subjective standard. Otherwise the information will not be integrated or unified with respect to the subject. Attention, I have argued

so far, is the process of prioritization by the subject or according to the subject. That is, attention is "subject-level" prioritization, where it will not matter for my purposes whether the subject is the actual mover of the activity of prioritization or whether the subject is like Confucius' pole star, which simply stands still in acceptance while the elements point back to it.[17] Either way, prioritization just is differentiation with respect to a standard. Thus, for the subject to differentiate those items it integrates, it will have to use subject-level prioritization or what I have called "attention."

The most basic form of integrated information, in my picture, can be found in a creature that perceives in terms of pleasure and pain (the most basic form of subject-level prioritization) without other sensory information. Let's imagine such a creature named "Michelle" who likes sitting on cool beaches (in Scotland, say) but detests sitting on hot beaches (in Florida, say). When presented with a particular beach in Scotland, Michelle might have a perceptual experience of pleasure at a particular intensity, shape, and duration, which she calls "West Sands." It may be that "West Sands" for Michelle applies to beaches other than what humans call "West Sands" so long as those beaches produce just the same experience of pleasure. When presented with a particular beach in Florida, Michelle might have a perceptual experience of pain at a particular intensity, shape, and duration, which she calls "Naples Bay." It may be that "Naples Bay" (for Michelle) applies to other beaches, or even to other things that Michelle detests, so long as those things produce just the same experience of pain. Michelle will not have a perceptual experience of anything that does not have a ranking of pleasure or pain for her. Michelle may have no feelings at all towards umbrellas, for example, in which case she would not perceive them. For Michelle to perceive something, that thing must be ranked according to her subjective

[17] "The Master said, 'He who exercises government by means of his virtue may be compared to the north polar star, which keeps its place and all the stars turn towards it.' " (Confucius, 2007, p. xi).

standard. Otherwise that thing would not be differentially integrated with respect to Michelle (qua subject). My claim is that all perceptual experience must be based on a subjective standard akin to Michelle's pleasure and pain, by virtue of which it has subject-level division and unity.

The story so far supplies only the most universal and basic form of integrated information. The conscious perception of most humans has layers of division and unity beyond the subjective standard. It is the division and unity of these further layers, I claim, that requires a common spatio-temporal framework.

Returning to an example used at the start of this section, there is a division and unity of green-ness in my visual field. How might this occur? My suggestion is that it occurs through the pre-subjective prioritization of sensory processing according to goals that reliably fit the subjective standard. That is, for stable subject-environment couplings it may be the case that the processing of sensations becomes tuned to goals that can reliably bring about value according to the subjective standard. Groups of neurons may become tuned, for example, to the inputs that normally enable the realization of subjective value and respond preferentially to them.[18] Thus, feature maps, or the mapping of inputs to pre-subjective standards that regularly contribute to the subjective standard, will take place without attention for long-term goals. This explains how it is that certain features of stimuli are pre-subjectively prioritized by the visual cortex – the inheritance of neural tuning created by millennia of other subjects matching goals like ours to a world like ours does much of the dividing work for us. Thus, one need not find attention necessary for each instance of division and unity, such as the division and unity of individual feature sets, but only for the unity of the perceptual field as a whole.[19] This roughly corresponds with the "maps model"

[18]I am indebted here to Tian Yu Cao whose lectures on the work of Ruth Millikan inspired the use of neural tuning in my theory.

[19]At the time of my prospectus, I had thought I would argue that attention is necessary for every

that Treisman offers above.

To achieve unity across multiple feature maps it will be necessary to have a framework on which to track the values of those maps. That is, the application of a unifying standard to multiple sets of feature values will require a structure through which to compare the values of each set without losing the internal structure of each set. This is the role of the common spatio-temporal framework found in nearly all perceptual experience. A common spatio-temporal framework is required so that the prioritizations found in the feature maps can be compared with respect to the subjective standard without losing the particular information contained in those feature maps. This common spatial framework need only be general enough to combine all the feature sets, and not so general as absolute or allocentric space-time. In such a system the different feature dimensions will share a framework or matrix that allows for prioritization across dimensions but that sacrifices neither the saliency values contained in each dimension nor the spatio-temporal location values of the information sources. This spatio-temporal framework helps us to solve the problem of subject-level unity in that it provides a mechanism of comparison for the application of a unifying standard. The accuracy yielded by egocentric spatio-temporal coding in the parietal cortex (as discussed by Robertson) also fits with this model: the common spatio-temporal framework will rely on egocentric information about the movement of the sensory organs to map the sensory data accurately through space and time. Even once we have a framework for comparison, though, we will still need attention to match the values from pre-subjective standards to the subjective standard to obtain subject-level unity.

Support for the Standard Theory comes from perceptual experience itself: per-

instance of division and unity. Christof Koch pointed out to me the error in this way of thinking during a meeting we had in April 2010, allowing for a complete shift in the way I went about these arguments.

ceptual experience always has the structure of theme to thematic field, or foreground to background, as was pointed out by Aron Gurwitsch in his *Field of Consciousness* (Gurwitsch, 1964).[20] That is, perceptual experience always contains a foreground of experience (the theme) that is delimited by contrast with its background (its thematic field), where the separation between foreground and background can be as rough-grained as that between two different sensory modalities. In "ganzfeld" experiences, for example, where a single sensory modality contains only homogenous information (as of a homogenous plane of light), that modality is yet experienced as being the foreground to the other sense modalities (Wackermann et al., 2008). Even in Balint's syndrome, where a patient with bilateral parietal lesions only recognizes one object at a time, that object is sometimes reported to be experienced against a background that serves as a contextual basis for that obect (Michel and Henaff, 2004, p.11). It is a strict result of prioritization against the subjective standard that a single theme will emerge for the perceiving subject at any given time (that one set of perceptual elements will emerge as having the highest priority for the subject). This theme may be an object, a type of motion, or some other set of perceptual elements. If there are cases of perceptual experience, perhaps involving Balint's patients, of only one theme at a time *without* a thematic field, I suspect that the theme will yet be experienced as informational for the subject. In that case, the residual structure of integrated information (the most basic "subject-world" structure) will still call for the prioritization work of attention in order for that theme to serve as perceptual content.

According to the Standard Theory, all perceptual experience minimally contains the structure yielded by prioritization with respect to the subjective standard, even

[20]I thus use Gurwitch's terms "theme" and "thematic field," as found in his books *Field of Consciousness* and *Marginal Consciousness* (Gurwitsch, 1964, 1985).

if it does not also contain other layers of prioritization that need to be fitted to a common spatio-temporal framework, as illustrated in the case of ganzfeld experience and in the case of Michelle, given above. Thus, although the provisional definition of conscious perception supplied at the start of this chapter involved the experience of sensory elements bound to spatio-temporal structure, the Standard Theory does not depend on this limited understanding of conscious perception to make its case. Any type of unified content, where unity is understood relative to the subject, will rely on attention, according to the Standard Theory.

4.5 Conclusion

Throughout this chapter I have illustrated how the content of conscious perception relies upon attention. I have claimed that although the process of division and unity for individual feature-sets can be automatized, such that neural structures can become tuned to particular types of features and collections of features according to the long-term goals of the subject, unity across features will only be achieved with the real-time application of the subjective standard through attention. Thus, if the Standard Theory is correct, attention is necessary for conscious perception as we know it, which minimally involves the division of theme from thematic field, but also normally involves deeper layers of division and unity.

The Standard Theory fits into a line of thinking that considers conscious perception the first step toward knowledge. As is previously mentioned in this chapter, James says that attention shapes our experience and provides it with "intelligible perspective." So, too, Merleau-Ponty finds that attention brings the unknown into the realm of the knowable, and Campbell shows that attention is necessary for reference, the building block of knowledge. The Standard Theory likewise holds that attention allows the subject to grasp the world by bringing the world into the space

of the subject.

This role for attention can be held in contrast to that which was discussed in the previous chapter. Namely, it was claimed there that attention is not necessary for mental selection, in particular the selection of sensory processing and action. The optimization of sensory and action selection discussed in that chapter can now be seen as an indirect result of attention's direct contribution to the mind, which is to allow for conscious perception and knowledge. That is, attention can now be seen to contribute, first, to the subject's grasp of the world, and second, only through that grasping, to the subject's ability to filter sensory information from the world and to act in the world.

A lingering worry that one might have with respect to the Standard Theory concerns conscious experience that does not have the structure of theme and thematic field (or, more basically, of subject and world): is attention necessary for this form of conscious experience? In the next chapter I will look at different forms of conscious experience that may exist outside the reach of attention.

Chapter 5

Conscious Experience Beyond Attention

ABSTRACT: Even if attention is necessary for conscious perception, a further question remains: are there conscious experiences beyond conscious perception and attention? In this chapter I examine some of the evidence for and against a positive answer, focusing on evidence of a second form of conscious experience that does not seem to have the structure of conscious perception. Before introducing this "immersion consciousness," I will review previous attempts to secure a "binary" theory of consciousness. Along the way I discuss the phenomenon of inattentional blindness as the most cited evidence for a "unitary" theory but argue (echoing several others) that it underdetermines the selection of a unitary or binary theory.

5.1 Introduction

In previous chapters I have argued that conscious perception depends upon attention to deliver its perceptual content, most notably resulting in the division of the perceptual field into a "theme," or highlighted foreground, and its "thematic field," or background. This thesis leaves open the possibility that there are conscious experiences that do not have the "integrated information" of conscious perception and so do not depend on attention. In order to address the question of this possibility it is useful to consider two types of theories concerning conscious experience that are currently used in the attention literature: *unitary* theories, according to which all conscious experience is assumed to be of a single type with respect to attention and *binary* theories, according to which conscious experience can be either of two

types, one of which does not rely on attention. Only if we side with the first can we echo William James' remark that "*my experience is what I agree to attend to*" (James, 1981, p. 402).[1] I will give examples of attention theorists on both sides of the unitary/binary divide before I go on to build the case for a binary theory in the sections to follow.

Some Unitary Theories In my understanding, a unitary theory of consciousness involves the assumption that all conscious experience is of the same type with respect to attention. Attention researchers reveal this assumption when they extend evidence that shows attention is (at best) sometimes necessary for conscious experience to conclude that attention is always necessary for conscious experience. Michael Posner was one of the first to draw this inference, saying of attentional mechanisms:

> I don't believe that any of these mechanisms are "consciousness" itself, just as DNA is not "life," but I do believe that an understanding of consciousness must rest on an appreciation of the brain networks that subserve attention, in much the same way as a scientific analysis of life without consideration of DNA would seem vacuous. (Posner, 1994, p. 7398)

Posner states explicitly that to understand any type of consciousness we must understand attention. In doing so, he presupposes that conscious experience is of a unitary type with respect to attention. Bernard Baars goes a step further than Posner in supplying a theoretical connection between attention and conscious experience via the Global Workspace Theory. The Global Workspace Theory is a theory of consciousness whereby conscious experience results from interconnectivity between brain areas, comprising a "global workspace" of high connectivity. Baars, like Posner, is careful to

[1] Valerie Hardcastle, for one, claims despite this passage that James "did not hold that what we pay attention to exhausts conscious experience" (Hardcastle, 1997, p. 102).

separate the concepts of attention and conscious experience while nonetheless giving attention a central role in conscious experience:

> Thus consciousness is the publicity organ of the brain, one that is used to access all of its functions. If this is the case, then attentional mechanisms exist to control access to this publicity organ, the bright spot on the stage of consciousness. (Baars, 1997, p. 370)

By assuming that consciousness is unitary in providing "publicity" (what I call "high connectivity"), Baars is able to move from the claim that attention is necessary for publicity to the claim that attention is necessary for conscious experience. Global Workspace theorist Stanislas Dehaene links attention and conscious experience more stringently by using a taxonomy of evidence to show that reports of conscious experience strongly correlate with the markers of serial (as opposed to parallel) processing, where publicity is marked by such processing:

> This framework postulates that, at any given time, many modular cerebral networks are active in parallel and process information in an unconscious manner. An information becomes conscious, however, if the neural population that represents it is mobilized by top-down attentional amplification into a brain-scale state of coherent activity that involves many neurons distributed throughout the brain. (Dehaene and Naccache, 2001, p. 1)

These Global Workspace theorists do not allow for conscious experience beyond the grip of attention in part because of their assumption that conscious experience is unitary. That is, they stipulate that conscious experience is *inherently* linked to publicity, where, for them, publicity relies on attention. In keeping with the work of these theorists, one might accurately say that within the cognitive science community

"it is commonly assumed that attending to a stimulus is a necessary precondition for becoming aware of it" (Pashler, 1999, p. 258).

