Perceptual Representation

Content, Constancies and Phenomenal Consciousness

Christian Pettersen Lains

Master Thesis

IFIKK, Faculty of Humanities
University of Oslo

Spring 2013

Perceptual Representation

Do perceptual constancies and phenomenal consciousness help provide an account of how subjects perceive things in the world? In particular, *how* can they help make sense of the idea that perceptual content is given in the perspective of a subject? In this essay I work towards understanding, in a scientifically grounded way, how perceptual constancies and phenomenal consciousness - jointly - can make sense of this idea in a very precise way.

2013

Christian Pettersen Lains

Perceptual Representation

http://www.duo.uio.no

Print: CopyCat, Forskningsparken

Summary

I investigate perceptual representation by taking Tyler Burge's account of perception as a point of departure. Drawing on cognitive psychology, Burge's account explains how the content of perception, partly understood in terms of perceptual constancies, can relate to objects in the environment. I argue, however, that Burge lacks a detailed account of how the perceptual content relates to the subject's perspective. By analyzing Burge's position we find that we need to supplement Burge's account to get a full understanding of how, specifically, the perceptual content relates to the subject's perspective. Following Burge's methodology, I want to ground such a supplementary account in cognitive psychology. In particular, and contrary to Burge, I argue that such a supplementary account should be given partly by appealing to the science of phenomenal consciousness. By supplementing Burge with this alternative account, I show that it is *possible* to give a scientifically grounded, rigorous account of how perception as given in the perspective of a subject can be about objects.

Preface

For as long as I can remember, the quality of perceptual experience has been a source of inspiration. In vivid dreams, beautiful art, virtual worlds and everyday life, the richness of our experience radiates a level of detail beyond words. In my teen years I followed psychology closely to get a glimpse of how such diamonds of experiences might fit reasonably into the scientific order. Only in my Bachelor years, however, did I begin to suspect one could understand such matters in a rigorous, scientific way. General theories of perception and studies on the correlates of consciousness were published in reputable, scientific journals. The impenetrable mist occluding the relation between mind and brain was beginning to evaporate.

More recent developments have only strengthened my belief that the mist is clearing. And yet, only the first few steps have been taken. The same is true of this essay, which represents only the first few steps towards the frontier it seeks to explore. Much more will need to be understood. In anchoring my essay in Tyler Burge's recent work I have also come to understand the immense work that is required to be faithful to aspects of both scientific and philosophical terminology, methodology and motivations.

I want to thank my supervisor, Einar Duenger Bohn, for his helpful comments and tireless support in straightening out my poetic proclivities. Without him, this essay would have been a lot less intelligible. I want to thank Joakim Haraldsen and Knut Dæhli for their endless supply of stimulating discussions. I want to thank Frank O. Barel, Bjørn T. Ramberg and Carsten Hansen for their passionate courses on Tyler Burge. In particular, I want to thank Frank and Carsten for their insightful comments on my essay. I also want to thank Jørgen Dyrstad for stimulating comments and discussions. Finally, I want to thank Tyler Burge for his generous, candid and illuminating replies to some exegetical doubts and arguments.

I also want to thank my friends, Espen Thoen, Jo Flottorp Ness, Matti Richoux and Lars Flottorp Ness, for being awesome. I want to thank my family, for being supportive (and also awesome). In particular I want to thank my mother Grete Pettersen, my brother Manuel A. Pettersen Lains and my sister-in-law Karin Sooväli for proofreading. Finally, I want to thank my girlfriend, Camilla Karlsen, for her endless love, affection and wonderful cuisine, as well as her uncanny ability to live unperturbed with a philosopher prone to rants and reclusion.

Contents

1 Perceptual Content

1.1 Introduction

I see. In my field of vision various objects are present and various things are going on. Scientifically, we know that when one sees, causal processes are involved; a surface reflects light, one's retina becomes active in a certain way, and the visual system becomes active in a certain way. If all goes well, one sees various objective things.[1]

Things do not always go well.

Case 1: Suppose I see a red apple in a bowl of fruit. I walk towards the bowl of fruit and pick up the red apple. Suppose my identical twin is in much the same situation. As my twin walks towards his bowl of fruit, however, he discovers that what seemed to him to be a red apple was actually a green one. By some quirky optical feature of the scene there seemed to be a red apple where there was in fact a green one.

Things can visually seem the same to a subject in cases where the object is as it seems (often called the good case of perception) and in cases where the object is not as it seems (often called the bad case of perception). One may interpret this in numerous ways according to ones favorite philosophical framework. On the one hand, we want to capture the intuition that, in the good case, we perceive various objective things. On the other hand, we want to capture the intuition that there can be bad cases where things (in some sense) seem the same to a subject as it would in a good case.

I will *assume* that there can be phenomenal experiences that are identical in cases of perception, illusion and hallucination.[2] Phenomenal experiences are experiences that are phenomenally conscious – there is something it is like to have such experiences. I restrict talk of experiences, and thus to phenomenal experiences, to experiences that are - or to the subject seems to be - perceptual. If I intend conscious experiences in general I will simply use the term 'phenomenal consciousness.' Further, one says that phenomenal consciousness and phenomenal experiences have phenomenal characters – "properties that type the experience

[1] Structure of the introduction is partly inspired by David Lewis's (Lewis, 1980)

[2] This is adapted from the *Common Factor Principle* found in (Fish, 2010, p. 4).

by what it is like to undergo them" (Fish, 2010, p. 39). There can both be identity of what it is like to have phenomenal experiences, and identity of specific phenomenal characters of such phenomenal experiences.

In relation to *Case 1* I want to highlight a *Subjective Indistinguishability* (SI): A phenomenal character, identical to that which is had in the good case of perception, can occur in bad cases of illusion/hallucination and make it seem to the subject that what is perceived in the good case is what is perceived in the bad illusory/hallucinatory case. I call it *Subjective Indistinguishability* because the subject, as it were, cannot distinguish the cases.

If my twin loves red apples, but not green ones, (SI) can help explain why my twin moved towards the green apple in the following way: A phenomenal character made it seem to him that what he perceived in the bad case is what he perceives in the good case. In particular, I assume that a phenomenal character is making it seem to the subject that a *red apple* is perceived, a phenomenal character that tends to occur when my twin sees red apples.

Phenomenal experience can stay the same even though what is actually perceived is different. Consider now the idea that it is possible to be presented with the same thing in different ways. Under a different light we can see the same things objectively. Phenomenal experience can differ while the same things are seen. How far can we stretch this?

Case 2: Suppose, again, that I see a red apple in a bowl of fruit. I walk toward the bowl of fruit and pick up the red apple. Suppose I have another identical twin. Suppose he is in much the same situation. This twin, however, has phenomenal experiences of colors that are inverted relative to mine, and so has a phenomenal experience of what I call green whenever he sees red things. In spite of this difference, it is plausible to suppose that my color-inverted twin and I could have roughly the same functional relations to the environment.[3]

Note that in addition to rough functional identity, it is plausible to suppose that my twin and I have roughly the same representational capacities. Note that I will assume that perceptual experiences are representational. I will not try to argue that perceptual experiences are

[3] One might question to what degree this twin and I (physically) could be functionally identical. I think this partly depends on resolving questions involving the nature of phenomenal consciousness that I will not dwell on here. The thought experiment is of course inspired by Ned Block and Sidney Shoemakers' inversion thought experiments (Block, 1990a; Shoemaker, 1982)

representational. Given that disclaimer, I think my twin and I have the same representational capacities in the following sense: Roughly, whenever I have a perceptual experience with the phenomenal character green, things will be as they seem if things are *equal to how* things are when they are as they seem to my twin whenever he has a perceptual experience with the phenomenal character red. On the representational view of perceptual experience, things are sometimes as they seem and sometimes not, because perceptual experience is representational. In particular, perceptual experience has representational content that types the conditions under which things are as they seem, as distinct from the conditions under which things are not as they seem. I call such conditions *accuracy conditions*. Perceptual representational content can be accurate or not accurate, depending on whether things are as they seem or not.[4] One can also think of accuracy and inaccuracy as being a matter of degree. I will not discuss such further details here. I further assume that it is in the very nature of perception to have perceptual representational content (that are either accurate or inaccurate).[5]

In contrast to Subjective Indistinguishability (SI), the thought experiment *Case 2* points the way to an *Objective Indistinguishability* (OI): Subjects can perceptually represent the same thing by means of different phenomenal characters, and these phenomenal characters can be used differently by different subjects (e.g. my twin and I use the phenomenal character *red* differently). I call it *Objective Indistinguishability* because an objective observer, as it were, cannot distinguish between the two cases.[6] In light of this it seems plausible that no particular phenomenal character of perceptual experience is *necessary* to type a given perceptual

[4] Susanna Siegel calls this *The Content View*: "All visual perceptual experiences have contents" where this content is understood as 1) as true or false, 2) conveyed to the subject by his/her experience (Siegel, 2010, p. 28). She then goes on to argue that content should be understood as accuracy conditions rather than as *simply* being true/false (Siegel, 2010, p. 30). When Susanna Shellenberg speaks of the *Content Thesis* she means the thesis "that perceptual experience is fundamentally a matter of representing the world as being a certain way" (Schellenberg, 2011, p. 714). This latter view is stronger insofar as she appeals to what perceptual experience *fundamentally* is. For our purposes, note that Susanna Siegel's view, while weaker in the sense just mentioned, is stronger insofar as it explicitly requires a 'subject' that the content is conveyed to. As we go along, I formulate and assume a view that is stronger in both senses.

[5] An article by (Ganson, Bronner, & Kerr, 2012, p. 1) use the "Content View," to mean the view that "Perceptual states are, at their core, representations with contents that are either accurate or inaccurate."

[6] I actually do believe that it is possible to distinguish the two cases by a sufficiently advanced theory of the brain and of phenomenal consciousness in particular. Perhaps I should have called them *First Person-*, and *Third Person Indistinguishability* respectively.

content, because another phenomenal character might equally be used as vehicle for that very same perceptual content.

In general, if a philosophical theory aims to give an explanation of perception it must explain how the *subject* that perceives can relate to the object perceived. This demand of perception is surely one most *philosophers* would accept. Part of the reason we would accept it, we should recognize, is that perception has been a point of focus in philosophy partly because it takes us from the subjective to the objective. Contrast: An artificial intelligence researcher who finds some deep unifying principle underlying general pattern-recognition might insist that the nature of perception is essentially just that underlying principle. For this person, the invocation of *subjects* might seem inessential to the nature of perception.

However, in philosophical circles, saying that a theory of perception must explain how the perceiving *subject* can relate to an object perceived is not saying much. Philosophers focusing on perception are typically motivated to explain the relation between the subjective and the objective. I submit that a philosophical theory of perception should aim to explain both of these essential[7] components:

Perceptual Duality (PD): First, perception is fundamentally attributed to a subject as a whole. Contrast this attribution to the *whole* subject with attribution of "perception" to the eyes, processing in the brain, or other subsystems or sub states of the individual. Second, perception is fundamentally about things in the world. Contrast this with "perception" being, for instance, about sense-data or phenomenal characters.