Some Binary Theories A growing group of theorists have nonetheless diverged from the commonly-assumed unitary theory in favor of a binary theory of conscious experience. These theorists hold that there are at least two types of conscious experience and at least one of these types does not require attention. Before Posner correlated the importance of research on attention with that on DNA, for example, DNA theorist Francis Crick and his student Christof Koch published an article distinguishing two types of conscious experience, only one of which relies on attention. That is, Crick and Koch allow that "what reaches visual awareness is usually the result of [an] attentional step" (Crick and Koch, 1990, p. 269) but also suggest the existence of "another form of awareness that is very transient, being associated with iconic memory and having a very large capacity at any one time" (Crick and Koch, 1990, p. 272). Victor Lamme, similarly, thinks that attention is only necessary to integrate experience with other functions, but not for experience itself:

> Depending on the extent to which recurrent interactions between visual areas incorporate interactions with action or memory-related areas, awareness evolves from phenomenal to access awareness. Whether this occurs depends on attentional selection mechanisms, via influences on both the feedforward sweep and recurrent interactions. (Lamme, 2003, p. 16)

Finally, philosophers John Campbell and Ned Block provide accounts of conscious experience that support a binary theory. What is common to their accounts is captured by Christopher Mole: "According to this alternative hypothesis attention isn't necessary for *consciousness*, but it is necessary if one's experience is to provide one

with knowledge of the sort probed by the experimenter's questions" (Mole, 2008, pp. 95-96).

In the following sections I review the most promising evidence for a dependence relation between conscious experience and attention – evidence of inattentional blindness – and then present the worry that this evidence does not suffice to determine whether we should agree with the unitary theorist that attention is necessary for all conscious experience or with the binary theorist that attention is only necessary for one type of conscious experience. I critically examine evidence that has been given for binary theory (primarily by Block and Campbell) before setting it aside in favor of a new type of evidence for binary theory. Namely, I suggest that a binary theory is better able to explain the phenomenon of what I call "immersion consciousness" than a unitary theory is. Thus, I contend that although attention is necessary for conscious perception in supplying its perceptual content, it is not necessary for all types of conscious experience, such as experiences of "immersion" or attunement with one's activity.

5.2 Evidential Underdetermination

In building the case that attention is necessary for conscious experience, one might cite any of several key studies, such as Merickle and Joordens (1997), Rensink et al. (1997), or Simons and Chabris (1999). However, the most cited evidence for this claim is the work of Arien Mack and Irvin Rock, whose studies on what they call "inattentional blindness" (IB) from 1988 to 1995 have changed the way that most people think about perceptual experience. These inattentional blindness studies reveal that a substantial minority of participants fail to detect a stimulus when otherwise engaged in an attentionally-demanding task. Thus, I discuss these studies by Mack and Rock

before showing what is missing from these studies if one wishes to make the case that attention is necessary for all conscious experience. Namely, although Mack and Rock say that "there seems to be *no conscious perception without attention*" (Mack and Rock, 1998b, p. ix), they do not say (and do not have reason to say) that there is no conscious experience without attention. This would take a further claim about whether conscious experience is limited to experiences of the perceptual type.

Mack and Rock's exhaustive set of experiments on IB are published in a book named after the phenomenon (Mack and Rock, 1998b). The most common paradigm in the book makes use of a cross with axes of unequal length, where the participants are asked to determine whether the horizontal or vertical axis of the cross is longer. (I call this the "main task.") After a few rounds of this task a "critical stimulus" (an unexpected visual object) is presented within a quadrant of the cross. The participants are then "asked whether they had seen anything on the screen other than the cross figure" (Mack and Rock, 1998b, p. 6), to which the majority normally answer in the affirmative, and are then even able to identify the critical stimulus. (I call this the "stimulus-detection task.") A significant portion of participants, however, answer that they did not see anything but the cross, and it is these participants' responses that give rise to the claim that conscious perception requires attention.

In order to examine the logic behind this claim, it may be helpful to review some details of the experiments. First, the main task is relatively difficult: the duration of the cross stimulus is short, at around 200 milliseconds, and is followed by a "mask" (used to prevent further retinal processing). Nonetheless, performance on that task is at around 70 to 80 percent (Mack and Rock, 1998b, p. 55). Second, the critical stimulus is small, at around 0.15 degree of visual angle (Mack and Rock, 1998b, p. 53) (less than one-sixth the size of the index finger's fingernail at arm's length), but it is close enough to fixation that one easily notices it in passive viewing conditions.

In fact, upwards of 95 percent of participants correctly identify the critical stimulus in this condition (Mack and Rock, 1998b, p. 62). Third, the critical stimulus varies across experiments by type and within experiments by value. The experiments make use of different types of critical stimuli to study the effects of different features, such as texture segregation, pop-out, location, color, number, motion, flicker, shape, and apparent size. Fourth, the critical stimulus has three presentations: the first is unexpected and in competition with the main task, the second is expected (because the participant has now been asked about the stimulus) and also in competition with the main task, but the third is both expected and free of competition for attention, since the main task is dropped for that presentation (Mack and Rock, 1998b, p. 54). It is this third presentation that serves as the "passive viewing condition," and the control. Finally, to be sure that the critical stimulus is unexpected for the first presentation, the grain of the data is only as fine as one participant. That is, 95 percent accuracy on a task means that 95 out of 100 participants correctly performed the task, rather than that all participants performed the task at an average of 95 percent accuracy.

With those details in mind, the most impressive result of these experiments is the difference between the first and third presentations. That is, for every type of stimulus tested there is a "robust" proportion of the participants who do not report seeing the critical stimulus in the first "inattentional" presentation but who do report the critical stimulus (and can identify it) in the third "passive" presentation. Note, however, that the first presentation differs from the third both in expectation and in competition for attention. For this reason the second presentation (where the critical stimulus is expected but still in competition with the main task) serves as a second control to determine whether the change in expectation alone is enough to get rid of the IB effect. Mack and Rock found that the effect lingers in the second presentation

(though to a lesser degree), and so it seems likely that IB results from the competition for attention and not from lack of expectation (Mack and Rock, 1998b, pp. 8 & 61). This evidence leads to the provisional conclusion that competition for attention can diminish the numbers of participants who report perceptual consciousness of a stimulus.

To get from the provisional conclusion to the full-fledged thesis (that attention is necessary for conscious perception), a few versions of this experiment are particularly helpful. First, to be sure that the effect of diminished report is not one of inattentional amnesia, or the forgetting of a stimulus due to attentional competition (Wolfe, 1999), Mack and Rock design a task where there is minimal delay between the presentation of the critical stimulus and the question of whether the participant saw the critical stimulus. Normally the question about the cross stimulus ("Which arm of the cross is longer?") is followed by the question about the critical stimulus ("Did you see anything in the display besides the cross figure?"). To reduce delay, Mack and Rock alter the main task to one where the participant needs to correctly move his or her eyes during the task (a saccade task) instead of answering a question after the task (Mack and Rock, 1998b, p. 112). IB was observed to be present *to the same extent* in the saccade task as in the question task, making it unlikely that IB is an effect of memory.

Second, to be sure that the effect of diminished report is not one of inattentional agnosia, or the failure to fully process perceptual content[2] due to attentional competition (Simons, 2000), Mack and Rock compare the amount of IB for the participant's name to that for slight misspellings of the name and for familiar words. They find a gross discrepancy between the amount of IB for the participant's own name and

[2]In order to be consistent with the papers I am discussing, from here to the end of the next section I use the term "content" to mean the more general "components" or "contained stuff," instead of the more technical "integrated information" understanding that I put forward in the last chapter.

for slight misspellings and familiar words. This discrepancy shows that the critical stimulus is deeply processed despite being reported as phenomenally absent by the participant (Mack and Rock, 1998b, p. 179). That is, the words are processed deeply enough to differentiate meaningful from non-meaningful words (i.e. one's own name from other familiar words) and at a high enough level of precision to rule out slight misspellings of one's own name, showing that the IB effect is not likely due to agnosia.

Third, Mack and Rock find that IB is likely modulated by attention. This finding is based upon experiments that show that IB (surprisingly) grows dramatically stronger when the critical stimulus is at fixation and the main task is in the periphery. Mack and Rock reason that this effect occurs because something appearing at fixation would normally be more distracting than something appearing in the periphery, and so attention inhibits items appearing at fixation more strongly than those appearing in the periphery. This inhibition effect does not work in all cases, however, and when the critical stimulus reaches a certain subjective size (for example) it is no longer subject to IB. Mack and Rock reason that, as with the participant's name and "smiley" faces, subjectively large objects reliably capture attention because of their universal importance to the perceiver, where being located at fixation does not have universal importance to the perceiver (Mack and Rock, 1998b, p. 161). Mack and Rock conclude that IB is evidence that attention stops the critical stimulus from entering perceptual experience in the case that attention is otherwise engaged with a main task and the critical stimulus is less important than the main task to the participant.

In more condensed terms, Mack and Rock's argument can be recomposed as follows. Assume an unexpected stimulus-detection task (SDT) where in the "dual-task" condition there is an attentionally-demanding main task and in the "single-task" condition there is no "main" task. As attention is resource-limited, the expected task

limits the attention available to the SDT in the dual but not the single-task condition. Thus, the SDT has less available attention in the dual than in single-task condition. Some significant minority of participants fail the SDT in the dual-task condition but almost none fail the SDT in the single-task condition. The (positive) change in attention from the dual to the single-task condition thus appears to be sufficient for a (positive) change in performance on the SDT. Now let's assume that a positive change in performance on the SDT indicates a mitigation of the initial failure of either memory, processing depth, or conscious perception. Since the positive change in performance on the SDT occurs even in cases of low memory requirements and even when there is evidence of adequate processing of the unexpected stimulus in the dual-task condition, it is not likely indicative of the mitigation of a failure of memory or processing but of conscious perception. Presuming that there is no other similarity (that is, other than less overall attention) shared among the "some participants" who fail to report the conscious percept or among the cases in which those participants fail,[3] it appears likely that a positive change in attention is generally sufficient for the mitigation of an initial failure of conscious perception. This leads Mack and Rock to claim that attention is necessary for conscious perception.