Note the following terminology. When attribution is to the *whole* subject in this demanding sense, I will say that it is *for* the whole subject, or for the subject as a whole.[8] When perception is about things in the world, I simply say that it is *about* objects, where this

[7] In the introduction I sometimes use the word essential. Later I introduce the term 'constitutive.' It would be too distracting to introduce here.

[8] One might use other kinds of terminology getting at roughly the same idea; e.g. that perception is *to* or *for* or *presented* or *conveyed* to the subject, with any number of mixes between them. The specificity of these terminological possibilities make an interesting contrast to the lack of any solid, scientific understanding of this relation – especially in perception, as we will see later in this essay. (The same can be said of the myriad of ways to speaks of consciousness, only a few of which give us any insight – the rest only contributing to the confusion. In that case, too, I simply chose a few prominent definitions and stick to them.)

includes relations, properties and kinds.[9] I use the word *fundamentally* to signal that perceptions need not always be for a subject as a whole, and likewise need not always be *about* objects. However, Perceptual Duality is the idea that these non-fundamental cases are relative to the explanation of the fundamental ones.

I call it Perceptual *Duality* because I think perception should roughly explain how it can be fundamentally *for* a subject, and fundamentally *about* things in the world. Metaphorically, the picture I mean to evoke by *duality* is that, in the *most* fundamental case,[10] the perceptual experience at the same time faces both the subject in the interior world and the object in the external world.[11]

Subjective and Objective Indistinguishability, (SI) and (OI), make it difficult to explain how a theory can satisfy (PD) by appealing to phenomenal consciousness, in particular, to phenomenal characters. From (SI) we saw that any given phenomenal character of perceptual experience seems *unnecessary* to type a given perceptual content. For instance, it seems that the phenomenal character red and the phenomenal character green could, in principle, type the same perceptual content. From (OI) we saw that the phenomenal character of a perceptual experience is *insufficient* to type a perceptual content. For instance, it seems that the phenomenal character red could, in principle, type two different perceptual contents.[12]

In much of this essay will explore the extent to which it is possible to explain how perception is fundamentally for a whole subject without appealing to phenomenal consciousness. Towards the end I develop the idea that phenomenal consciousness plays a fundamental part in explaining how perception is fundamentally for a whole subject. Somewhat against the spirit of (SI) and (OI), but still consistent with them, I will also hold that phenomenal consciousness plays a fundamental part in typing perceptual representational content.

---

[9] One might distinguish between *intentionality* or *aboutness* on the one hand, and representation on the other. The literature is fraught with terminological battles. I try my best not to engage in them here.

[10] The most fundamental case is of course when it is jointly the case that the whole subject perceives and the subject perceives objects (nothing much rests on the plural here) in the external world.

[11] Some people find this picture engaging and interesting. Others find it worthless and ambiguous. For me, it is one of those striking puzzles we find immediately when studying perception; the distance between those highly intelligent thinkers who adopt a sense-datum theory, to those highly intelligent thinkers who adopt a naïve realist theory of perception signal its force in splitting people as to its interpretation at the very outset.

[12] Of course I have in mind here two contents in two different individuals as in the example above.

1.2 What will I discuss?

The overarching goal is to explain how Perceptual Duality (PD) might be possible in the context of several assumptions. I assume that there can be phenomenal experiences that are identical in cases of perception, illusion and hallucination. I assume that Subjective and Objective Indistinguishability capture important intuitions about the nature of perceptual experience and phenomenal experience in particular. I assume that it is in the very nature of perception to have perceptual representational content (that are either accurate or inaccurate). I will also assume that having perceptual representational content is part of the very nature of perception. In particular, that an explanation of perceptual representational content partly explains the nature of the perspective of the subject perceiving. I will also assume that identical phenomenal experiences have an underlying psychological state in common. In sum, the overarching goal is to explain how (PD) might be possible, given all these assumptions.

In particular, I will focus on perceptual representational content ('perceptual content' for short) as the "mediator" between the object and the subject. Further, I will assume that Burge explains, in a scientifically grounded way, how perceptual content is possible insofar as it is *about* objects. In explaining Burge's account I will assume that he gets most of it right. I will only take issue over his lack of a scientifically grounded explanation of how the perceptual representational content is possible insofar as it is fundamentally *for* the subject as a whole.

My particular focus is to explain, in a scientifically grounded way, how perceptual representational content is possible insofar as it is *for* a whole subject. In particular, I want to supplement Burge's account with the thesis that phenomenal consciousness is fundamental to perception. By doing this, the motivation is to contribute to the overarching goal of explaining how PD might be possible, given all these various assumptions.

1.3 How will I discuss it?

In giving a theory of how a whole subject can perceive objects through perceptual representational content, we have a decidedly philosophical outlook. However, in giving such a theory, there are some scientific constraints I will stick to. In particular, I will accept two constraints on a philosophical theory aiming to explain some aspect of mental representation. Both of these constraints are adapted from Robert Cummins' constraints with the same name (R. Cummins, 1996, p. 2).

Explanatory Constraint (EC)**:** In a minimal sense, we should make sure that our philosophical theory is not in tension with findings in cognitive psychology. In an ideal sense, we should help make philosophical sense of the explanatory appeals that cognitive psychology makes to mental representation. *Implementation Constraint* (IC)*:* In particular, we should make sure that - minimally - our philosophical theory of mental representation is not in conflict with the best stories about what actually does the representing in the mind/brain. And ideally, we should help make philosophical sense of what cognitive science takes to be implementing representations in the mind/brain.

I hasten to note that the adherence to both of these principles should be tempered by the fact that cognitive psychology is still in its infancy.[13] For instance, there is no generally accepted story about what actually does the representing in the mind/brain, and the notion of 'mental representation' is still very open.

Given these constraints it is natural to start with Tyler Burge, a philosopher who is deeply interested in the philosophy of perception and the psychology of perception. In particular, drawing on his grasp of both these fields, Burge has given what has been hailed as "the most sustained and sophisticated defense of the content view to date" (Ganson et al., 2012, p. 2). To make his philosophical points about perception, Burge interacts deeply with cognitive psychology, and does so in a principled, methodological fashion that is in accordance with (EC) and (IC). For the most part I adopt Burge's framework and methodology.

We can conceptually carve up perception into a relation between *content and subject* on the one hand, and *content and object* on the other hand. I will argue that Burge does a beautiful job of accounting for the latter but lacks a good account of the former. In particular, I will argue that phenomenal experience is more tightly knit to the explanation of why we attribute perception to a *whole* subject than Burge thinks.

[13] Broadly construed, psychological science is in its infancy. Cognitive psychology, as such, is a particular frontier that I focus a lot on in this essay. The best stories of what actually does the representing in the mind/brain might, and I think will, partly come from other scientific and mathematical fields. I believe many of the scientific fields that will eventually contribute to this story do not exist as independent disciplines yet.

1.3.1 Structure of the Thesis

This introduction is the first chapter. In chapter two I introduce Burge and explain a central notion of perceptual psychology that he makes use of, namely what are called 'perceptual constancies'. I then show how Burge wants to make use of this perceptual psychological notion in his account of perceptual content. In particular, Burge uses the notion to explain how perception can be *about* things in the external world. I then investigate how Burge relates perceptual content to the subject as a whole. Burge, in his discussion of so-called Individual Subjects, uses phenomenal consciousness to demarcate the difference between the subject as a whole and mere subsystems of such individuals. Partly based on this, I wonder to what extent Burge dissociates phenomenal consciousness and perceptual content. Burge does not think phenomenal consciousness is essential to perceptual content. I discuss what positions his framework "allows" him to take on the relation that holds between them.

In chapter three I look at some of the reasons *why* Burge does not seem to think that it is a big deal to dissociate the phenomenal experience and perceptual content. In particular, I look at how Burge understands the relation between perceptual constancies and the subject as a whole in the absence of phenomenal experience. As it turns out, Burge *argues* against the essential connection between perception and phenomenal experience. Against this I argue that Burge does not have a complete explanation of how the perceptual constancies give rise to a perceptual representational content for the *whole* subject. In contrast to Burge I think that this incompleteness must at least partly be filled by phenomenal consciousness. But how can phenomenal consciousness, and phenomenal experience in particular, be of help?

In the fourth chapter I sketch a view of phenomenal experience as being a kind of phenomenal information. I develop a strong and a weak interpretation of phenomenal information. I then outline Paul Churchland's Map-Indexing Theory of Perception. This theory shows, in fine-grained way, how perceptual constancies might relate to a definite perceptual content. It also makes phenomenal predictions. Finally, I show that Churchland's theory, in combination with both strong and weak phenomenal information, helps solves the incompleteness. Most speculatively, on the strong interpretation, phenomenal consciousness *just is* Integrated Information, which is required to fully make sense of Churchland's theory, in particular, how '*indexing*' is possible.

In the fifth chapter I discuss objections and conclude the essay.

2 Tyler Burge: Perceptual Constancies and Representation

2.1 Introduction

Burge agrees that perceptual content is representational. Terminologically, Burge uses the terminology 'perceptual content' to mean 'perceptual representational content.' Burge also agrees with the Explanatory and Implementation Constraint ((EC) and (IC)). In particular, Burge agrees that a philosophical theory of perception should not be in tension with findings in cognitive psychology, and that it should not be in conflict with the best stories about what sorts of things actually do the representing in the mind/brain. Burge's project is also partly that of making philosophical sense of the explanatory appeals that cognitive psychology makes to mental representation, and the things invoked to explain mental representation.

I will begin with Burge's views on (CV), and then proceed to his views on (EC) and (IC).

2.1.1 The Content View

Burge's recent 2010 book *Origins of Objectivity* has been hailed as giving "the most sustained and sophisticated defense of the content view to date" (Ganson et al., 2012, p. 2). The Content View (CV) is the idea that perceptual content is representational. In particular, this means that perceptual content has accuracy conditions. Burge thinks the very nature of perceptual content to be representational. Burge thinks that perceptual contents types perceptual states.

Burge often lumps intentionality with representation. Burge has in mind something quite strong when he speaks of 'representation' and 'intentionality' (I will use 'representation' for short). In the introduction to *Origins of Objectivity* Burge holds that "Representation - intentionality - is, along with consciousness, the most striking feature of mind," (Burge, 2010, p. 4). This feature of mind is striking because it difficult to understand how mere events in the brain or in our perceptual experience can represent (or be about) things in the world. Burge is interested in the idea that perceptual content is representational only insofar as this is a striking feature of mind. We do not want a definition of representation that trivializes the striking feature of mind. For instance, if a theory claims that thermostats are said to represent or (intentionally) be about the temperature in a room, we might be unsatisfied *insofar* as we

want ‘representation’ or ‘intentionality’ to be important terms in a theory that explains a striking feature of mind. For instance, insofar as we want representation to be attributed to subjects as a whole, it is clear that a thermostat is not a subject in any meaningful sense. Metaphorically, we call such theories *deflationary* because they aim to deflate notions that many believe are filled with deep and interesting explanatory potential.

Further, Burge’s account of perceptual content aims to type the specific *way* in which perception represents. On the one hand, Burge holds that perceptual representational contents type differences in the perspectival *ways* we perceptually represent something (even if *what* is represented is the same in different cases). On the other hand, Burge holds that perceptual representational contents type differences in “context-bound applications - demonstrative occurrences” (Burge, 2010, pp. 390-391, 394, 412). The analogy is to demonstratives like ‘this’ that have contents reflecting what they refers to on each occasion of use.[14]

2.1.2 The Explanatory and Implementation Constraint

Part of the reason why Burge’s *Origins of Objectivity* is hailed as the most sophisticated defense of the content view to date is that Burge takes perceptual psychology seriously. Perceptual psychology is a branch of cognitive psychology and, as such, makes essential use of ‘mental representations’ as (EC) constrains us to accommodate. Note that Burge thinks serious mainstream perceptual psychology *just is* cognitive perceptual psychology.[15]

In particular, Burge appeals to perceptual psychology in the lineage of the cognitive psychologist David Marr (Daivd Marr, 1982).[16] Cognitive psychology partly reduces mental

[14] Burge thinks “representation of physical entities in language and thought is the way it is largely because representation in perception is the way it is” (Burge, 2009b, p. 293). “For example, the Kripke-Donnellan points about reference of names have rather obvious counterparts about perception” (Burge, 2009b, p. 293). Burge is a systematic philosopher who is investigating the nature of perceptual representation partly to understand the nature of representation in general.

[15] For instance, Tyler Burge notes that “the central mode of explanation in this science [mainstream work in perceptual psychology] takes representational state and transformation that produces representational states to be the central explanatory notions” (Burge, 2010, p. 298). Brackets are mine unless noted otherwise.

[16] Other proponents of cognitivism include Noam Chomsky (Chomsky, 1995), Jerry Fodor (Fodor, 1990), and, more recently, Paul Churchland (P. M. Churchland, 2012) who thinks of representation in terms of neural nets. While I focus on cognitive psychology, it is clear that it is not exhaustive. For instance, ecological or enactive approaches that focus on interaction between the environment and mental states, such as Andy

states to states that involve cognitive processing and representation. It is typically couched in a functionalist framework taken to be computational.[17] For instance, David Marr's classic *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (Daivd Marr, 1982) proclaims by its very title the cognitive approach to perception. Burge argues persuasively that perceptual psychology has become a "serious" and "well-supported, mathematically rigorous, mature science" (Burge, 2010, pp. 87, 297).

The Implementation Constraint (IC) urges us to take into account what sorts of things actually do the representing in the mind/brain. As noted, there is still no generally accepted story about what sorts of things actually do the representing in the mind/brain. However, this is not to say that there is not progress. For instance, Burge is skeptical of giving an explanation of perceptual content in terms of linguistic structures. In particular, Burge holds that we should not force perceptual content to "speak" in terms of propositional content or propositional attitudes.[18] To illustrate how perceptual content could well be structured unlike any type of proposition, Burge has us suppose that perceptual contents "are organized in ways that are structurally isomorphic with a topological or geometrical structure. Think of a map. A map does not have a sentence-like structure. It is a singular noun-like representation that functions to correspond to a piece of geography" (Burge, 2010, p. 540).[19] The idea is that there can be a mathematical, one to one correspondence between the structure of the perceptual content and a geometrical structure.[20] Although there is no generally accepted story of what does the

---

Clark's extended mind thesis (Clark, 2008), or Alva Noë's externalistic approach to understanding phenomenal characters (Noë, 2006). There's also the embodied cognition view. For instance, (Varela, Thompson, & Rosch, 1991), who hold that mental states depend essentially on the body, or, more cautiously (Gallagher, 2005), investigations of how the body shapes the mind.

[17] While it is often taken to be computational, functionalism about mental states is broader than computationalism. Gualtiero Piccinini notes in his (Piccinini, 2004) that computationalism about mental states is an empirical claims about the specific kind of functionalism that is relevant to explain mental states.

[18] Burge states outright that "elements in perception are organized non-propositionally" (Burge, 2010, p. 540).

[19] There is also a positive suggestion here. For instance, Burge comments that elements in "perceptual contents are organized in the structures of various magnitudes, most prominently spatial magnitudes (both topological and geometrical), though the groupings indicated in perceptual systems can be at various levels of abstraction" (Burge, 2010, p. 381).

[20] Specifically, Burge talks about the singular and attributive elements in perception. Perceptual representational contents have singular and attributive elements. I will discuss Burge's terminology in section 2.3.

representing in the mind/brain, certain possibilities are excluded as impossible or at least implausible by ongoing progress in the various relevant sciences.

Burge draws from perceptual psychology the key idea that something called 'perceptual constancy' plays a key role in the explanation of perceptual content. Perceptual constancy, also sometimes called representational constancy, object constancy or constancy phenomenon, is the ability to see familiar objects as constant, or as having constant properties, despite fluctuating variations in, for instance, color, illumination, shape or size. Findings in perceptual psychology substantiate this familiar observation by holding that there are mechanisms in the brain that give rise to such perceptual constancies.

I agree with the general approach Burge takes. I think the assimilation of perceptual psychology works philosophically not only by constraining the space of theories, but also by being active in catalyzing new ways of seeing old problems, such as seeing in new ways how something can represent something else.

2.2 Perceptual Constancies

2.2.1 Introduction

To understand Burge's defense of the Content View one must understand perceptual constancies, i.e. representational constancies. The goal of Burge's *Origins of Objectivity* is to explain the origins of representation. Burge takes "as a working hypothesis" the idea that "representation begins with perception" (Burge, 2010, p. 316). Particularly, that perception is "the most primitive kind of (non-deflated) representation" (Burge, 2010, p. 316).

Burge thinks of the criteria that distinguishes real perception from deflated versions as a "distinctive sort of objectification" (Burge, 2010, p. 10). Further, Burge thinks that these objectifying capacities are best exemplified by a certain finding in perceptual psychology. These objectifying capacities are, according to Burge, best exemplified by perceptual constancies (Burge, 2005, p. 10).[21]

[21] For our purposes we might take objectification to coincide with perceptual constancies. Burge (in private correspondence) notes that there are many higher level forms of objectification that do not rely on perceptual constancies. However, if I understand Burge correctly, the claim in this context is the (possibly a priori, but

For us, the perceptual constancies take care of (EC) and (IC) as they relate to perception. They also go some distance towards establishing the Content View and help explain how Perceptual Duality is possible.

2.2.2 What is a Perceptual Constancy?

A fundamental perceptual psychological fact about human and animal perception is that it holds familiar objects, and properties of objects, constant despite huge variations in the sensory input. In general, we can understand that this is so by considering the fact that our eyes always take in information roughly on a 2 dimensional (2D) surface (our retina), while our perceptual content often features 3 dimensional (3D) objects populating a 3D scene. Since any 2D input stimulus to our eyes vastly underdetermine the 3D object causing this input stimulus, our perceptual system integrates information contained in several such 2D inputs to help determine the 3D object or property causing the various inputs. To get a grip on our 3D world, our perception fundamentally integrates different perspectives.

Fortunately, we have two eyes. But this helps us only a bit.[22] We get more help when our eyes, several times a second, are shifting their position and focus. We know scientifically that each time we shift our gaze, a new set of light rays are absorbed by our eyes. If we look at a screen with some fast action video content running on it, from the side, we can get a sense of just how quickly the pixels on the screen are changing. If we walk in the night and catch eye of a room with the TV on, we get a sense of just how quickly the illumination in the room changes up and down. One pertinent question given that our eyes are constantly absorbing all this fluctuating information is the following; why are our phenomenal experiences, and their underlying states, so stable and constant?[23]

---

empirically defeasible) claim that perceptual constancies are the basic kind of perceptual objectification.

[22] For depth perception the triangulation of point of focus and the two eyes "will be much greater when the observation distance is small; a fact that has led many researchers to conclude that stereopsis is only useful in near/personal or interaction/action space" (Palmisano, Gillam, Govan, Allison, & Harris, 2010).

[23] Scientifically it is commonly known that the underlying states too are relatively constant. This will emerge in the course of the discussion. A related question that I will not go into is the following question: Why do some objectively similar types of stimuli cause vastly different representational states, while some objectively different types of stimuli cause very similar representational states?

In addition to this question, there is a further, related question concerning the relation between two ways one might interpret phenomenal experience. Consider looking at a round table. On the one hand, from what Sean Kelly (Kelly, 2008) and others call the 'painterly attitude,' we can understand what is meant if, looking at the table sideways, someone proclaims that it looks elliptical. And yet, from the very same perspective, one someone might hold that it is perfectly circular. However, notice that we could hold that the table was round from a great variety of different perspectives. In contrast, the 'painterly attitude' will give us a new result for most shifts of perspective.

On the *first* interpretation, we have *shape constancy*; we can recognize an invariant shape from variously different angles. In addition to *shape constancy* there is also distance constancy, size constancy, color constancies, motion constancies, and so on. For instance, Tyler Burge explains size constancy as "the capacity to represent an object's size as the same even as the stimulus from the object affects a smaller or larger proportion of the visual field – for example, while it moves closer to or farther away from the viewer" (Burge, 2010, p. 409). On the *second* interpretation, we have a fine-grained perspective on the table that shifts whenever the perspective shifts even a little. And yet this fine-grained perspective seems fundamentally related to the *shape constancy*.[24]

The first question was, why does our perceptual experience seem so stable and constant? The second question is how do the two interpretations, 'shape constancy' and 'painterly attitude,' relate? Let us look at some examples.

[24] For instance, the table seems to look round because it (from the painterly attitude) looks oval in a certain way, just like the table might look distant because it (from the painterly attitude) looks small in a certain way.

2.2.3 How do Perceptual Constancies Work?

A good illustrative case to explain how perceptual constancies work is the case of *texture constancy*. To understand this we first need to understand *texture density*. Texture density is simply how densely packed together the texture in any given area is. In general, if we recede from something, say a painting on a wall, so that the area of the painting seems to shrink from the 'painterly attitude,' there will be a greater density of texture within the frame, because the texture content of the painting is now crammed into a smaller area relative to our viewpoint.