Even if we accept this argument, we are still left without an answer as to whether we should adopt the unitary theory and extend the evidence from IB to all conscious experience. That is, since the argument concludes that attention is necessary for

[3]Less overall attention is thought to exclusively explain these failures. It is presumed that the participants who fail to report the conscious percept in the dual-task condition but successfully report it in the single-task condition have less overall attention at the time of the experiment than those who successfully report the percept in both conditions. The idea is that decreased overall attention is enough to induce perceptual failure for these participants when there is competition with the main task. If the group of participants who fail to report in the dual-task condition have some other, non-attentional deficit then this presumption could prove fatal to the argument. However, no such non-attentional deficit has been suggested, and it is reasonable to suppose that a deficit of attention is the primary difference between these groups. Similarly, if the dual-task condition can be differentiated from the single-task condition in some other way then this could discredit the inference. As with the last worry, such a difference has not yet been discovered.

conscious perception, *even if it succeeds* it does not tell us whether attention is necessary for all conscious experience. For this, one would need a further story of why all conscious experience is of the perceptual type. A unitary theorist might offer such a story through the taxonomy work of Dehaene et al. – all tested conscious experience has been perceptual or informational in character (Dehaene et al., 2006). However, as Mole suggests, if there is a form of conscious experience left out of these tests then it remains an open question whether attention is only necessary for conscious perception or for conscious experience, in general.[4] I first look at the evidence given by Campbell and Block for a second form of conscious experience before introducing a new type of evidence for a binary theory of consciousness. I conclude that this latter evidence is sufficient to side with the binary theorist in the interpretation of experiments like those of Mack and Rock.

5.3 Imagistic and Phenomenal Consciousness

Both Campbell and Block put forward accounts of conscious experience that need not involve attention. As a second form of conscious content in addition to "propositional content," Campbell proposes "imagistic content," which involves only spatial (and not conceptual) relations between stimuli. As a second form of consciousness in addition to "access consciousness," Block puts forward "phenomenal consciousness," which involves only the experiential aspect of consciousness without any of its attendant functionality. These theorists' accounts diverge at several points, but importantly Campbell (in contrast to Block) does not argue directly for the existence of a second form of consciousness, and Block (in contrast to Campbell) does not explicitly tie attention to the first form of consciousness. Nonetheless, both Campbell and

[4] I will here set aside the argument, which I discuss elsewhere (Suchy-Dicey, 2011), that the *conceivability* of this second form of consciousness is enough to stave off the unitary theorist and instead examine positive evidence for the binary picture.

Block are helpful guides in laying out an argument for a binary theory. I assume in this section that Block's division of types of consciousness can be described in terms of types of conscious content and that Campbell's division of types of conscious content can be described in terms of types of consciousness, but I am not sure whether either Block or Campbell would accept this assumption. Thus, this section is not addressing Campbell's and Block's views directly, but two forms of binary theory based on their views.

Campbell was not the first theorist to differentiate so-called "imagistic" from "propositional" conscious content (see, for example, Kosslyn (1980, p. 366)). He does, however, provide a likely account of why propositional, but not imagistic, content might depend upon attention.[5] Namely, where propositional content is built on reference, attention allows the subject to separate the target stimulus from other stimuli within experience, and so allows the subject to identify and act on the referent. Campbell's account, I will argue, leaves open the question as to whether or not we should take the stage prior to propositional content (that is, prior to reference) to be conscious or subconscious. Thus, although Campbell is no doubt right to consider attention necessary for reference (and thus for propositional content), it is far from obvious that what facilitates attention and reference must be pre-attentional conscious content, which he calls "imagistic content."

Campbell's several writings on attention and reference appear to have the same central goal of explaining the role of conscious experience in bringing about reference to physical objects. That is, instead of collapsing the link between concepts and their physical referents (as with Wittgenstein's behavioral concepts (Wittgenstein, 2009,

[5]By "attention" here I mean Campbell's "conscious attention," or the experiential highlighting of a target, rather than his notion of subpersonal attention. Only Campbell's concept of conscious attention can be seen as a subset of the broader notion of attention as prioritization by the subject, the operative notion in this dissertation.

p. 108e) and McDowell's conceptual referents (McDowell, 1996, p. 46)), Campbell wants to explain how conscious experience can transform non-propositional content (yielded by the physical referents) into propositional content (the referring concepts). Campbell's explanation is that conscious attention (or "experiential highlighting") allows imagistic content to indicate a target, which makes the relevant underlying information processing available to the subject, thereby transforming the attended imagistic content into propositional content (Campbell, 1997, p. 55) (Campbell, 2002, p. 5).

In order to argue that conscious experience plays a role in bringing about propositional content, Campbell puts forward cases where intuition tells us that other indicators of reference, such as behavioral responsiveness and conceptual knowledge, are insufficient for reference in the absence of conscious experience. The first case is that of super-blindsight. Where blindsight is a condition yielding partial visual functionality in the absence of visual experience (as introduced in Chapter Three), super-blindsight is an imagined form of blindsight where the participant retains total visual functionality despite lacking conscious visual experience (Block, 1995). Thus, the super-blindsighter can behave as though he or she knows which stimulus is being referred to without consciously perceiving the stimulus. Campbell claims that intuition reveals to us that the super-blindsighter lacks an understanding of the demonstrative in this case, or lacks reference. The second case is that of a color-blind color expert, as found in Locke (2004a,b, p. 24 & 88) and Jackson (1982, 1986), where scholarly knowledge of a color without corresponding experience of the color does not suffice for an understanding of the demonstrative for that color, or for the subject being capable of referring to that color.

Whatever one thinks of these oft-debated examples, one can likely see that there is a more conservative lesson to be drawn from them than a binary theory draws.

That is, one could agree that these examples show that conscious experience plays a key role in reference but disagree that this shows that there are two types of conscious experience, one of which comes before attention in the causal order and is separable from reference.[6] If, for example, attention is the cause of both conscious experience and demonstrative reference (i.e. if attention is the common cause), then this could satisfy the intuitions at play in the cases above. The superblindsighter might, for instance, lack conscious experience and reference to the object because he or she is not attending[7] to the object. Perhaps the superblindsighter attends to certain visual features that allow for accurate behavioral response without attending to the object itself (as in visual-form agnosia (Goodale and Milner, 2004)). Alternatively, the complex behavior of the superblindsighter might be achieved through pre-subjective mechanisms, as discussed in Chapter Three. The color-blind color expert, similarly, might not consciously experience and refer to color because he or she does not attend to color. The expert may attend only to the effects of color, for example, without attending to color itself. Thus, intuition may tell us that reference is absent in these cases because both conscious experience and attention are absent. In this more conservative story, there need be only one form of consciousness – consciousness that is brought about by attention and which provides for demonstrative reference.

One significant worry about this conservative story is known as the Problem of Richness (as discussed, e.g., in Noë and O'Regan (2000); Coates (2004); Schwitzgebel (2007)): how can a unitary theory account for the richness of experience outside the focus of attention?[8] In Campbell's terms, how can a unitary theory account for the

[6]Some attention theorists, such as Mole (2008, p. 96), nonetheless take Campbell's findings to establish his binary theory.

[7]This use of "attention" is as subject-level prioritization, which is not necessarily conscious. If I were to use Campbell's "conscious attention," the superblindsighter and the colorblind color expert would not be attending to the visual stimulus by definition.

[8]A second form of the Problem of Richness argued for by Coates and Mole (among others) mistakenly supposes that attention must be either volitionally-felt or volitionally-directed: Mole

conscious content surrounding the highlighted target? This worry is sound so long as we see attention as operating exclusively on the "highlighted" or most-prioritized target, but has less traction when we allow attention to operate both on the most-prioritized target and on the less-prioritized surround, which is the understanding of attention that I have been using throughout this dissertation.[9] It would then be available to the unitary theorist to claim, for instance, that there is no conscious experience outside of the *thematic field* of attention (where the thematic field includes the theme and the background surrounding the theme), instead of claiming that there is no conscious experience outside the *focus* of attention. Once we allow this more nuanced version of the unitary theory it is difficult to see whether the Problem of Richness has bite. That is, it is difficult to see whether there are conscious experiences beyond the thematic field of attention.[10] Without further reason to hold the Problem of Richness as evidence against the unitary theory, we are left with an impasse with respect to the comparative explanatory merits of the unitary and binary theory of conscious content.

In an attempt to overcome this impasse, Block reinvigorates the Problem of Richness by meshing richness intuitions with behavioral and neural evidence. Before addressing his "Mesh Argument," I want to note a couple of differences between his and

says that "one does not pay attention to all the things that one is conscious of" (Mole, 2008, p. 86) and Coates says "there is usually too much that is present in my experience in a single moment for me to attend to and conceptualize every aspect" (Coates, 2004, p. 14). As discussed in Chapter Two, attention need not correspond to either volitional direction or felt volition, so these examples are fundamentally flawed. In fact, as I argued in Chapter Two, phenomenological observations on their own do not always track attention, so intuitions from such self-reflection cannot tell us for sure whether attention is sometimes absent from conscious experience. More evidence is required to make such claims – namely, behavioral and neural evidence.

[9]Note that Eric Schwitzgebel sets aside the possibility of any sort of "diffuse attention," whether of the highlighted target or of the surrounding attentional field, and defines attention as either focally or volitionally directed, which colors his findings that "every participant reported experience of unattended objects or in unattended modalities" (Schwitzgebel, 2007, p. 24).

[10]In other words, it is difficult to see whether we consciously experience anything like Gurwitsch's "margins" or Husserl's "horizons," at least on some interpretations of them (interpretations that take these to exist beyond the limit of the thematic field, rather than comprising that limit).

Campbell's accounts. First, instead of distinguishing imagistic from propositional content, Block distinguishes phenomenal from access consciousness, where (unlike Campbell) he does not assume outright that these two types of content are naturally separable, but intends to establish their separability through his Mesh Argument. Block describes *phenomenal consciousness* as the experiential content of consciousness: "phenomenally conscious content is what differs between experiences as of red and green" (Block, 2005, p. 46). This quote refers to the Inverted Spectrum Argument, where it is argued that behavioral and functional accounts of consciousness fail to distinguish between a normal subject experiencing red and an inverted-spectrum subject experiencing green (Block, 1990). Thus, phenomenal consciousness is the experiential content that is left out of behavioral and functional accounts. Block describes *access consciousness*, on the other hand, as the functional content of consciousness. The functional content of consciousness is contrasted with functionality outside of consciousness (e.g. sleepwalking), which evinces "indirect control...but not direct control of the sort that happens when a representation is poised for free use as a premise in reasoning and can be freely reported. (It is this free use that characterizes access-consciousness.)" (Block, 1995). Block argues that conscious functionality occurs when we are able to access our mental states, in the sense that we are able to reflect and act on them. A second difference between the two accounts lies in the fact that, in contrast to Campbell, Block does not explicitly discuss attention with respect to the distinction between these two forms of consciousness, so application to the current discussion will involve some interpolation. Specfically, I will assume that what Block understands by phenomenal consciousness is separable from attention, while access consciousness is not separable from attention.

As to the Mesh Argument, Block claims that the natural separability of phenomenal and access consciousness makes the best sense of certain behavioral and

neural evidence. The relevant neural evidence is that visual processing takes place in lower-level brain areas than the processing associated with control of access (Block, 2008a, p. 498). This evidence allows for two possibilities: either visual processing in lower-level areas is subconscious and only becomes conscious with access, which is controlled by higher-level areas, or this lower-level processing is by itself conscious and just takes on a new type of content with access. Block contends that particular behavioral evidence sides with the latter. Namely, he contends that particular behavioral evidence indicates that we have a greater capacity for conscious experience than for access, suggesting that these are processed separately. In an early paper on the topic, for example, he cites the work of Diana Raffman, a music theorist who shows that whereas we distinguish 1400 frequencies in sound, we report on only 80 pitches (Block, 1998, p. 34). Block argues that whereas access consciousness has the capacity of only around 80 pitches, phenomenal consciousness has the capacity of 1400 or more. Thus, by meshing the phenomenal, behavioral, and neural evidence, Block contends, we arrive at what I call a "binary" theory of consciousness.