Since there tends to be a constant relation between texture-density and distance, the part of a slanted textured surface that is further away will typically, from the 'painterly attitude,' look to have denser texture. Thus, the perceptual system can work out the slant from the change in texture. The perceptual system will "know" that the denser texture is further away.[25] Imagine that our perceptual system is determining the orientation of a 2D surface in 3D space (Burge, 2010, p. 357). For instance, consider the yellow text opening each Star Wars movie. It appears to us that this text is rolling on a slanted invisible surface up and away "into" the screen.[26] However, as we all know, the text on the 2D screen is actually only getting smaller and more narrow across as one looks upwards on the movie screen. However, this is just to say that the text is more densely packed together as we look upwards on the screen. We get the illusion of depth because the perceptual system "assumes" that the denser texture is further away. The perceptual system keeps the density and letters constant, and varies the spatial representation of orientation instead.[27]

Burge calls the principles that give rise to constancies *formation principles*. For instance, "The formation principles describe law-like transformations that produce perceptual representation that privileges representation as of a slanted surface with regularly distributed

[25] Both Burge and Searle occasionally talk in this way, as if the perceptual system, or any subpersonal process for that matter, is an entity that could know, act, will etc. There is no harm in it so long as we are careful. See Burge's "Vision and Intentional Content" and Searle's reply in (LePore, 1991, pp. 195-225).

[26] If the reader has not seen this, imagine looking down on a book page as someone drags the book away. As it gets further away the text gets smaller. In addition, the smaller text is further away. But we neither think of the characters overall as shrinking in size, or of the characters further away as smaller than those closer by, except when considering the perceptual content in the 'painterly attitude.'

[27] Of course, in this case our perceptual system also makes use of shape constancy and letter constancy mechanisms that contribute to the perceptual system keeping them constant in this particular way. In general the perceptual system is very good at integrating different such statistical "assumptions." See footnote 29.

textural elements over representation as of a straight-on surface with irregularly distributed textural elements" (Burge, 2010, p. 356). The system *favors* one particular interpretation over another. In this case, the particular formation principles leave the perceptual state free to make transformations that accord "with the principle that an increase in average density is roughly proportional to an increase in distance" (Burge, 2010, p. 356, fn.56). Since the perceptual system in this instance favors the regularity of textural elements, it will increase apparent distance, as in the Star Wars text example, as our eyes move vertically upwards, since the density increases.[28]

While some formation principles can be formulated in strict mathematical terms, we know that the perceptual system operates on a complex set of different cues simultaneously. Burge notes that there are "as many as twelve basic capacities for determining distance and depth" plus "combinations among these capacities and input from other senses [that] complicate distance and depth determination further" (Burge, 2010, p. 347).[29] Instead of further isolating models of specific formation principles leading to specific constancies, some researchers have been motivated to unlock more general principles. Burge devotes little attention in his philosophical papers to such speculative *general* formation principles, but his methodology would on the face of it commit him to using 'perceptual constancy' in a way that is consistent with the best explanations in perceptual psychology.[30] Contrast Burge's comment that perceptual psychology is a "serious" and "well-supported, mathematically rigorous, mature science" (Burge, 2010, pp. 87, 297), to his comment that "the science of perceptual psychology is in its early maturity" (Burge, 2010, p. 98).

[28] This is not strictly true. But it serves as illustration. The perceptual system always integrates different cues. We also see the text as the same size because the perceptual system 'knows' how the letters are proportioned relative to each other, independently of the density of the texture they make up.

[29] He also notes that "recent work suggests that the division into roughly a dozen basic capacities [to determine depth and distance] is probably artificial, and the number of types of cues that provide absolute or relative distance and depth information may be much greater" (Burge, 2010, p. 347). Note that brackets are always mine unless noted otherwise.

[30] I doubt that Burge has a very specific notion of 'perceptual constancy' in mind. An open conception might be a virtue insofar as perceptual psychology is only in its *early* maturity. See (Ganson et al., 2012) for a critical discussion.

Let us therefore race ahead and see to what extent it is possible to have a unified view of constancies. This unified view appeals to a general method that may subsume the explanation of at least a range of constancy phenomena under it as special cases.

2.2.4 Attempt at a Unified View of Constancies

It has become clear that the problem of explaining how perceptual constancies work is not just a problem for perceptual psychologists. This is because the perceptual problem of how we see the world as constant, despite fluctuating sensory stimuli, turns out to be related to several other problems: how our perceptual system is able to compress information, how it is able to recognize patterns, and how it is able to predict what will happen next.[31] Already with David Marr, the early pioneer in cognitive approaches to vision (D. Marr, 1983), the problem was shifted from the study of purely biological mechanisms of vision, to be about how best to model the process of vision computationally. Since then, the problem has attracted everything from computer engineers, information-theorists, artificial intelligence researchers and theoretical neurologists, and they all have something to contribute.[32]

While I will not go into detail on these developments, there is simple and rewarding insight on the nature of perceptual constancies to be stolen here. The rewarding insight is perhaps most easily made sensible by Jeff Hawkins theory in his (Hawkins, 2004) and (George & Hawkins, 2009). His theory, like that of cognitive psychology from the time of David Marr onwards, is partly meant to explain how we get from sensory input to a *3D object-centered description*s, that is, object descriptions "that allow the object to be recognized from any

[31] It is also related to the capacity for orchestrated action, or motor-action. See footnote 32 for a list of examples.

[32] Jürgen Schmidhuber (information-theorist) and Churchland (neurophilosopher) believe that there is a ReNNaissance for neural nets, that is, the "**N**eural **N**etwork Renaissance." They explore Neural Networks that are able to distil various constancies automatically from natural sources of data. Schmidhuber and his team have won several artificial intelligence pattern recognition competitions. For a collection of papers, see (Schmidhuber, 2012). Paul Churchland has been championing neural nets as a way to understand mental content for a long time, most recently in his (P. M. Churchland, 2012). Rodolfo Llinas, a giant in neuroscience, claims in his (Llinás, 2002, p. 21) that prediction is the fundamental function of the whole brain. He also claims (Llinás, 2002, p. 133), similarly to Churchland (P. M. Churchland, 2012, pp. 139-143) that motor-action can be understood as reverse perception. The idea that prediction is the function of the brain has recently been formalized by brain-imaging pioneer Karl Friston in (Friston, 2010). These are just a few examples of the size of the terrain.

angle (i.e. independent of the viewpoint of the observer)" (Edgar, 2005, p. 91).[33] Like Marr, Hawkins treats the *how* of perceptual constancies purely as an information-processing problem.[34] They both think of the mind with an analogy to computers and imagine that the problem of constancies has to do with finding the right algorithm that can transform the stimuli that hits the eyes into a useful representation. Like Marr, Hawkins starts by assigning to each of the possible elementary changes in sensory stimuli a classification that is regarded as a "primitive element."[35] In the case of vision (shape) we can imagine the primitive elements as dots, edges, lines and the various orientations of these; in audition we can imagine it as pitches, intervals, and volume, and so on.[36] The idea in both Marr and Hawkins theories is that these primitive elements are hierarchically processed to gradually construct a more and more time and space invariant representation of features in the environment.

Jeff Hawkins calls his idea the Hierachical Temporal-Memory (HTM) theory. Part of the motivation for a *unified* account comes from the fact that the neo-cortex, literally the new cortex or rind, the seat of all our higher perception and intelligence, has a columnar or modular organization (Mountcastle, 2003). In particular, it is uniformly hierarchically structured.[37] Given this fact, Hawkins thinks it is plausible that there is a common set of principles at work in the neo-cortex.

---

[33] Of course, the HTM theory is meant to be a unified theory of perception in general, not just how to get from sensory input to 3D object-centred descriptions in particular. In general, unifying theories tend to explain a lot more than was initially posed as the explananda, and the same is true if something like what Jeff Hawkins proposes turns out to be correct.

[34] "The detection of physical invariants ... is exactly and precisely and information-processing problem, in modern terminology ... The only way to treat it is as an information-processing problem." See (Marr, 1982, p.30) cited from (Edgar, 2005, p. 90). Also see (Poggio, 1981)

[35] For each type of change, he assigned an edge-segment, a bar, a termination, or a blob. See, (Marr, 1982, p.74) cited from (Edgar, 2005, p. 94). Also see "it is not the details of Marr's theory which have so far stood the test of time, but the approach itself," (Wade & Bruce, 2001, p. 97) and "Marr's model of vision which, in our view, marked the start of the modern era of vision research" (Wade & Bruce, 2001, p. 1).

[36] It is drastically more complicated than this. For one thing, I ignore top-down processes. And the idea that the 'top' of the hierarchy, as in (P. M. Churchland, 2012, p. 62), is usually composed of several nodes, each capturing statistical regularities in the whole perceptual field, that only jointly compose what is actually perceived, such as a given face. For an indepth study of how this might make sense of how we process and represent *music*, see Rob Snyder's *Music and Memory* (Burge, 2003b). I also ignore the sense in which these "primitive elements" are themselves a function of some limited computation or encoding.

[37] See (Hawkins, George, & Niemasik, 2009, p. 1092) and (Goldman-Rakic & Rakic, 1991). Also see, "we find a

Hierarchical models of object recognition in the cortex have been around for quite some time (Fukushima, 1988; Riesenhuber & Poggio, 1999). Inventor and futurist Ray Kurzweil makes similar points in his recent book *How to Construct a Mind* (Kurzweil, 2012), where he is fronting what he calls the Pattern Recognition Theory of Mind (PRTM). All of these theories are quite similar to (HTM).[38] I will focus on HTM here.

Jeff Hawkins attempts to account for what we know in terms of a simple, but powerful model. While there are disputes about precise details of the model and of exactly how general such a model can be, the hierarchical uniformity of the neo-cortex is a fact. It is also a fact that single neurons, or small groups of neurons, in some way manage to encode highly abstract and specific features.[39] Interestingly, Hawkins thinks it plausible that our hierarchical neo-cortex is in some way mirrors the hierarchical informational structure of the world (Hawkins, 2004, p. 125).[40] According to (Shagrir, 2010, p. 19) "a careful examination reveals that Marr's explanation appeals to similarity between the internal mapping relations and external relations between the features that are being represented." Thus, HTM might also ultimately help with explaining the representational content – accuracy conditions - of perception partly in terms of a kind of mirroring.[41]

---

remarkable degree of transcriptional uniformity compared to other brain regions, apparently reflecting the similarity in laminar architecture across the entire neocortex" (Hawrylycz et al., 2012, p. 397), and "the great majority of cortical regions being six-layered isocortex, including visual, auditory and somatosensory cortices, and many prefrontal and frontal areas" (Raizada & Grossberg, 2003, p. 105).

[38] This recently published book (December 2012) only helps confirm what that it may be that we are close to getting a unified view of what a large component of our intelligence fundamentally is, and that perception (pattern recognition) is essentially related to it.

[39] See, for instance, (Hung, Kreiman, Poggio, & DiCarlo, 2005). For an overview, see (Gross, 2002). For a critical discussion, see (Bowers, 2009). Interestingly, the discussion on this topic is rife with philosophical discussion. For instance, in a reply to Bower's article, the authors note, "To continue the last example, since the specific piece of toast was only ever encountered once, Bowers's recognition of it (as a piece of toast) cannot be attributed to previous experience with that particular object but must depend on previous experiences with other objects—other pieces of toast he presumably has encountered at other times and places" (Plaut & McClelland, 2010, p. 287).