However, as with Campbell's evidence, there is a more conservative rendering of the evidence that is available to us. The evidence that participants can distinguish any two of 1400 frequencies, for example, is arguably based on low-level access rather than on no access: the participants report that they experience a difference between any two of 1400 frequencies, although they cannot identify the precise difference. That is, it may be that the participants can freely report on and use as a premise in reasoning *that* there is a difference between any two of the 1400 frequencies without being able to report on and use as a premise in reasoning *the precise difference*. In fact, Raffman's description of ineffable musical experience involves access: "some sensory perceptual states have contents–and I mean legitimate representational contents–that are consciously accessible but not reportable" (Raffman, 1993, p. 9). This "low-level

access" interpretation of Raffman's evidence, that at least some access is found in even the most ineffable cases of conscious experience, fits better with the first account of the neural data, in which account it is access that pulls the subconscious visual content into the space of conscious experience.

This more conservative interpretation applies just as well to Block's better-known Sperling example (as used in Block (2008b,a)). As Block discusses, the Sperling experiments show that when a participant is flashed a random set of up to twelve letters for a fraction of a second, the participant reports experiencing all twelve letters but is only able to access and report around four individual letters (Sperling, 1960, p. 7). Block concludes, as with the Raffman example, that the capacity of phenomenal consciousness overflows that of access consciousness, where phenomenal consciousness has a limit upwards of twelve letters and access consciousness is limited to around four (Block, 2008a, p. 487).

Whether the Sperling experiments suggest such overflow, however, depends on how we understand these limits. There is good reason to think that the limits should not be understood in simple numerical terms: if a participant familiar with reading the English language is briefly flashed "four-letter word," he or she will probably be able to report all fourteen letters, showing that the limit of access consciousness can be extended beyond four letters. Similarly, most adult humans, unlike infants, can correctly choose the fuller of two buckets of graham crackers after watching how many crackers have been put into each bucket, even if more than four crackers are placed in either of the two buckets (Feigenson et al., 2002, p. 152); the adult can reduce multiple crackers to a single number concept, just as an English speaker can group "f," "o," "u," and "r" into a single word. Thus, to know that the capacity of phenomenal consciousness overflows that of access consciousness in the Sperling experiments we have to ask ourselves, "What four things are being accessed?" and "What twelve

things are being experienced?" If the answers to these two questions involve different "things," we may not have a case for overflow. If, for example, the subject has access to only four fine-grained letters but experiences twelve rough-grained letters, no overflow is evident; it may be that the subject both experiences and accesses twelve rough-grained letters and then experiences and accesses four fine-grained letters. Along these lines, Sid Kouider demonstrates that Sperling participants say they experience twelve normal letters even when some letters are reversed or replaced with non-letters (Kouider et al., 2008, p. 510), which suggests that the participants have only rough-grained information about all of the letters until accessing them as particular letters. The behavioral data, then, does not make a clear case for overflow.[11]

Without clear evidence of overflow, Block's Mesh Argument cannot establish a case for binary theory. That is, the neural evidence only meshes to a binary theory if we have phenomenal or behavioral evidence of overflow, but without such evidence we have no reason to believe that phenomenal consciousness takes place without access, even if visual data are processed separately from the control of access in the brain. If we add to this my suggestion that phenomenal consciousness is separable from attention but access consciousness is inseparable from attention, the lack of evidence for overflow with respect to access consciousness is also a lack of evidence for conscious experience outside the range of attention. Thus, binary and unitary theories appear equally able to account for the evidence put forward by Block, and his Mesh Argument leaves us with the same unenviable impasse regarding the comparative merits of these theories.

To review, after arguing that attention is necessary for reference Campbell asks the question, "What content drives our capacity to refer to a stimulus?" and finds an answer in imagistic conscious content. I suggest, to the contrary, that an answer

[11] A similar point is made about Block's Landman example by Byrne et al. (2008, p. 501).

to that question could involve conscious content, subconscious content, or a mixture of the two, leaving the relationship between attention and conscious experience up for grabs. Block asks the question, "What explains the richness of experience?" and finds an answer in phenomenal consciousness. I suggest, to the contrary, that the multi-grained character of phenomenal consciousness may be matched by the multiple levels of access consciousness, leaving no room for overflow. In this case, the richness of experience would not point to a relevant difference between phenomenal and access consciousness and would thus not establish the existence of conscious experience beyond attention. In short, the question of whether there are conscious experiences outside the reach of attention has not been settled by these accounts and the debates surrounding them.

One underlying problem is that these debates do not start from a full understanding of attention: Campbell looks at only one feature of attention (the fact that it highlights the stimulus) and Block ignores the topic altogether. This is not a problem with Campbell's and Block's work *per se*, but only with using that work to try to make a case for binary theory. Without a clear understanding of attention the task of finding conscious experience outside of attention cannot get off the ground. Thus, in the next section I will start by looking at the hallmarks of attention that I have argued for in the dissertation so far and then look to see if there is evidence of conscious experience beyond those hallmarks. I think I have found a candidate form of this sort of conscious experience in the phenomenon of "immersion consciousness," which I discuss in the next section, though more research on this phenomenon is needed to make a firm case for binary theory.

5.4 Immersion Consciousness

I claimed in Chapter Four that attention is necessary for the perceptual content of conscious perception, or for integrated information, where integrated information can be understood as information-for-a-subject. As I argued there, integrated information requires the differentiation of elements within the perceptual field as well as their integration with respect to a subjective standard. The usual form of integrated information is the division of the perceptual field into a theme and its thematic field, where the theme is the most prioritized activity for the subject and the thematic field serves as a less-prioritized background to that theme. Thus, to find evidence of conscious experience beyond the reach of attention we should look for conscious experience that does not contain this sort of integrated information.

What of the neural and behavioral markers that I discussed in Chapters Two and Three? In Chapter Two I argued that attention is subject-level prioritization, where subject-level activity is best picked out through neuropsychological markers. Specifically, I argued that subject-level activity is best picked out through markers of global recurrency. Although such markers enabled me there to separate attention as a subject-level activity from other pre-subjective activities, they may not be helpful for the question I pose in this chapter. Namely, if conscious experience is sometimes the result of *non-attentional* subject-level activity, then these neuropsychological markers will not help us because they confound attentional and non-attentional subject-level activities. I try to allay this confusion in Chapter Two by claiming that attention should be understood as *prioritization* that takes place through global recurrency (where I call such prioritization "subject-level neural governance"), but this point does not help us here without further neural markers of prioritization.

In Chapter Three I argued that attention corresponds with the neuropsychological

markers of serial resource-pooling, or performance interaction between attentionally-demanding tasks. However, because the markers of global recurrency are required to specify a first attention-involving task, these behavioral markers are subject to the same problem of potentially confounding attentional and non-attentional subject-level activities. That is, parallel-resource pooling may indeed indicate an absence of attention, but it may do so only by indicating an absence of (the more general) subject-level involvement. In that case, *serial* resource-pooling would not necessarily indicate the presence of attention but of any subject-level involvement. For these reasons, I do not use the neuropsychological markers discussed in Chapters Two and Three in this chapter.

A candidate for non-attentional conscious experience is what I call "immersion consciousness." Where the integrated information brought about by attention involves a separation between the experiencing subject and the experienced object (the object is experienced as being informational *for* the subject), immersion consciousness does not involve this separation. Immersion consciousness is a form of experience where the subject is fully immersed in its object rather than consuming information from it. (Note that if immersion consciousness truly operates beyond the reach of attention then it is a candidate for non-attentional subject-level activity, as discussed in the paragraph above.)

Something akin to the phenomenon of immersion consciousness can be found among the phenomena analyzed by Heidegger in *Being and Time*. Heidegger is intent on exposing the way that human beings find themselves immersed in the world.[12] These sorts of immersions are experienced and yet are, for Heidegger, more fundamental than the sort of subject-object relationship found in perception (Heidegger,

[12]He calls it "being-in-the-world," but I find this phrasing confusing because it makes it sound as though persons are physically encapsulated by the world, which Heidegger himself discusses as a problematic reading of the phrase (Heidegger, 1978, pp. 50-51).

1978, p. 49). In other words, Heidegger finds the separation of the perceiver from the perceived that shows up in observation of the object (as seen in the structure of integrated information) as secondary to various unions that take place in human activities.

To illustrate this "more fundamental" relationship Heidegger discusses the case of using a hammer, where

> the less we just stare at the thing called hammer, the more actively we use it, the more original our relation to it becomes...the act of hammering itself discovers the specific "handiness" of the hammer...No matter how keenly we just *look* at the "outward appearance" of things constituted in one way or another we cannot discover handiness. When we just look at things "theoretically," we lack an understanding of handiness. (Heidegger, 1978, p. 65)

Heidegger describes this way of relating to the world, this paradigmatic activity of being immersed, as a pre-theoretical involvement with the hammer's "handiness." He further illustrates this sort of phenomenon through examples where the activity and corresponding "handiness" go awry. That is, when a tool like a hammer breaks, is noticeably absent, or gets in the way, its "handiness" is forfeited (as is at once and inseparably the subject's use of it). As a result, the tool becomes an object, taking on a categorical nature, where the subject is forced into the subject-object relation with it again (Heidegger, 1978, pp. 68-69). A subject may be hammering, for example, when the head of the hammer suddenly flies off, leaving the subject faced with the hammer as a mere object instead of something bound to the subject in its use.

Such examples nicely illustrate a way of interacting with the world that has neither the structure of theme and thematic field nor the more basic perceptual structure of

subject and object. Heidegger himself generally eschews any talk of consciousness, and so does not directly address the question of the relationship between attention and consciousness. Nonetheless, I hope to show how it is helpful to extend his examples to find evidence of immersion consciousness. Such evidence can at best reveal an apparent lack of subject-object structure, and so more work will be needed beyond this chapter to show that the underlying processes behind these experiences are outside the reach of attention.

Although I think I have found a candidate form of experience, I will have to rely on the familiarity of the experience that I intend to invoke since this form of experience has not (as far as I know) received much attention in scientific research. Thus, I will proceed by offering two anecdotes in the hopes that I stir in the reader the memory of something familiar. Only the second of these anecdotes properly corresponds with immersion consciousness, but the first will be used to distinguish immersion consciousness from a closely related type of experience.

Example 1: David grew up on the edge of Yosemite National Park and is thus fond of all mountain-related activities, including mountain biking. He reports that he once had an experience when he was at his peak (not long before he flipped the bike over a flat rock, landing on and splitting his helmet) where the foreground of his visual experience receded into the background. That is, he experienced the bike trail, tree line, mountain horizon, and sky at once, without experiencing the prominence of one over the others.

Since one effect of attention on conscious experience is to divide it into foreground and background, the lack of such division (as in the example above) may seem to indicate a lack of attention. However, another reading of such experiences is that they involve low-effort attention to the foreground activity. That is, when the foreground

recedes into the background there are three potential causes. First, it might be caused by the subject *failing to divide* the perceptual field through attention. Second, it might be caused by the *purposeful attempt* to widen the theme of attention to the entire perceptual field (as might be achieved through certain kinds of meditation practice). Third, it might be caused by the minimal effort required for the foreground activity, resulting in a similarity of effort for foreground and background activity that is *felt* as a lack of division between foreground and background. It seems as though the third of these is the most likely cause of the conscious absence of division between foreground and background in this example. That is, we might say of the above example that the activity of biking involves the prioritization of biking-related visual information, but that the use of such information is so practiced that the information is experienced as though at the same level of priority as the background. Thus, this example is not clear evidence for immersion consciousness because the structuring work of attention may still be present, even though it is difficult to detect through the subjective phenomena.

Example 2: Carolyn has drawn friends from life since the age of five or six (although those early renderings looked more like each other than like her friends). She is now very practiced at life drawing. She sometimes finds that while she is drawing she enters something of a conscious trance where she is not interested in anything but the present moment. Once, for example, she found after finishing a drawing that her jeans were entirely covered in charcoal dust because of the ferocity of her application. When in this trance-like state she emerges from it suddenly, as though waking.[13]

[13]The previous example derives from the report of my fiancé, David Jennings, and this example derives from my own experience. I decided to use the third-person to make the examples more digestible as case-studies for the reader.

Carolyn is not alone in this type of experience. Jackson Pollock, for instance, has said about the process of painting: "when I am *in* my painting I am not aware of what I am doing" (Janson and Janson, 2003, p. 974). In fact, a catalogue of similar examples has been collected by Mihaly Csikszentmihalyi as part of his work on what he calls "flow consciousness" (as in Csikszentmihalyi and Csikszentmihalyi (1992)). Hunter and Csikszentmihalyi write: "Flow is a state of experience where a person, totally absorbed, feels tremendous amounts of exhiliration, control, and enjoyment. In flow states people...experience a merging of action and awareness" (Hunter and Csikszentmihalyi, 2000, p. 12). While Csikszentmihalyi's flow consciousness has a positive emotional element, which is his target in studying such experiences, I am focused here on the respect in which subjects no longer see their environment objectively, or apart from themselves, but nonetheless experience their own activity within it. That is, Csikszentmihalyi's examples illustrate not a mere collapse of the division between foreground and background but a collapse of the division between subject and world.

Immersion consciousness, as illustrated in the second example, has three characteristics that are relevant to this chapter (besides the apparent collapse of the division between subject and world). First, the relationship between the subject and the world is one of attunement rather than attention. I first discuss the concept of attunement in Chapter Three as an alternative type of action, borrowing the concept from Wu. Attunement, as I say there, is goal-based coupling of percept to motor response directed by pre-subjective mechanisms that are neither necessarily conscious, attentive, nor within the subject's control. In other words, the attuned subject is for all appearances skillfully responding to particular visual stimuli, but not in a way that *requires* the subject's volition.

Second, the immersed subject is nonetheless conscious, and not just in the sense

of being awake. The subject may not be able to recount just what occurred while in a state of immersion consciousness, but remembers *being* conscious. The subject may experience his or her immersion activity as though in a dream. Many people consciously experience dreams that they are unable to remember upon waking, where their dream experiences slip through their grasp as soon as they wake up. So long as we accept the account of the dreamer, it seems reasonable to accept that of the subject who undergoes immersion consciousness. That is, immersion consciousness appears to be a type of conscious experience that, in lacking the integrated information of conscious perception, cannot be grasped by the subject in a way that allows for memory and contentful report.

Third, the jump from immersion consciousness to conscious perception is a discrete one. In the example of Carolyn, the life-drawing artist, I wrote: "when in this trance-like state she emerges from it suddenly, as though waking." The observation that such experiences, though made possible by the gradual attainment of a skill, involve discrete shifts from perceptual experience indicates that immersion consciousness is different in kind from perceptual consciousness. Further evidence that immersion consciousness is different in kind is the fact that when subjects emerge from it they often find it difficult to retain control of their activity.

The above points indicate that immersion consciousness does not involve the most basic level of integrated information (the division of subject from world) but does involve conscious experience. Thus, we have a candidate form of conscious experience outside of attention. The problem with such experience is that it is difficult to document, making it a challenge for us to move away from mere anecdotes. Without consistent documentation one may reasonably object that we do not know for sure whether these experiences operate outside of attention, especially noting the likely

effect of both biased report[14] and biased report gathering. It could be, for example, that the apparent loss of the subject-world divide in immersion consciousness derives from an objectification of the entire perceptual field at once, where the division between the subject and the perceptual field is nevertheless maintained through contrast with the subject's history. Further, one could wield the conservative story of the last section against immersion consciousness and ask whether the immersed subjects are using minimal attention, rather than no attention. The subjects are able, after all, to report something about their experiences, which implicates residual attention. Evidence against both of these worries is that the subjects in question experience a merging with their environment, rather than a further distancing and objectification. It is not as though, for instance, the immersed subject is merely *observing* his or her environment and the activity occurring within it. Moreover, as noted above, the shift from immersion consciousness to perceptual consciousness is felt as discrete, where the attentional shifts within perceptual consciousness are normally felt as relatively gradual. However, these concerns can only be fully resolved with a more thorough study of the phenomena.

To answer this problem, more scientific study centered around immersion consciousness is needed so that a more definitive answer on the binary/unitary debate can be found. Such study could look for correlations between phenomenological report and behavioral or neural patterns, perhaps by examining experts in the middle of their craft. Some useful contrasts might be made with non-conscious immersion in skilled activity and with emergence behavior following a state of immersion. It may also be possible to study immersion consciousness through the observation of meditation practice, which has long been connected to the loss of subjective presence. However, the current model of non-focused meditation ("open monitoring meditation")

[14]Such biased report may even be due, in part, to Heidegger's influence on popular psychology.

reviewed by Antoine Lutz, et al. is more akin to the objectification of all experience at once than to the immersion of the subject within his or her environment (Lutz et al., 2008).

5.5 Conclusion

In sum, the evidence I find most convincing for binary consciousness, or for conscious experience outside of attention, is that of immersion consciousness, where immersion consciousness lacks the subject-world structure of perceptual experience. In place of this structure, immersion consciousness involves a union of subject and world in the activity connecting the two. However, the evidence that I offer for immersion consciousness is incomplete, leaving the question of whether there is conscious experience beyond attention in need of further support.

If immersion consciousness if found to exist as described here, it would make conscious experience an area of research that must transcend the study of attention. That is, if immersion consciousness exists and operates outside of attention, then we have no reason to think that what makes consciousness unique is due to attention. What attention provides, I argued in the last chapter, is the structure of perception. What attention might not provide is the quality of "presence" that so many have sought to understand through consciousness research. Thus, a fuller investigation into immersion consciousness may help us to understand what allows for conscious experience outside of attention, which would take us one step closer to an understanding of conscious experience itself.

Chapter 6

Gist Perception: A Case Study

ABSTRACT: Because the role of attention in the phenomenon called "gist perception" is controversial, it provides a valuable test case for the Standard Theory and its contention that attention is necessary for the perceptual content of conscious perception. By discussing gist perception I hope to explicate some of the claims made throughout the dissertation in more practical terms. I find that gist perception, although it may sometimes undergo attentional modulation, has been shown to operate without attentional modulation in some cases.[1] However, when gist perception is non-attentional it is not yet clear whether the percept is consciously experienced. Thus, I conclude that more research is needed to settle whether the case of gist perception is problematic for the Standard Theory.

6.1 Introduction

Gist perception, or the perception of stimuli with very short exposure times resulting in minimal perceptual content, has served as a central point of criticism against theories claiming that attention is necessary for conscious perception (of which the Standard Theory is an instance). Philosopher Paul Coates writes in 2004: "But there is further evidence that shows that data from the unattended parts of the visual field enter into conscious experience....Subjects are able very rapidly to ascertain the gist or 'general sense of a complete scene'" (Coates, 2004, p. 16). Scientists Christof Koch

[1] "Attentional modulation," in my understanding, is a change in the percept's underlying neural processing that is due to the neural processing of attention, and which results in a change to the percept.

and Naotsugu Tsuchiya write in 2006: "We are always aware of some aspects of the world that surrounds us, such as its gist" (Koch and Tsuchiya, 2007, p. 18). These theorists focus on gist perception largely because of its speed: it is faster, they suppose, than could allow for attentional involvement. However, some research suggests that the speed of gist perception does not prevent attentional modulation. Thus, I look at evidence both for and against the claim that gist perception is modulated by attention. I conclude that gist perception sometimes operates outside of attention. With this conclusion in hand, I then look at whether the gist percept is consciously experienced in these cases. I find the empirical evidence lacking to resolve the latter question.

6.2 Defining Gist

In everyday speech, "gist" refers to certain global properties of a stimulus. The gist of a story may be its plot, for example, because the plot can give a sense of the story without all of its details. In cognitive science, "gist perception" has come to refer to the perception of certain properties of a stimulus before others (normally the perception of some of its global properties before finer-grained details). When I flip through the channels of a television, for example, I might have the gist percept of a young girl drinking juice, without perceiving other details of the scene, such as the color of the juice and whether the girl is sitting or standing. One must not confuse the gist percept with knowledge of the gist percept – the idea is that the gist percept is itself limited in this way. It may seem impossible that one could perceive juice without perceiving its color or perceive a (fully-embodied) girl without perceiving the position of her legs, but gist percepts are strange. In my experience, the gist percept is as of a cloud of visual information in that it is neither fully determinate nor filled in (where what is left out is experienced as blackness or "subjective gray"). A similar

experience to gist perception occurs soon after one stands up too quickly and the visual field goes dark. In such a condition I sometimes find that, as the details of the visual field return, I think that I can see everything. It is only when the details continue to fill back in that I realize my earlier experience must have been relatively bare, by comparison.

Gist perception is differentiated from other perceptual phenomena by both the (abbreviated) length of stimulus exposure time and the resulting limitations on perceptual content. First, gist perception occurs with very short exposure times, which some researchers consider to be what defines this sort of phenomenon (see, e.g., Fei-Fei et al. (2007)). Mary Potter first observes the remarkable capacity of gist perception in 1976 by showing participants images for less than 125 ms and then testing their knowledge of those images (Potter, 1976). Aude Oliva demonstrates gist perception for as little as 30 ms of exposure (Oliva, 2005), and more recent work claims to show gist perception for 16 ms exposure times (Pavlopoulou and Yu, 2010).

Second, gist perception comprises limitations on perceptual content. Oliva demonstrates one such limitation through the different spatial frequencies of visual stimuli, where low spatial frequencies capture changes in light intensity over long distances and high spatial frequencies capture changes in light intensity over short distances (see figure 6.1 below).[2] As Oliva shows, normal conscious perception picks up both the low and high frequencies of an image, but gist perception includes only one or the other (Oliva and Torralba, 2006). Specifically, if an image contains two scenes where one scene is filtered for low frequencies and one is filtered for high frequencies

[2] "Spatial frequency" is short for "frequency of change over a particular distance or region of space." For example, if we wanted to know the spatial frequency of an image, we would look at how often the image changes from white to black over each region of space, where the size of a region is held constant. We can filter an image for a particular frequency by allowing only one change of white to black for a particular region size. A high frequency version of an image, for example, might represent only the changes that occur over every two pixels, whereas a low frequency version of an image might represent only the changes that occur over every 20 pixels.

(a "hybrid" image – see figure 6.1 below), normal conscious perception involves both scenes but gist perception involves only one of the scenes (normally the low-frequency scene) (Schyns and Oliva, 1994). Thus, the short exposure times of gist perception limit the percept to one or the other spatial frequency.