[40] Kurzweil also make similar points (Kurzweil, 2012, pp. 5-6, 24, 35-61). I think this it an interesting claim. In particular, there is an interesting empirical question that perhaps could be solved by a joint effort by information theorists and evolutionary psychologist to uncover whether or not the neo-cortex evolved at least partially for this reason or not.

[41] Burge is dismissive of proposals that take mirroring in this sense as a – ceteris paribus- sufficient condition on

Understanding HTM

To understand HTM, we should imagine a multi-layered hierarchy of nodes, for instance neurons. In this hierarchy we imagine each neuron performing roughly the same type of function. What each of the neurons does is to *detect a given type of sequence*. Think of the sequence as a sequence of electrical signals coming from one or more neurons on the level below in the hierarchy. As soon as a neuron detects the given sequence, we say it becomes active. Just as we can detect a song and remember what it is called, the neurons detect particular sequences and become active in a certain way. Think of a neuron becoming active in a certain way as a classification or a name given to the sequence it becomes active in response to.

Now, imagine a neuron, call it N1, at level one in the hierarchy, becoming active in response to a given sequence. For instance, N1 may become active in a certain way when it detects a sequence that we observers perhaps would call a "line segment." Another neuron at level one, call it N1*, may become active in response to what we call "an edge." Now imagine a neuron at level two, call it N2, taking inputs from N1 and N1*. Imagine that N2 only becomes active in response to a certain *sequence* of inputs from N1 and N1*. In effect, N2 only becomes active in response to a sequence of the sequences that N1 and N* classify. For instance, we can imagine that N2 only responds to a sequence from N1 and N1* that corresponds to the pattern an observer might call "L." This "L" is thus classified by N2 in response to a certain spatial-temporal succession of N1 and N1* becoming active in a certain way.[42]

Whenever N2 detects that this sequence is in progress, it becomes active in a certain way. To the neuron even higher up, however, call it N3, "L" is just a simple unit, just like "an edge" is to N2. Accordingly, N3 will classify sequences of the sequences that N2 and other neurons on

---

representation. However, Burge sometimes suggest that it is a useful notion such as when he says that "explanation tends to operate on categorizational (perceptual-attributive) capacities whose structure is that of various magnitudes. The most prominent magnitude structures in perceptual representational content map onto structures of spatial magnitudes in nature"(Burge, 2010, p. 104). The reader might want to keep this conception in mind.

[42] This example is strikingly similar to the way Kurzweil explains his theory in (Kurzweil, 2012, pp. 41-48). The common root is Hawkins's book (Hawkins, 2004), although Kurzweil independently discovered the basic principles in his work on speech recognition and synthesis in using Hidden Markov Models (HMM). See (Kurzweil, 2012, pp. 131-146) for an explanation of HMMs, Kurzweil's early developments and his work on speech recognition and synthesis.

level two can classify. For instance, if another neuron at level two classifies "E," then a sequence detected by N3 might be "EEL." If there are an additional two neurons at level two that classify "A" and "P", we can get a single neuron at level three to respond to the sequence "APPLE." While each neuron does the same simple thing, each rise in the hierarchy (and this method should be current across most of the neo-cortex) implies that the neurons in effect classifies sequences of the sequences of the level below.

The upshot is that, because this process of classification treats spatial and temporal patterns equally as subject of abstraction and compression[43], and because the neurons activate only approximately to the pattern they respond to, the neurons can activate, i.e. classify, even with significant aberrations and omissions in the underlying details of the full sequences they detect.[44] We can for instance, forgive my spelling, raed wrods taht are srcmabld up. Here the underlying sequences in phonology and morphology is approximately correct, the letters in the words are plausibly instances of the word we interpret it to be, and the overall meaning and grammar of the sentence implies a certain sequence of words.

This latter point is also illustrated by the fact that, whenever we read a sentence, we predict, at the start of each word, whatever letters will be at the end of it, and at the start of a given sentence, whatever word will be at the end of it. Since all neurons becomes active as soon as they get stimuli that corresponds to their sequence, the neurons higher up in the hierarchy, that become active at the beginning of a large sequence, can constantly "expect" what will happen next, as when we know how a song goes.

Similarly, the argument from this model goes, when we see a table as round despite looking elliptical to the 'painterly attitude,' it means that the perceptual system subsumed our perspective under an object-centered 3D model of the table.

Indeed, on this theory, it makes as much sense to speak of a perceptual surprise that a song does not go as expected, as it makes sense to speak of a surprise if it turns out the that table is just an illusory 2D surface at so-and-so distance. In both cases, we can objectively make sense

[43] "Abstract" and "compression" because what a node higher up in the hierarchy treats as a simple, is actually an elaborate spatio-temporal sequence. So it both treats them in the abstract, and compresses information.

[44] I simplify. For the details of how this might work, that they respond only 'roughly' to the spatio-temporal pattern, see the distinction between 'simple' and 'complex' cells in theories like (Fukushima, 1980; Riesenhuber & Poggio, 1999; Serre & Riesenhuber, 2004).

of the idea that the perceptual system "expected" something else. In both cases, we (our perceptual system) had a complex spatial-temporal sequence "in mind."

Some Worries

One might worry that this account differs from Burge's account of perceptual constancies. In Burge's Origins of Objectivity, it is not perfectly clear what Burge means by a perceptual constancy. He introduces the notion by example and by appeal to the perceptual psychological community. As I see it, Burge is committed to the relevant concept that *successful* perceptual psychological science uses in their explanations of representation.

Much work needs to be done on this topic. However, one might distinguish two kinds inherent in Burge's usage of the term 'perceptual constancy.' On the one hand there are rigid transformations according to mathematical laws, such as the translation-laws that govern depth perception by computing, geometrically, the distance from the space between the eyes and the angle that they form in looking at something.[45] On the other hand, we have a memory-governed hierarchical, pattern-recognizing, "constructivist" notion of perceptual constancies along the lines of HTM. It seems to me that Burge does not clearly distinguish between them. While the former might give us better reason to suppose that the perceptual constancies are inherently objective, it is the second version that can more plausibly account for the idea that it is the *subject as a whole* that perceives.

A recent article critical of Burge's use of the notion of perceptual constancies attribute to him (and David Marr) the second, "constructivist" version explicitly, and hold that Burge is simply wrong to use examples like depth-perception by convergence to explain what a perceptual constancy is (Ganson et al., 2012, pp. 8-10). I don't want to take a stand on this issue here.

I chose to focus on the HTM version of perceptual constancies because it is coherent with fundamental, unifying principles of general brain function, such as predictive coding.[46] In

[45] This seems to me largely uncharted territory. However, Farid Masrour seems to make fruitful use of such a mathematical notion of perceptual constancies to argue for a kind of phenomenal objectivity based on the "dynamical unity" that such constancies display (Masrour, 2013).

[46] This is, if possible, an even more general approach than HTM. A review is given by (Bastos et al., 2012) where the heading 'predictive coding' is understood as a form of Bayesian inference implemented in hierarchical

addition, as I am motivated by the fact that it is the *subject* that perceives, I focus on the kind of mechanisms that I can conceive not only as contributing to, but as potentially entering into the perspective of the subject as a whole. For instance, a HTM-like view of perceptual constancies might help make some sense of the distinction between a table both looking round, and yet from the 'painterly attitude' looking elliptical.[47] It is exactly in the hierarchical structure of the theory that two different looks of the table might be present simultaneously. If we simply focus on the lowest level of the hierarchy, such as perhaps computation of distance, and segmentations of lines and so on, it is clear that the various ways in which the perceptual system represents a given distance or line, do not even potentially enter into the perspective of the subject as a whole, as they are, by stipulation, lowest in the hierarchy.

To see the difference more concretely, consider what Erwin Schrödinger notes in his book *What Is Life*: "if waves of 760μμ, which by themselves produce the sensation of red, are mixed in a definite proportion with waves of 535μμ, which by themselves produce the sensation of green, this mixture produces a yellow that is indistinguishable from the one produced by 590μμ" (Schrödinger, 1992, p. 154). The unit (μμ) is nanometers. It is clear that this difference does not potentially enter into the perspective of the subject. However, consider that we often see the same shade of color (say, a yellow flower) as being the same even under varying light conditions (inside and outside). In this case, we can apply a dichotomy analogous to that between 'painterly attitude' and constancy. The varying light enters into the very perspective of the subject as a whole, in a way that the varying light in Schrödingers case did not. In particular, the difference between (760nm and 535nm) and (535nm) cannot enter into the perspective of the whole subject.[48] I will discuss the notion of 'perspective' in this sense and 'whole subject' in more detail in 2.3.6 and in chapter 3.

---

cortical areas.

[47] In this sense we also get theoretical continuity with full-blown visual object recognition. See (DiCarlo, Zoccolan, & Rust, 2012) for a (very good) review of primate level 'core object recognition,' the ability to rapidly recognize objects despite substantial appearance variation.

[48] One need not stop there. A subject typical computer screen only consists of – roughly – photons with three distinct wavelengths. In this sense, almost all of the colors on our screen are illusions, if we regard most of the colors we see as purporting to be about the full range of what is called 'the visible spectrum' of light.

2.3 Perceptual Representation

2.3.1 Introduction

According to Burge, "Perception is representational in that its nature is both to purport to be about something and to represent it as being a certain way" (Burge, 2005, p. 5). The intentional aboutness aspect and the representational aspect stand or fall together on Burge's account. In addition, Burge holds that "Perception requires perceptual constancies" (Burge, 2010, p. 399). Perceptual constancies are meant to partly explain *how* perception is representational. To understand this we should understand what Burge means when he says that perceptual content is representational and intentional (I use representational for short).

To properly see how perceptual constancies are related to perceptual representation, I will explain what Burge means by representation. Then I will explain what Burge means by perceptual representation, and why constancies are central to understanding it. Finally, I will look at how and in what sense the perceptual constancies underlie the perceptual representational content. These explanations are not exhaustive. Our main goal is to understand Burge's views on the relation between perceptual constancies, perceptual content and the perspective of the subject's as a whole.

2.3.2 Burge on Representation

To explain representation, Burge appeals to two concepts. One he calls reference, and the other he calls indication.

Reference is both a relation to "an entity in a subjectmatter" and a "function (or exercise of a function) of a state, event or activity to establish a reference relation" (Burge, 2010, p. 31). The paradigmatic example Burge uses to explain *reference* is a case of singular reference: If someone looks at a chair and thinks *that chair*, then that someone will on that occasion be referring to the particular chair. A common linguistic example of *reference* in this sense is the relation established by a proper name and what it names. Another common example is the relation between a demonstrative, such as *this* and *that*, and what is demonstrated.