Figure 6.1: A Low-Frequency Image, A High-Frequency Image, & A Hybrid Image

6.3 Attentional Gist

Koch and Tsuchiya claim that the exposure times of gist stimuli are too short to involve significant modulation by attention: "In a mere 30 ms presentation time, the gist of a scene can be apprehended. This is insufficient time for top-down attention to play much of a role" (Koch and Tsuchiya, 2007, p. 18). "Top-down attention" is described by the authors as being "task-dependent" (Koch and Tsuchiya, 2007, p. 16). However, at least one study shows that task can play a significant role in gist perception. Namely, a paper by Philippe Schyns and Aude Oliva reveals the task-dependency of the selection of a spatial frequency in gist perception when they show that the selection of one frequency or the other is related to the participant's current task. When, for example, the participant is given a hybrid image consisting of a high-frequency face and a low-frequency face and is asked whether the face is expressive or non-expressive (remember that participants see only one of the two faces because of the short duration of the stimulus), the participant will tend to rate the expressiveness of the high-frequency face. On the other hand, if the participant is

asked to *identify* the facial expression, the participant tends to identify the expression of the low-frequency face (Schyns and Oliva, 1999, pp. 253-254). Since the very same participants see the high-frequency version of a particular image in one task and the low-frequency version of the same image in a different task, there is good reason to suppose that this task-dependency reflects the involvement of subject-level prioritization, or attention.[3]

6.4 Non-Attentional Gist

On the other hand, dual-task studies from Koch's lab have demonstrated that gist perception can be processed in parallel with other, attentionally-demanding tasks. One such study by Leila Reddy, Patrick Wilken, and Christof Koch compares the results from two tasks run separately and together (see figure 6.2) (Reddy et al., 2004). The first task (the "T/L task") involves five letters at fixation turned to different angles that are masked after 200 ms, where the participant has to report whether the letters are five L's, five T's, or a mix of T's and L's. The second task (the "gist task") involves identifying the gender of a face at any of eight peripheral locations, where the face is shown for 26 ms and then masked. When run separately the performance on the two tasks is close to 80 percent; when run together the

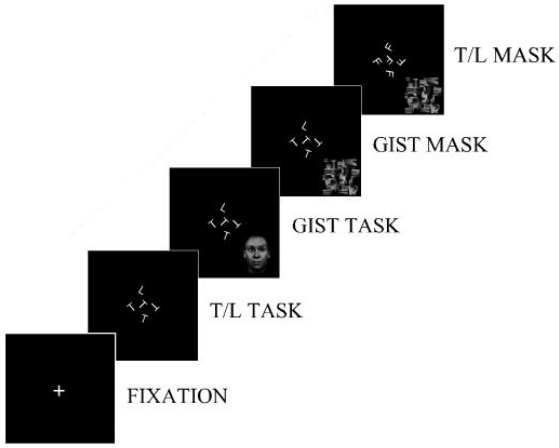

Figure 6.2: The Dual-Task Experiment

[3]There is also some work under review claiming that an increased perceptual load can eliminate gist perception and that this is evidence that gist perception is modulated by attention. Unfortunately, the perceptual load demands in this work are so high that gist may be filtered out pre-attentionally, and so I find this route less promising.

performance does not drop significantly for either task (Reddy et al., 2004, pp. 109-110). Thus, the two tasks seem to be processed in parallel, or without interference between the tasks. When the researchers ran a similar experiment with the T/L task at fixation but a different task in the periphery (where this "new task" had the same single-task performance as the gist task), performance dropped dramatically for both tasks (the T/L task and the new task) in the dual-task setting, showing that the T/L task requires attention (Reddy et al., 2004, p. 113). According to the behavioral markers of attention that I discussed in Chapter Three, this indicates that gist perception need not always involve attention. That is, since there is no apparent interference between the T/L task and the gist task, we have reason to believe that these tasks are not pooling resources. The fact that the T/L task interferes with the new task indicates that the T/L task is attentionally-demanding, and thus any other attentionally-demanding tasks should pool resources with this task. Ergo, the gist task in this study does not seem to be an attentionally-demanding task.[4]

6.5 Non-Attentional Gist and Conscious Perception

The above studies seem to show that gist perception is sometimes attentional and sometimes non-attentional. The key question in assessing whether gist perception is a counter-example to the Standard Theory will thus be whether the gist percept is conscious when it is non-attentionally processed. Unfortunately an answer to this question cannot be found just by looking just at the performance in the experiments above or at reports of participants who say that they experienced the gist percept (if these exist). That is, a participant's ability to answer questions about the global properties of a stimulus is insufficient to establish conscious perception of that stimulus, since the ability to answer questions about a stimulus has been shown to exist

[4]Other similar experiments can be found in Li et al. (2002), Li et al. (2003) and Reddy et al. (2006).

for non-conscious stimuli (Debner and Jacoby, 1994; Kentridge et al., 1999, 2004). Furthermore, a participant's ability to answer questions about the global properties of a stimulus together with conscious experience of the stimulus is insufficient to establish conscious perception of the stimulus, since what we need to show is that there is information available to the subject in the conscious gist percept that either drives the ability to answer questions about the stimulus or is informationally equivalent to whatever drives the ability to answer questions about the stimulus. If, for example, the participant experiences only a homogenous gray shape where the gist percept should be, this experience is likely insufficient to account for the participant's behavioral performance in categorizing the percept as, say, an animal.

A couple of studies examining gist perception in the single-task setting find it to be inconsistent with normal perceptual experience. Daniel Levin et al. find that jumbling the features of animals has very little impact on the gist perception of animals (Levin et al., 2001), even though normal conscious perception of animals is severely affected by jumbling their features. Evans and Treisman find that participants are very poor at localizing the gist percept of an animal in a natural scene when given three location options (Evans and Treisman, 2005), even though conscious percepts are normally localized. On the other hand, that gist perception is unlike normal conscious perception does not show that it is not conscious. In fact, it is likely that gist percepts are at least sometimes unlike normal percepts, which have color, shape, and location, since gist percepts can include properties like average size (Alvarez and Oliva, 2009), where such properties do not have color, shape, and location. Thus, further study on the topic is needed.

I suggest that one route to an answer might be found in looking at the statistical information available in participants' descriptions of their own experiences of gist percepts while in the dual-task setting. If participants' descriptions of their own

experiences carry information equivalent to that of their categorization performance then we might infer that the participants' performance is driven by conscious perception. Three approaches attempting to capture and systematize subjective reports are relevant. While none of these approaches has yet been applied to gist perception in the dual-task setting (from what I understand, such experiments are in progress), I review them to show what is possible in the future.

The first approach for measuring (reports of) conscious perception uses confidence ratings, where participants assign values to their confidence in particular perceptual judgments. Participants in the dual-task experiment might, for example, say that they are "very confident" or "not confident at all" in their judgment that a particular stimulus contains an animal. The advantage of confidence ratings is that they reveal whether participants see themselves as guessing, where the experience that one is "just guessing" despite performing at above average accuracy is a hallmark feature of blindsight and other subconsciously-driven behavior.

The confidence ratings approach is criticized in a 2007 study which claims that its participants found such ratings artificial. This study offers the more "intuitive" approach of gambling on one's perceptual judgments (Persaud et al., 2007). Participants in the dual-task experiment might, for example, be given the option to bet either 50 cents or a dollar on their judgment that a particular stimulus contains an animal. If the participant's willingness to bet money does not correspond with the participant's performance then this might be a sign that the participant's performance is subconsciously-driven. This "post-decision wagering" is claimed to be successful in experimentally differentiating the subconsciously-driven high performance of blindsight patients, Iowa Gambling participants, and artificial grammar learners from the high performance of normal perceptual experience.[5]

[5]On the other hand, Zoltan Dienes and Anil Seth found confidence ratings and post-decision

Against both confidence ratings and post-decision wagering, second-order valuations of first-order perceptual judgments are unlikely to fully capture perceptual experience, since perceptual experience is separable from judgments resulting from that experience. This worry can be leveled against both approaches discussed so far to the extent that both require the participant to rate the participant's confidence in a *decision*, whether explicitly or through betting, rather than rating the quality of the percept itself. A recent review compares the above two approaches to a third approach, discovered by asking participants to create their own method of assessing the quality of their conscious percepts, and finds that this third approach is more direct than the others (Sandberg et al., 2010). This third approach, the "perceptual awareness scale," requires participants to describe and then assign a numerical value to the quality of their perceptual experience. Kristian Sandberg, et al. show that this perceptual awareness scale captures borderline conscious experience that is not captured by the other approaches. It is the fact that gist perception is a borderline case that makes it attractive as a test case for the Standard Theory. Thus, the perceptual awareness scale seems more promising than either confidence ratings or post-decision wagering as a possible test of whether gist perception is conscious in the dual-task setting.

However, even if the perceptual awareness scale does not discover the presence of a conscious gist percept in the dual-task setting, this is not sufficient evidence that these gist percepts are not consciously experienced. Any description by the participant of his or her experience will rely on the participant's ability to describe his or her experience, or to metacognate that experience. As Anil Seth argues, any test of consciousness that relies on metacognition is open to the criticism that it leaves out some conscious experience (Seth, 2008). The proposed experiment, which seeks

---

wagering to yield equivalent values for the artificial grammar case (Dienes and Seth, 2010).

to discover whether gist perception is conscious by means of a statistical analysis of participant reports, would thus only be a first step in discovering whether the gist percept in the dual-task setting is consciously experienced.

6.6 Conclusion

In conclusion, the test case of gist perception is inconclusive on the question of whether it constitutes conscious perception without attention. On the one hand, gist perception is sometimes modulated by attention. On the other, we do not have evidence that the gist percept is conscious in the cases where it appears to be processed outside of attention. Thus, more evidence is needed to determine the fate of the Standard Theory in the face of such cases.

Bibliography

Allport, D. A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action. In Heuer, H. and Sanders, A. F., editors, *Perspectives on Perception and Action*. Lawrence Erlbaum.

Allport, D. A. (1992). Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In Meyer, D. and Kornblum, S., editors, *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*. MIT Press.

Allport, D. A. (2011). Attention and integration. In Mole, C., Smithies, D., and Wu, W., editors, *Attention: Philosophical and Psychological Essays*. Oxford University Press.

Alvarez, G. and Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention. *Proceedings of the National Academy of Sciences*, 106(18):7345.

Augustine, S. (2003). *City of God*. Penguin Classics. (Trans. H. Bettenson).

Baars, B. J. (1997). Some essential differences between consciousness and attention, perception, and working memory. *Consciousness and Cognition*, 6(2-3):363–371.

Bartels, A. and Zeki, S. (1998). The theory of multistage integration in the visual brain. *Proceedings of the Royal Society B*, 265:2327–2332.

Bergson, H. (1920). *Mind-Energy: Lectures and Essays*. McMillan and Company. (Trans. H. W. Carr).

Bergson, H. (2007). *Matter and Memory*. Cosimo Classics. (Trans. N. M. Paul and W. S. Palmer).

Bichot, N. and Desimone, R. (2006). Finding a face in the crowd: parallel and serial neural mechanisms of visual selection. *Progress in Brain Research*, 155:147–156.

Block, N. (1990). Inverted Earth. *Philosophical Perspectives*, 4:53–79.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18:227–47.

Block, N. (1998). How to find the neural correlate of consciousness. *Royal Institute of Philosophy Supplement*, 43:23–34.

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9(2):46–52.

Block, N. (2008a). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5-6):481–499.

Block, N. (2008b). Consciousness and cognitive access. In *Proceedings of the Aristotelian Society*, volume 108, pages 289–317.

Block, N. (2010). Attention and mental paint. *Philosophical Issues*, 20.

Bottcher-Gandor, C. and Ullsperger, P. (1992). Mismatch negativity in event-related potentials to auditory stimuli as a function of varying interstimulus interval. *Psychophysiology*, 29(5):546–550.

Brigard, F. D. and Prinz, J. (2010). Attention and consciousness. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1:51–59.

Brisson, B., Leblanc, E., and Jolicoeur, P. (2009). Contingent capture of visual-spatial attention depends on capacity-limited central mechanisms: Evidence from human electrophysiology and the psychological refractory period. *Biological Psychology*, 80(2):218–225.

Broadbent, D. E. (1958). *Perception and Communication*. Pergamon Press.

Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.

Busse, L., Katzner, S., Tillmann, C., and Treue, S. (2008). Effects of attention on perceptual direction tuning curves in the human visual system. *Journal of Vision*, 8(9).

Byrne, A., Hilbert, D., and Siegel, S. (2008). Do we see more than we can access? *Behavioral and Brain Sciences*, 30(5-6):501–502.

Campbell, J. (1997). Sense, reference and selective attention. In *Proceedings of the Aristotelian Society*, volume 71, pages 55–98. Blackwell Publishing; The Aristotelian Society.

Campbell, J. (2002). *Reference and Consciousness*. Oxford University Press.

Campbell, J. (2004). Reference as attention. *Philosophical Studies*, 120(1):265–276.

Carrasco, M., Ling, S., and Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, 7(3):308–313.

Chua, K. P. and Chun, M. M. (2003). Implicit scene learning is viewpoint dependent. *Perception & Psychophysics*, 65(1):72.

Coates, P. (2004). Wilfrid Sellars, perceptual consciousness and theories of attention. *Essays in Philosophy*, 5(1):8.

Confucius (2007). *The Analects*. NuVision Publications.

Crick, F. and Koch, C. (1990). Towards a neurobiological theory of consciousness. In *Seminars in the Neurosciences*, volume 2, pages 263–275.

Csikszentmihalyi, M. and Csikszentmihalyi, I. (1992). *Optimal Experience: Psychological Studies of Flow in Consciousness*. Cambridge University Press.

de Gelder, B., Tamietto, M., van Boxtel, G., Goebel, R., Sahraie, A., van den Stock, J., Stienen, B., Weiskrantz, L., and Pegna, A. (2008). Intact navigation skills after bilateral loss of striate cortex. *Current Biology*, 18(24):1128–1129.

de Vries, W. A. (1988). *Hegel's Theory of Mental Activity*. Cornell University Press.

Debner, J. A. and Jacoby, L. L. (1994). Unconscious perception: Attention, awareness, and control. *Journal of Experimental Psychology*, 20:304–304.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., and Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences*, 10(5):204–211.

Dehaene, S. and Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79:1–37.

Desimone, R. and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18(1):193–222.

Deutsch, J. A. and Deutsch, D. (1963). Attention: Some theoretical considerations. *Psychological Review*, 70(1):80–90.

Dienes, Z. and Seth, A. (2010). Gambling on the unconscious: A comparison of wagering and confidence ratings as measures of awareness in an artificial grammar task. *Consciousness and Cognition*, 19(2):674–681.

Dretske, F. (1993). Misrepresentation. In Goldman, A., editor, *Readings in Philosophy and Cognitive Science*. The MIT Press.

Engel, A., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nature Reviews Neuroscience*, 2:704–716.

Evans, K. and Treisman, A. (2005). Perception of objects in natural scenes: Is it really attention free?. *Journal of Experimental Psychology: Human Perception and Performance*, 31(6):1476.

Fei-Fei, L., Iyer, A., Koch, C., and Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1).

Feigenson, L., Carey, S., and Hauser, M. (2002). The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological Science*, 13(2):150–156.

Floridi, L. (2009). Philosophical conceptions of information. In Sommaruga, G., editor, *Formal Theories of Information*. Springer-Verlag.

Frankfurt, H. G. (1969). Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66(23):829–839.

Goodale, M. and Milner, D. (2004). *Sight Unseen*. Oxford University Press.

Gurwitsch, A. (1964). *Field of Consciousness*. Duquesne University Press.

Gurwitsch, A. (1985). *Marginal Consciousness*. Ohio University Press. Embree, L. E.

Haggard, P. (2001). The psychology of action. *British Journal of Psychology*, 92(1):113–128.

Hardcastle, V. G. (1997). Attention versus consciousness: A distinction with a difference. *Cognitive Studies*, 4(3):356–366.

Harel, J. and Koch, C. (2009). On the optimality of spatial attention for object detection. *Attention in Cognitive Systems*, pages 1–14.

Hegel, G. W. F. (1971). *Philosophy of Mind*. Clarendon Press. (Trans. W. Wallace).

Heidegger, M. (1978). *Being and Time*. Wiley-Blackwell.

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., and Heeger, D. (2010). When size matters: attention affects performance by contrast or response gain. *Nature Neuroscience*, 13(12):1554–1559.

Hirst, W., Spelke, E., Reaves, C., Caharack, G., and Neisser, U. (1980). Dividing attention without alternation or automaticity. *Journal of Experimental Psychology*, 109(1):98–117.

Humphrey, N. K. (1974). Vision in a monkey without striate cortex: a case study. *Perception*, 3(3):241–255.

Hunter, J. and Csikszentmihalyi, M. (2000). The phenomenology of body-mind: The contrasting cases of flow in sports and contemplation. *Anthropology of Consciousness*, 11(3-4):5–24.

Hurley, S. (2002). *Consciousness in Action*. Harvard University Press.

Husserl, E. (1975). *Experience and Judgment*. Northwestern University Press. (Trans. J. S. Churchill and K. Ameriks).

Husserl, E. (2001). *Analyses Concerning Passive and Active Synthesis. Lectures on Transcendental Logic*. Dordrecht: Kluwer Academic Publishers. (Trans. A. J. Steinbock).

Jackson, F. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32:127–136.

Jackson, F. (1986). What Mary didn't know. *Journal of Philosophy*, 83:291–295.

James, W. (1981). *The Principles of Psychology*. Harvard University Press.

James, W. (2010). *The Principles of Psychology*. H. Holt and company.

Janson, H. and Janson, A. (2003). *History of Art: The Western Tradition*. Prentice Hall.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In *Attention and Performance IX*, volume 9, pages 187–203.

Kahneman, D. (1973). *Attention and effort*. Prentice-Hall, Inc.

Kant, I. (1996). *Critique of Pure Reason*. Hackett Publishing. (Trans. W. S. Pluhar).

Kentridge, R. W., Heywood, C. A., and Weiskrantz, L. (1999). Attention without awareness in blindsight. *Proceedings of the Royal Society B: Biological Sciences*, 266(1430):1805.

Kentridge, R. W., Heywood, C. A., and Weiskrantz, L. (2004). Spatial attention speeds discrimination without awareness in blindsight. *Neuropsychologia*, 42(6):831–835.

Kim, M. S. and Cave, K. R. (2001). Perceptual grouping via spatial selection in a focused-attention task. *Vision Research*, 41(5):611–624.

Kinsbourne, M. (1987). Mechanisms of unilateral neglect. In Jeannerod, M., editor, *Neurophysiological and Neuropsychological Aspects of Spatial Neglect*. Elsevier.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Sciences*, 4(4):138 – 147.

Knau, H. and Spillman, L. (1997). Brightness fading during *Ganzfeld* adaptation. *Journal of the Optical Society of America*, 14(6):1213–1222.

Koch, C. and Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11(1):16–22.

Kosslyn, S. (1980). *Image and mind*. Harvard University Press.

Kouider, S., de Gardelle, V., and Dupoux, E. (2008). Partial awareness and the illusion of phenomenal consciousness. *Behavioral and Brain Sciences*, 30(5-6):510–511.

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1):12–18.

Lamme, V. A. F. and Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11):571–579.

Lapoujade, D. (2005). The normal and the pathological in Bergson. *Modern Language Notes*, 120(5).

Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology*, 21(3):451–468.

Levin, D., Takarae, Y., Miner, A., and Keil, F. (2001). Efficient visual search by category: Specifying the features that mark the difference between artifacts and animals in preattentive vision. *Perception & Psychophysics*, 63(4):676.

Li, F., VanRullen, R., Koch, C., and Perona, P. (2003). Natural scene categorization in the near absence of attention: further explorations. *Journal of Vision*, 3(9):331.

Li, F. F., VanRullen, R., Koch, C., and Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences*, 99(14):9596.

Ling, S. and Carrasco, M. (2007). Transient covert attention does alter appearance: A reply to Schneider (2006). *Attention, Perception, & Psychophysics*, 69(6):1051–1058.

Locke, J. (2004a). *An Essay Concerning Human Understanding*, volume 2. Kessinger Publishing.

Locke, J. (2004b). *An Essay Concerning Human Understanding*, volume 1. Kessinger Publishing.

Lutz, A., Slagter, H., Dunne, J., and Davidson, R. (2008). Attention regulation and monitoring in meditation. *Trends in Cognitive Sciences*, 12(4):163–169.

Mack, A. and Rock, I. (1998a). *Inattentional Blindness*. MIT Press.

Mack, A. and Rock, I. (1998b). Inattentional blindness: Perception without attention. In Wright, R. D., editor, *Visual Attention*, volume 8, pages 55–76. Oxford University Press.

Mansell, W., Harvey, A., Watkins, E. R., and Shafran, R. (2008). Cognitive behavioral processes across psychological disorders: a review of the utility and validity of the transdiagnostic approach. *International Journal of Cognitive Therapy*, 1:181–191.

McDowell, J. H. (1996). *Mind and World*. Harvard University Press.

Merickle, P. M. and Joordens, S. (1997). Parallels between perception without attention and perception without awareness. *Consciousness and Cognition*, 6(2-3):219–236.

Merleau-Ponty, M. (1962). *Phenomenology of Perception*. Routledge. (Trans. C. Smith).

Michel, F. and Henaff, M. (2004). Seeing without the occipito-parietal cortex: Simultagnosia as a shrinkage of the attentional visual field. *Behavioral Neurology*, 15.

Milliken, B. and Tipper, S. P. (1998). Attention and inhibition. In Pashler, H., editor, *Attention*. Psychology Press.

Minati, L., Rosazza, C., D'Incerti, L., Pietrocini, E., Valentini, L., Scaioli, V., Loveday, C., and Bruzzone, M. (2008). FMRI/ERP of musical syntax: comparison of melodies and unstructured note sequences. *Neuroreport*, 19(14):1381.

Minati, L., Salvatoni, L., Rosazza, C., Pietrocini, E., Visani, E., Panzica, F., Scaioli, V., Avanzini, G., and Franceschetti, S. (2010). Event-related potential (ERP) markers of melodic processing: The N2 component is modulated by structural complexity, not by melodic 'meaningfulness'. *Brain Research Bulletin*.

Mole, C. (2008). Attention and consciousness. *Journal of Consciousness Studies*, 15(4):86–104.

Mole, C. (2011a). *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford University Press.

Mole, C. (2011b). The metaphysics of attention. In Mole, C., Smithies, D., and Wu, W., editors, *Attention: Philosophical and Psychological Essays*. Oxford University Press.

Neumann, O. (1987). Beyond capacity: A functional view of attention. In Heuer, H. and Sanders, A. F., editors, *Perspectives on Perception and Action*. Lawrence Erlbaum.