Indication strays quite a bit further from common usage.[49] Consider the example, 'That apple is red.' Here 'that apple' refers to a certain apple. Instead of saying that 'is red' refers to the property *redness* Burge says that it *indicates redness*. The key difference between 'is red' and 'that apple' is that 'is red' functions attributively; it attributes redness to something – in this case the particular apple.[50] In general, we can think of predicates and concepts as essentially attributive. *Indication* is "reference with the further constitutive representational function of attribution (or functional application)" (Burge, 2010, p. 31). To say that something is a *constitutive* condition, such as the representational function of attribution is to *indication* in this case, is not quite to say that they are part of, or elements in, the nature of the thing, or that they are necessary properties of the thing. Rather, "Constitutive conditions for a nature help ground explanations of what it is to be that nature" (Burge, 2010, p. 533). To say that something is a constitutive condition is to say that this something enters into the best explanation of what the nature of the thing to be explained fundamentally is.

Notice that neither reference in a case of singular reference or indication in a case of attribution is necessarily a relation to an entity/property. This is because reference and indication are both defined either as *a relation* to an entity/property **or** *as an exercise/application of a function* to establish a reference. They can *function* to establish a reference on an occasion of use without necessarily succeeding. Burge offers the following analogy: "Representation is rather like shooting. Some shots do not hit anything, but they remain shootings" (Burge, 2010, p. 45).

Thus, while there needs to be something that functions to establish a reference relation, (and this function needs ultimately to be associated with a function that successfully does establish a reference), an occasion of use does not necessarily need to be successful. For instance, an occasion of use need not successfully establish a reference.

2.3.3 Burge on Perceptual Representation

Let us now relate these terms to what Burge calls *perceptual* representation, and begin to transition into how this relations to perceptual constancies. Consider a perception of a

[49] For instance, Robert Cummins in his (R. C. Cummins & Poirier, 2004) uses 'indication' quite differently, to contrast sharply with representation (a usage more in line with what Burge means by reference here.)

[50] One might notice that 'that apple' also picks out the particular apple with the concept 'apple.'

particular apple as red. According to Burge, such perceptual representational content consists of attribution and a singular element of reference. The particular apple is referred to by a singular element of reference, while redness is attributed to it. In addition, the apple is referred to by means of the attribute 'apple.' We attribute the concept *apple* to the particular apple referred to. On the one hand, attributes "can be instantiated by various particulars" (Burge, 2010, p. 24 , fn.11). On the other hand, singular reference is always with respect to a given context and application.

Perceptual constancies can primarily be understood as an explanation of how perceptual attributives are possible. In attribution, as we have seen, indication is important. However, another important aspect of perceptual attribution is their function to group or categorize particulars.[51] For instance, in harmony with our HTM account of constancies earlier, Burge says that "the attributive elements categorize at various levels of abstraction" (Burge, 2010, p. 104).[52] From what we have seen of perceptual constancies, this seems to be the most relevant function that perceptual constancies can help us understand.

In a daunting passage, Burge states that the "Perceptual representational contents 'semantically' determine their representata, if any. In particular, perceptual attributives semantically determine, or specify, the attributes that they attribute" (Burge, 2010, p. 76). This somewhat cryptic comment should begin to make sense once we notice that there is a "distinction between the perceptual attributive (a certain type of representational content), what it indicates and attributes (a kind, property or relation – an attribute), and what it attributes something to (a particular)" (Burge, 2009a, p. 7, fn.5). Thus, "the attributes that they attribute" means the kind, property or relation that is indicated and attributed (to a particular); but now how can perceptual representational content 'semantically' determine, or specify these? I will make a distinction between how the perceptual content can *determine*,

[51] How to get from "grouping" and "categorization" to indication is a complex issue. Most immediately, the objectification present in perceptual constancies is some help. In addition, there are further anti-individualist considerations that make this transition possible.

[52] Burge often uses 'grouping' often when speaking of perceptual attribution; "Grouping, or perceptual attribution, is a constitutive aspect of perception" (Burge, 2010, p. 305, fn.20). "Perceptual grouping discriminates a kind from other kinds in the environmental context" (Burge, 2010, p. 198). Burge sometimes also uses "group or characterize" is such contexts. See (Burge, 2010, pp. 364, 380-381, 453, fn.35). While the exegetical work on Burge too comprehensive to get into, notice that a theory like HTM helps makes sense of *why* 'discrimination,' 'grouping' and 'categorization' might be aspects of the same phenomena.

and how it can *specify* the attributes that they attribute. Before we get to that, let me say some more about how the perceptual constancies are related to perceptual representation.

2.3.4 Burge on Perceptual Representation and Perceptual Constancies

The appeal to perceptual constancies is *part* of Burge's solution to the problem of how perceptual attributives 'semantically' can determine the "attributes that they attribute." Burge gives the following example: "The aspect of a perception that groups something as a body, or as cyclindrical [sic], indicates the kind *body* or the shape *cylindricality*, if it indicates anything – and does so in every context of use and regardless of what possible situation is under consideration" (Burge, 2010, p. 76). We can think of a perceptual constancy as grouping something as say, *cylindricality*; thus understood, the perceptual constancy is understood to be an attributive that has the constant function of indicating *cylindicality* (as actually succeeding to indicate and/or just functioning to indicate).

Burge uses perceptual constancies to help explain how perceptual attributives are possible.[53] This is because the explanation of perceptual constancies can be understood partly as an explanation of how it is possible that attributes can be of the same attribute in every context of use and with regard to any possible situation; in every context of use and with regard to any possible situation a perceptual constancy can attribute something *cylindrical* as indicating *cylindricality*. Also in this sense the HTM model outlined earlier fits quite well. The hierarchy of ever more abstract meanings arising from primitives fits the idea that perception groups something as *cylindrical* or as *body* and does this in a way that might, at least conceivably, begin to have some 'semantic' import.[54]

[53] Recall that perceptual attributives are a certain kind of perceptual content. Thus, perceptual constancies at least explain how a certain kind of perceptual content is possible.

[54] It is not clear how or in what sense it fits with the law-like mathematical conception of constancies as exploiting regularities in the environment, though, of course, the most primitive elements in the hierarchy might be determined by input in such a fashion. A related question is whether one should adopt an Ideal Observer Analysis, e.g. (Geisler, 2003), where one ascertains the statistically optimal way of perceiving certain things in the environment; also this seems to be a step away from semantic 'grouping' in perception.

2.3.5 Perceptual Representational Content – Determination and Specification

Burge tacitly uses a principle that I call Conditional Principle of Determination (CPD): *if* something is insufficient to determine the *representata* of perceptual states, it is insufficient to determine their representational content. For instance, Burge uses this implicit principle when he says that "*if* a phenomenal feature is insufficient to determine the *representata* of perceptual states, it is insufficient to determine their representational content" (Burge, 2010, p. 77). Thus, to find a candidate for perceptual representational content, we need something that is sufficient to determine the *representata* of the content. We should also keep in mind that if we want perceptual constancies to explain perceptual (representational) content, we must understand to what extent constancies are sufficient to determine a perceptual representational content. To make sense of this we should distinguish between two things that I will call *specification* and *determination* respectively. By *specification* I mean that perceptual constancies might be said to rigidly designate a kind, property or relation, regardless of context. By *determination* I mean the extent to which perceptual constancies are a capacity to discriminate particulars such that they respond only to a certain range of things. As perceptual constancies are partly capacities to discriminate (and group) particulars, they exclude a certain range of alternatives and thus determine what they are about.

Plausibly, perceptual constancies by themselves can determine a single attribute under varying conditions. For instance, in each case I look at something cylindrical, the perceptual constancy may, to a certain extent, determine cylindricality. However, it is clear that perceptual constancies may be exposed to illusions. Thus, the grouping or discriminative capacity of a perceptual constancy can only go so far.[55]

If we want to *specify,* it is clear that no amount of determination, as capacity to group or capacity for discrimination, will do. These can only exclude certain possibilities, but never specify the environmentally relevant ones. For instance, one might worry whether the

[55] Note, however, that even in attributing something cylindrical as being cylindrical, in a way that it indicates cylindricality in a way that specifies the representata, this does not determine the perceptual representational content to the finest level of explanatory grain. As Burge puts it, "There are many representational contents for any given kind, property, relation, or particular that is represented in these 'as' expressions. The representational content is *always* more fine-grained than the as locution suggests.[my emphasis]" (Burge, 2010, p. 40).

perceptual constancy in a frog's perceptual system is specifying flies, sensory information hitting the senses, physically indistinguishable undetached parts of flies or temporal slices of flies; more realistically, one might wonder whether a frog's perceptual constancy determines a black, flying, buzzing dot, or whether it determines something else. In light of this, we should be asking two questions; first, how fine-grained do we want the *determination* to be? Second, how fine-grained do we want the *specification* to be?

First, note that Burge defends the content view partly by appealing to perceptual constancies. For this reason, we might expect that Burge sets the sufficient condition of how fine-grained we want perceptual representation to be around what perceptual constancies are able, in principle, to determine, as it were, on their own. In one sense, Burge does this by limiting our attributes to what our perceptual constancies are able to determine and represent on their own.

Burge also strengthens the fineness of grain by individuating perceptual constancies in terms of an externalist framework he calls anti-individualism. With this, Burge can hold that perceptual constancies *specify* what they represent in a fine-grained way. Here we simply note these two aspects of content individuation; on the one hand there is a certain specification at work, on the other hand a kind of determination.

In the following we first look at how Burge limits the sense in which perceptual constancies can determine content. Second, we look at the added specificity that his anti-individualist framework brings.

Limited Determination

On Burge's account, "perceptual constancies are capacities for objectification" (Burge, 2010, p. 399), and objectification "hinges on distinguishing and contrasting, in the operations of the system, what concerns the individual's receptors and what concerns a receptor-independent reality and doing so in an attribute-specific way" (Burge, 2010, p. 398). Burge illustrates how perceptual representational content can be attribute-specific by appealing to constancies. By means of constancies, the perceptual system processes the immediate effects of sensory stimulations "to provide a perceptual model of the world" (Burge, 2010, p. 398). In this sense perceptual constancies are able to *determine* with a certain fineness of grain say a particular elliptical object as the surface of a table of a certain sort. We might see a particular table as being at a certain approximate distance. Here we singularly refer to the particular table or

distance, by means of the attribute that determines the table as being of a certain sort, or the distance as being of a certain approximate distance.[56]

Burge notes that the idea of a perfect determination is an idealization: "Perceptual representational contents commonly represent attributes approximately," that is, "within some range. Thus states represent a distance plus or minus" (Burge, 2010, p. 380). I think this is plausible. For our purposes, we can simply see this as a natural consequence of how perceptual constancies work. Perceptual constancies respond to general, abstract patterns, not specific objects in the actual world.[57] Thus, to get at specific objects in the actual world we need a further condition.

However, notice that the limitation we get on determination will have practical consequences in the perspective of the animal. It influences action. It is not just that we are fallible in our attributions; it is that our attributions have a certain resolution. In particular, our perceptual perspective will operate with certain approximations. The added specificity we will look at next can, in some sense, make up for this, but the added specificity cannot make up for the fact that we represent some distance plus or minus in the sense that follows from reflection on how our perceptual system works.