Nobre, A. C., Sebestyen, G. N., Gitelman, D. R., Mesulam, M. M., Frackowiak, R. S., and Frith, C. D. (1997). Functional localization of the system for visuospatial attention using positron emission tomography. *Brain*, 120(3):515.

Noë, A. and O'Regan, J. K. (2000). Perception, attention, and the grand illusion. *Psyche*, 6(15):6–15.

Norman, D. and Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In Davidson, R. J., Schwarz, G. E., and Shapiro, D., editors, *Consciousness and Self-regulation: Advances in Research and Theory*, volume 4, pages 1–18. Plenum Publishing Corporation.

O'Craven, K. (2005). Neural evidence for object-based attention. In Itti, L., Rees, G., and Tsotsos, J. K., editors, *Neurobiology of Attention*. Academic Press.

Oliva, A. (2005). Gist of the scene. In *Neurobiology of Attention*. Elsevier Academic Press.

Oliva, A. and Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in Brain Research*, 155:23–36.

Pashler, H. E. (1999). *The Psychology of Attention*. MIT Press.

Pavlopoulou, C. and Yu, S. (2010). Classification and feature selection with human performance data. In *IEEE International Conference on Image Processing*, pages 1557–1560. IEEE.

Peelen, M. V., Heslenfeld, D. J., and Theeuwes, J. (2004). Endogenous and exogenous attention shifts are mediated by the same large-scale neural network. *Neuroimage*, 22(2):822–830.

Persaud, N., McLeod, P., and Cowey, A. (2007). Post-decision wagering objectively measures awareness. *Nature Neuroscience*, 10(2):257–261.

Posner, M. I. (1994). Attention: The mechanisms of consciousness. *Proceedings of the National Academy of Sciences*, 91:7398–7398.

Posner, M. I. and Cohen, Y. (1984). Components of visual orienting. In *Attention and Performance X*, volume 32, pages 531–556.

Potter, M. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology*, 2(5):509.

Prinz, J. (2000a). The ins and outs of consciousness. *Brain and Mind*, 1(2):245–256.

Prinz, J. (2000b). A neurofunctional theory of visual consciousness. *Consciousness and Cognition*, 9(2):243–259.

Prinz, J. (2011). On the (dis)unity of consciousness. *Online Consciousness Conference*. URL: http://consciousnessonline.wordpress.com/2011/02/14/on-the-disunity-of-consciousness/.

Raffman, D. (1993). *Language, Music, and Mind*. MIT Press.

Reddy, L., Reddy, L., and Koch, C. (2006). Face identification in the near-absence of focal attention. *Vision Research*, 46(15):2336–2343.

Reddy, L., Wilken, P., and Koch, C. (2004). Face-gender discrimination is possible in the near-absence of attention. *Journal of Vision*, 4(2).

Rensink, R. A., O'Regan, J. K., and Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8:368–373.

Ro, T. and Rafal, R. (1996). Perception of geometric illusions in hemispatial neglect. *Neuropsychologia*, 34(10):973–978.

Robertson, I. H. and Marshall, J. C. (1993). *Unilateral Neglect: Clinical and Experimental Studies*. Psychology Press.

Robertson, L. C. (2003). Binding, spatial attention and perceptual awareness. *Nature Reviews Neuroscience*, 4(2):93–102.

Rosen, A. C., Rao, S. M., Caffarra, P., Scaglioni, A., Bobholz, J. A., Woodley, S. J., Hammeke, T. A., Cunningham, J. M., Prieto, T. E., and Binder, J. R. (1999). Neural basis of endogenous and exogenous spatial orienting: A functional MRI study. *Journal of Cognitive Neuroscience*, 11(2):135–152.

Sandberg, K., Timmermans, B., Overgaard, M., and Cleeremans, A. (2010). Measuring consciousness: Is one measure better than the other? *Consciousness and Cognition*.

Schneider, K. A. (2006). Does attention alter appearance? *Perception & Psychophysics*, 68(5):800.

Schwitzgebel, E. (2007). Do you have constant tactile experience of your feet in your shoes?: Or is experience limited to whats in attention? *Journal of Consciousness Studies*, 14(3):5–35.

Schyns, P. and Oliva, A. (1994). From blobs to boundary edges: Evidence for time- and spatial-scale-dependent scene recognition. *Psychological Science*, 5(4):195.

Schyns, P. and Oliva, A. (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, 69(3):243–265.

Seth, A. (2008). Post-decision wagering measures metacognitive content, not sensory consciousness. *Consciousness and Cognition*, 17(3):981–983.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423.

Simons, D. J. (2000). Attentional capture and inattentional blindness. *Trends in Cognitive Sciences*, 4(4):147–155.

Simons, D. J. and Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28:1059–1074.

Singer, W. and Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18:555–586.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74(11).

Spratling, M. W. and Johnson, M. H. (2004). A feedback model of visual attention. *Journal of Cognitive Neuroscience*, 16(2):219–237.

Suchy-Dicey, C. (2011). Inductive parsimony and the Methodological Argument. *Consciousness and Cognition*, In Press.

Titchener, E. B. (1910). Attention as sensory clearness. *The Journal of Philosophy, Psychology and Scientific Methods*, 7(7):180–182.

Tononi, G. (2005). Consciousness, information integration, and the brain. In Laureys, S., editor, *Progress in Brain Research*. Elsevier.

Treisman, A. M. (1988). Features and objects: The Fourteenth Bartlett Memorial Lecture. *Journal of Experimental Psychology*, 40(2):201–237.

Treisman, A. M. and Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136.

Treue, S. (2004). Perceptual enhancement of contrast by attention. *Trends in Cognitive Sciences*, 8(10):435–437.

Wackermann, J., Putz, P., and Allefeld, C. (2008). Ganzfeld-induced hallucinatory experience, its phenomenology and cerebral electrophysiology. *Cortex*, 44:1364–1378.

Watzl, S. (2011). Attention as structuring the stream of consciousness. In Mole, C., Smithies, D., and Wu, W., editors, *Attention: Philosophical and Psychological Essays*. Oxford University Press.

Wittgenstein, L. (2009). *Philosophical Investigations*. John Wiley and Sons. (Trans. G.E.M. Anscombe, P.M.S. Hacker, and J. Schulte).

Wolfe, J. (1999). Inattentional amnesia. *Journal of Mental Imagery*, 29(2-4):71–94.

Wolfe, J. M., Cave, K. R., and Franzel, S. L. (1989). Guided search: An alternative to the feature integration model for visual search. *Journal of Experimental Psychology*, 15(3):419–433.

Wolfe, J. M. and Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5:495–501.

Wu, W. (2011a). Attention as selection for action. In Mole, C., Smithies, D., and Wu, W., editors, *Attention: Philosophical and Psychological Essays*. Oxford University Press.

Wu, W. (2011b). Confronting many-many problems: Attention and agentive control. *Nous*, 45(1):50–76.

Yantis, S. (1998). Control of visual attention. In Pashler, H., editor, *Attention*. Psychology Press.

Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Sciences*, 7:214–218.

Curriculum Vitae

Carolyn Suchy-Dicey

Philosophy Department
Boston University
745 Commonwealth Avenue
Boston, MA 02215

phone: (443) 535-3777
email: suchyd@bu.edu
url: http://people.bu.edu/suchyd

Areas of Research

AOS: Philosophy of Mind and Philosophy of Cognitive Science

AOC: Philosophy of Science, Philosophy of Physics, Phenomenology, Philosophy of Perception, and Medical Ethics

Education

2006-2012 PhD Philosophy candidate, Boston University
Dissertation Title: *Attention, Subject, and World*
Defense Date: September 12th, 2011
Committee: Daniel Dahlstrom, Walter Hopp, Ned Block, John Campbell

MS Psychology expected, Boston University

2000-2004 MA Philosophy, University of St Andrews

Publications

"The Standard Theory of Conscious Perception," forthcoming
Frontiers in Consciousness Research

"Inductive Parsimony and the Methodological Argument," 2011
Consciousness and Cognition

"Review of Thomas Metzinger's *The Ego Tunnel*," 2010
Journal of Consciousness Studies

"It Takes Two: Ethical Dualism and the Vegetative State," 2009
Neuroethics

Invited Talks

2010

"Re-thinking the Active-Passive Distinction in Attention"
Association for the Scientific Study of Consciousness
Vision Sciences Society
Boston University Mind and Brain Society

"Experiential Richness: Beyond the Grasp of Attention?"
Society for Philosophy and Psychology

"Inductive Skepticism and the Methodological Argument"
Online Consciousness Conference

2009

"Epistemic Restraint: An Antidote to Zombie Poison"
Association for the Scientific Study of Consciousness
Metaphysics of Mind Graduate Conference
Rocky Mountain Graduate Conference

"Attention and Consciousness"
Boston University Undergraduate Philosophy Association

2008

"What the Gist? A Case Study in Perception and Attention"
Association for the Scientific Study of Consciousness

Teaching

As Instructor

Introduction to Philosophy, Northeastern University, 2008-2010
Philosophy of Science, Boston University, 2009

As Teaching Fellow

Introduction to Logic, with Jaakko Hintikka, 2009
Medical Ethics, with Alfred Tauber, 2008
Introduction to Logic, with Judson Webb, 2008
Introduction to Ethics, with Susanne Sreedhar, 2007
Critical Reasoning, with Alisa Bokulich, 2007
Virtue Ethics, with Charles Griswold, 2006

Service to the Profession

Founder and Organizer of Neuphi (http://www.neuphi.com), 2007-2012
Founder and Organizer of IGCC (http://www.bu.edu/conscious), 2008-2012
Associate Member of Takeo Watanabe's Vision Lab, 2009-2012
Reviewer, *Philosophy of Science Journal*, 2011
Organizer of Sean Kelly's Perception Lab, 2010
Speaker, Boston University's Professional Development Panel, 2010
Editor, *Re-examining the Quantum-Classical Relation* by Alisa Bokulich, 2008

Grants, Honors, & Awards

Boston University Research Assistanceship, 2011-2012
Boston University Dissertation Fellowship, 2009-2011
Boston University Travel Grant, 2008-2011
Shimony Prize, 2010
Boston University Teaching Fellowship, 2006-2009
Saltire Essay Prize, 2004
P.E.O. Janet H. Griswold Scholarship, 2003
University of St Andrews Academic Bursary, 2000-2004

Graduate Courses

Seminar on Consciousness*, Daniel Dennett (Tufts)
A Systems Neuroscience Approach to Consciousness*, Dan Pollen (Harvard)
Philosophy of Cognitive Science*, Tian Yu Cao
Philosophy of Psychology*, Sean Kelly (Harvard)
Philosophy of Mind, Peter Bokulich
Philosophy of Perception, Susanna Siegel and Alex Byrne (Harvard and MIT)
Psychology of Perception, Takeo Watanabe
Phenomenology and Perception*, Daniel Dahlstrom
Phenomenology*, Walter Hopp
Seminar on Meaning, From Analytic Philosophy to Phenomenology*, Daniel Dahlstrom
Epistemology*, Walter Hopp
Seminar on Free Will, Daniel Dahlstrom
Philosophical Foundations of Quantum Theory, Alisa Bokulich and Gregg Jaeger
Wittgenstein, Jaakko Hintikka
Philosophy of Physics, Tian Yu Cao
Aristotle's Metaphysics, David Roochnik
Rights, Natural Law, and their Discontents, Aaron Garrett
(* indicates audited courses)

Language Proficiency

French (reading)