Added specificity

Burge's discussion at many points centrally involve his anti-individualistic framework; anti-individualism, for empirically based psychological states that we are discussing here, is the idea that "states are what they are partly by virtue of non-psychological, causal relations between individuals and a wider environment" (Burge, 2010, p. 11). Thus, Burge makes use of these wider relations between the individual and the environment to help type (and constitutively ground) psychological states in general, and perceptual representational states and the specificity of representational contents in particular.

Burge is concerned with what, say, the perceptual attribute 'table' indicates. Burge wants 'table' just to indicate *table*, and not also, for instance, a certain 'table'-pattern of sensory

[56] The singular element in perception is fleshed out by Burge in terms of indexing, an indexical-demonstrative like tracking relation that relates to a particular object during a given perceptual event.

[57] This is not meant as presupposing what the nature of the perceptual constancies is; as we will see, it might be in the very nature of a perceptual constancy to respond to specific representata in the environment.

stimulations, a hallucination as of a 'table,' or an evil demon projecting a 'table' into the mind. To specify the indication of the attribution at this fine level of grain, Burge has to appeal to the individual's interaction with an external environment in which the attribution was formed.[58] For instance, when I see a table, I assume that it is round and solid, even if it is possible that the same perceptual experience of a table, or pattern of stimuli reaching me, is consistent with there being nothing solid and nothing round there. In a similar vein, the perceptual system is of a round and solid table, as opposed to all these other possibilities, because the perceptual system was formed, historically, to respond to, in response to, things at least associated with such things as tables,[59] that existed in the environment natural to the given individual having the perceptual system, or the ancestors of this individual.[60]

Once these anti-individualistic conditions are in place however, the perceptual attribution will be *specific* in a way that lets Burge say that an attribution as of a table, on a singular occasion of use, will *fail* if the table merely seems to be a table, but is in fact something else. The perceptual content, however, like a shooting, is still the same in cases of failure and success. It fails to be of a *table*, if there is no *table* there, just because the content is specifically of a *table*, and not what merely seems like a table.

What all of this means is that the full individuation of perceptual content that Burge appeals to is partly specified by the environmental interaction an individual has had in forming the representational perceptual state (Burge, 2010, p. 68), and partly determined by the perceptual

[58] Or appeal to the individual's ancestors' interaction with the external environment in which an associated attribution was formed. Anti-individualism is meant to capture, among other things, broadly causal, evolutionary, environmental, formative conditions that can enter into good explanations of perceptual states.

[59] The exact nature of this *association* is not clear. Burge says, (Burge, 2003a, p. 307) "The notion of association here is a technical one. A development of the notion, and an exact statement and defense of the thesis that I am broaching, will have to be postponed to another occasion." The most radical way we can interpret this *association* according to Burge that "A perceptual attributive may partly depend for its being the type of attributive that it is on employment in a perceptual system in the system's evolutionary history, before the individual was born. An individual frog might have been given only illusory, non-veridical perceptions as of moving bodies of such and such a size. The frog can have such illusions with such representational content because its perceptual system had evolved from ancestors in which relevant veridical perceptions occurred" (Burge, 2010, p. 69).

[60] This is a short paraphrase, for our purposes, of Burge's very detailed Anti-Individualistic framework. The extent to which something can be for an individual, and still be part of an anti-individualist program, is highly interesting and I think worth pursuing, but outside the scope of this essay.

constancies as grouping or discriminating a certain range of things. Burge concludes such considerations (and in this passage the example is visual perception) by the following: "Thus, although seeing is in a sense a natural kind, it is a hybrid kind" (Burge, 2010, pp. 389-390).[61]

While we assume with Burge here that perceptual constancies are individuated anti-individualistically, this is by no means a trivial point.[62] Along with the idea that the very nature of perception should be understood partly in terms of representational content, the assumptions are beginning to stack up. However, I am motivated to understand how it is possible, in a scientifically informed way, to see how Perceptual Duality can be true. It certainly seems possible that Burge's account of perceptual constancies as individuated anti-individualistically is right, and that it is compatible with and informed by scientific considerations. The same is true of Burge's account of the nature of perception as understood partly in terms of perceptual content, and perceptual content as understood partly in terms of perceptual constancies.

Conclusion

Burge recognizes at the same time a limit to how well perceptual content is determined by perceptual constancies alone, and the added specificity, as given by anti-individualistic considerations. Naturally, there is a limit to the perceptual constancies because they cannot, as it were by themselves, recognize the difference between a spatio-temporal pattern matching what they respond to, and the real object that they have been conditioned on to respond in that particular way. However, when Burge takes the nature of perceptual constancies to be individuated anti-individualistically, they can be said to specify what they have been formed to respond to, or in response to. Thus, Burge has given us this hybrid account of perceptual representation partly in terms of a hybrid account of perceptual constancies.

[61] One might object to Burge's use of the word 'determination,' insofar as it is ambiguous between specification of the one hybrid half, e.g. specification in anti-individualist terms, and sometimes means determination, i.e. determination of what the perceptual constancies respond to, of the other hybrid half.

[62] There has been decades of controversy, at least since (Burge, 1986) over this issue. See for instance (Francescotti, 1991) , (Shagrir, 2001), (Kroustallis, 2006) and (Shapiro, 1993) on the issue of whether perceptual psychology is individuated with respect to the environment or not.

So far we have mainly been preoccupied with relating the perceptual content to the perceptual constancies, and what they are about. Next, we look at how Burge's account might relate in the other direction, to the subject as a whole.

2.3.6 Perception for a Subject?

We have not yet explained how the perceptual representational capacities, as partly explained in terms of perceptual constancies, can be *for* the subject. Typically, we think of subjects as perceiving through phenomenal experience. In an epistemic sense, we have access to the representational content through our phenomenal experience. On the face of it, when we recall that perceptual representation is for a perspective of a whole subject, it seems plausible that phenomenal experience is a constitutive condition on perceptual content.[63] However, as we will see, it might not be. And, as we will see in the next chapter, there are other things than phenomenal experience that might explain the idea, or related ones, that perception is constitutively for a subject.[64]

Unlike Burge, I do think phenomenal experience is crucial to explaining how perception is for a whole subject. To give some taste of why, consider that there is emerging consensus that consciousness occurs when the brain is processing information in a coordinated, global way (G. M. Edelman, Gally, & Baars, 2011). In contrast, stimuli that are registered unconsciously typically only reach limited regions of the brain, and more easily understood as pertaining to the subsystem of the individual (Andrade, 2012, p. 602).[65] As it turns out, Burge is more impressed by reflections on unconscious perception than on the link between the science of consciousness and whole individual agency. As Burge is concerned with establishing the minimum conditions on objective perception, cases of unconscious perception is yet another scientific reason to think that philosophers are exaggerated the requirements for objective

[63] Some philosophers might disagree with giving prior plausibility to the idea that to phenomenal experiences is constitutively associated with the perspective of a whole subject. I think most would not.

[64] It is not a solution per se to say that conscious experience is what relates the content to the subject, but it is one immediate choice we must take in explaining *how* the content relates to the subject. If we focus on consciousness, we must explain *how* it relates the content to the subject as a whole. If we want to find some other explanation, we must explain *how* that other explanation relates the content to the subject as a whole.

[65] I will come back to these points in chapter 3. Particularly section 3.2.2

perception. While I agree with Burge on most of the other requirements that he strives to get rid of, I think this is to go one step too far.[66]

Note that while I have spoken in terms of a *subject*, to contrast naturally with an *object*, Burge tends to use the word 'individual' in his discussions of minimal conditions on objective perceptual representation. Burge reserves the joint term 'individual subject' for individuals that are conscious. In his own words, "Being an individual subject requires phenomenal consciousness" (Burge, 2007, p. 395). Burge holds that phenomenal consciousness is required for being a certain type of individual, namely an individual subject; he also uses phenomenal consciousness, in *such* individuals, to demarcate the difference between what is attributable to the individual as a whole, and what is attributed merely to subsystems of the individual subject. This is important insofar as we want the perceptual representational content to be fundamentally attributed a whole individual. Later I want to use phenomenal consciousness to explain the difference between representations that are attributed to an individual *as a whole*, in contrast to perceptual representations that are attributed to mere subsystems of the individual. Naturally, some of the dispute will center on whether the individual's that seem to have perception are individual subjects or not, i.e. whether they are phenomenally conscious or not. Thus, Burge understands why someone might make the move to invoke phenomenal consciousness in an explanation of how perceptual content can be for a whole individual.

As I have insisted, perceptual representational content, according to Burge, "is a perspectival way of representing at the finest explanatory grain" (Burge, 2010, p. 38). One might worry that Burge uses the notion of 'perspective' in different ways. In the most abstract sense "all representation is necessarily from some perspective or standpoint" (Burge, 2010, p. 51). However, In a passage in (Burge, 2009a, p. 4) it seems quite clear that when he speaks of perspectival ways of representing *in perception*, something more concrete is implied. For instance he says that in these perceptual cases, the notion is "concrete, commonly spatial-directional, sometimes phenomenological" (Burge, 2009a, p. 4). Burge further holds that the perceptual constancies, in helping to make sense of perceptual representational content at the finest explanatory grain, "give empirical point to a distinction between perspective and

[66] In doing this I do not intellectualize conditions for perceptual representation. I think that even humbler animals, such as bees, have phenomenal experience. However, I do think consciousness has an important function in relation to objective representation that we are only beginning to understand. I stress its importance in making sense of attribution to the *whole* individual. I think it is important in other ways too.

subject matter" (Burge, 2010, p. 411) even in the concrete, commonly spatial-directional, sometimes phenomenological sense.[67]

There are two points. First, we want to understand *how* perceptual representational content is for the individual as a whole. A priori we assume it must be. Just like we a priori assume that perception has representational content, and want to find in a scientifically informed way *how* perception can have representational content, we now want to find in a scientifically informed way *how* this representational content can be for the individual as a whole. Second, we want to understand the representational content down to the finest level of explanatory grain. A priori we assume that it must help explain the concrete perspective of the individual as a whole. We ultimately want to explain, in a scientifically informed way, *how* perceptual constancies can help explain this concrete perspective of the individual as a whole.

Take the case of color constancies. Burge says that "… the modes of presentation involved in perceptual constancies are different. The different ways in which a given color shade are perceived as the same shade are usually available to the individual's phenomenal consciousness" (Burge, 2010, p. 412). This is analogous to the case mentioned above in section 2.2.4. It is clear that the different ways in which the given shade is perceived is differences that, as it were, show up in the perspective of the subject as a whole.

This example involving phenomenal consciousness is of course to highlight a certain fineness of grain in a perspective for the whole animal. However, with such illustrations, it becomes difficult to see *how*, for instance, the HTM model we have given of perceptual constancies, or any other model, can help explain the perspective of the subject as this fine level of grain. It also becomes difficult to see exactly *how* a model of perceptual constancies such as HTM can help give empirical point to the idea that content types the perspective of the individual as a whole. I will come back to these concerns in chapter 3.

[67] As I noted previously, Burge uses 'perceptual constancy' in a very general way, such that he can claim the benefits both of the hierarchical structure that gives empirical point to the distinction between 'painterly attitude' and for instance shape constancy, *and* the law-like structure that gives empirical point to the distinction between different subsystem transformations leading to the same invariant, say, distance. In the latter case, only the *result* enters into the concrete perspective of the whole individual and shape action and so on. In the former, *both* the 'painterly attitude' and the constancy can enter into the concrete perspective of the whole individual; e.g. I see something square also as rectangular (Burge, 2010, p. 381).

In addition to holding that phenomenal experience is not a sufficient condition on perception, Burge also think that phenomenal experience is not necessary for perceptual representational content. This could be understood quite trivially. In light of the fact that Burge is out to explain one of the most striking features of mind, and is out to explain perceptual representational content as a "perspectival way of representing at the finest explanatory grain," I think it is not (Burge, 2010, p. 38). While Burge uses phenomenal consciousness to help demarcate what is attributed to an individual subject and what is attributed to the subsystem in an individual subject, it turns out that he does not think phenomenal consciousness is needed to make this distinction in more primitive *individuals*.[68] As Burge holds that perception with perceptual representational content is possible even in primitive *whole* individual, he has to flesh out both this fineness of grain and the attribution to the whole individual without appeals to phenomenal experience.[69]

In light of the fact that Burge makes this non-trivial claim, and the fact that I ultimately want to supplement his account with phenomenal consciousness as being constitutively associated with perception, let us now look at exactly how Burge characterizes his view that that phenomenal consciousness is not constitutively associated with perception; in particular, we will look at a possible tension in his views on these matters and some possible ways that he might respond. In the next chapter, I more deeply go into *why* Burge holds the view that he does, and see investigate whether there might be the reasons to supplement his view.

2.4 Burge on Phenomenal Consciousness

Burge does not think that phenomenal consciousness is a constitutive condition on perceptual representation. In spite of this, Burge thinks that "understanding consciousness and the nature of phenomenal aspects of experience is an important and difficult enterprise" (Burge, 2003d, p. 405). This genuine interest in how phenomenal experience relates to perceptual

[68] Again, the assumption is partly that these primitive individuals are not conscious.

[69] One might worry, for instance, that Burge shifts his notion of 'perspective' when he talks of these primitive individual's to mean something different. I think he cannot. Whole individual action, which is constitutive to perception on Burge's account, requires very fine content individuation to make sense of specific actions. Thus, I agree with Burge that primitive individual's like bees have perception with fine grained perspectival representational content, I only differ in saying that this must be explained partly by appealing to phenomenal consciousness, and phenomenal experience (conscious perceptual experiences) in particular.

representation is apparent in his *replies* to Ned Block and Brian Loar, which I will be focusing on in this section. While I focus on Burge's replies, the reader might want to keep in mind that most of them are made in the context of replying to Block and Loar.

Ned Block is famous for his 'inverted earth' thought experiment, where we invert the phenomenal characters of color experiences while keeping the functional relation to the environment the same (Block, 1990a). Brian Loar is one of the explicit promoters of 'phenomenal intentionality' (Loar, 2003, p. 129), a view holding that phenomenal consciousness is the key to fundamentally understanding intentionality - how something can be *about* something else. Of Ned Block Burge hopes "that investigating the relations between Block's projects and mine will be fruitful" (Burge, 2003d, p. 405). Of Brian Loar Burge thinks "that Loar is on to something that needs better understanding" (Burge, 2003c, p. 448).

Objective and Subjective Indistinguishability, (OI) and (SI), are crucial to both Block and Loar. In one thought experiment that Loar is fond of, we imagine ourselves as brains in vats. It certainly seems possible to be phenomenally conscious even if this were the case. This is what I called (SI); roughly, different objects can be represented with the same phenomenal characters. On the other hand, we have the inverted spectrum of Block; "that things we agree are red look to you the way things we agree are green look to me (and we are functionally identical)" (Block, 1990b, p. 53). This is what I called the (OI); roughly, the objects can be represented as the same even as phenomenal characters differ. Recall that phenomenal characters are properties that type experiences by what it is like to undergo them.

Burge agrees with these intuitions, and thus (SI) and (OI). Further, Burge follows the common sense intuition that "we are directly aware of qualia" (Burge, 2003d, p. 406). That is, that we are directly aware of our own phenomenal characters. Moreover, Burge conjectures that "certain qualitative aspects of the mind depend purely on the underlying chemistry" (Burge, 2003c, p. 444). Thus, we can simply imagine this underlying chemistry being held constant, while varying the functional relations to the environment.[70] Contrast this to prominent thinkers like Dretske (Dretske, 1995b) and Tye (Byrne & Tye, 2006), who hold that qualitative aspects of mind at least partly depend on elements that are outside the head;

[70] Of course he does not literally mean that only those things that chemistry studies are relevant; the point is simply that there is some physical substratum of the brain that is sufficient.

they are externalists about consciousness, a view widely held to be counter-intuitive, even by Dretske himself (Dretske, 1995a, 1998).[71]

We want to understand perceptual representational content as being for the individual as a whole, and we want to understand the representational content down to the finest level of explanatory grain. As I find it plausible that phenomenal experience has a big role to play in the explanation of these things I want to see to what extent Burge dissociates or denies the mirroring link between phenomenal experience and the contents of perception.

One way of analyzing the relation between phenomenal experience and perceptual content is through the phenomenal characters of perceptual experience. In particular, I question Burge's account on the following two points: To what extent is sameness of phenomenal character compatible with difference in representational content? To what extent is sameness of representational content compatible with difference in phenomenal character? Recall that I typically use 'representational' instead of 'representational intentional' for short.

2.4.1 Burge: Is sameness of phenomenal character compatible with difference in perceptual representational content?

Burge holds that a given phenomenal character is dissociable from differences in perceptual representational content in two ways. In the first way, "A given phenomenal character could in principle have been associated with any of various intentional representational contents" (Burge, 2003d, p. 412). In the second way, Burge claims that a given phenomenal character can even (at least contingently) be non-representational.

In the first way, the phenomenal character is not 'enough' to determine or specify representational content because the phenomenal character could be associated with any "number of nonindexical recognitional representations and recognitional abilities with different referents" (Burge, 2003d, p. 412). This could be so by varying either the historical

[71] Recently, Bill Fish argues that while his terminology commits him to deny that hallucinations have 'phenomenal characters' he does think it *is* something it is like to have a hallucination, only that that this is "entirely parasitic on the contribution made by veridical experiences" (Fish, 2013). This lands him in a position like Dretske where he must presumably deny that it would be something it is like to be a man, physically identical to himself, arising from a swamp struck by lightning, undergoing a hallucination (Because it could not be parasitic on veridical experience). Example of the swamp man taken from Donald Davidson's (Davidson, 1987).

associations that formed the perceptual state, or by varying the functional role of the perceptual state as in the color-inversion thought experiment, or by perhaps by varying the way in which the perceptual constancies themselves associate with phenomenal characters.

In the second way, the phenomenal character can be non-representational full stop. In part, this is also the view of Ned Block, which Burge agrees with. Phenomenal characters can be 'non-representational' and 'non-intentional', where these terms are used in a similar way. This is argued by pointing out the aspects of phenomenal experience least likely to be representational. For instance, Burge argues that visual blur is not representational, but rather "a defect or a noise in the medium of representation, not an application of a perceptual category by the perceptual system" (Burge, 2003d, p. 407).[72]

In denying visual blur as an intentional feature of perceptual representation, Burge at the same time denies that it is a part of perceptual content, by appealing to the idea that it is not a perceptual category. Presumably, that it is not an attribute. The idea that it is not a perceptual attribute may, in turn, reduce to the idea that it lacks the right sort of specificity inherent in anti-individualistically individuated representational content, or that it lacks the right sort of functional role.[73]

Burge's view is the following: On the one hand we *can* see phenomenal experience (insofar as they only *seem* perceptual to the subject) as staying the same while anti-individualist relations differ.[74] And on the other hand we can find phenomenal characters, e.g. visual blur, which are not representational contents, because they do not function as a perceptual category (or perceptual attribute) in the perceptual system.

[72] I think Burge means that visual blur at least contingently is non-representational, in the sense that the same phenomenal character, in some other perceptual system, might have some representational function.

[73] Functional role broadly understood. Not just to the environment but also having the right functional role within the cognitive system of the individual – e.g. relating to perceptual constancies in the right way.

[74] Burge clearly accepts that even a brain in a vat, if it was identical to his own brain physically, would have a feeling of perceptual perspective and so on (Burge, 2003c, p. 440). He does not think of this truth as particularly fundamental, but he does accept it, unlike, for instance Bill Fish or Dretske, mentioned above in footnote 71.

2.4.2 Burge: Is sameness of perceptual representational content compatible with a difference in phenomenal character?

On the one hand, Burge thinks that "The perceptual intentional content, the perceptual modes of presentation, will commonly be in some way, at some level, different, if phenomenal character is different" (Burge, 2003d, p. 412). On the other hand Burge thinks that "Differences in the 'shades' of qualia associated with the standard sample seem to be irrelevant for typing verbal and even discriminative similarities between the individuals" (Burge, 2003d, p. 414). The difference is striking when we add that this last comment is made in the context of deciding what is and is not relevant for typing representational content.

How does this mesh with the idea that we are out to explain a "perspectival way of representing at the finest explanatory grain"? Burge explains that "at the most fine-grained level, there are as many perceptual contents as there are phenomenal discriminations ... that the individual uses, or can use, in perception" (Burge, 2003d, p. 413). On a more coarse-grained level, we lump together different phenomenal discriminations under a particular kind of representational content.[75] Note that this most fine-grained level of representation thus enters into phenomenal consciousness.

But there are now two notions of fine-grained perceptual content. The first was used above, in the section on *added specificity*, to note perceptual content that is partly specified by anti-individualist conditions. Above, Burge used this specificity to argue that perceptual constancies were sufficient to *specify* a representata. But now there seems to be a *determination* in phenomenal experience "irrelevant to typing verbal and even discriminative similarities between the individuals" (Burge, 2003d, p. 414).

One might wonder whether the kind of determination that comes with phenomenal experience is necessary for perceptual representational content. Burge answers the following; "I think that it is not sufficient for conceptual or perceptual intentionality. But it may be necessary" (Burge, 2003c, p. 448). In the context and full passage it is clear that Burge, for a moment,

[75] Also see, "Differences in the 'shades' of qualia associated with the standard sample seem to be irrelevant for typing verbal and even dis-criminative similarities between the individuals. They are irrelevant to individuating this sense of 'looks red'. Here our understanding of 'looks red' cuts through the presumed qualitative differences between the individuals' qualia and counts them as instances of the same representative type" (Burge, 2003d, p. 414).

considers the possibility that phenomenal consciousness is necessary to perceptual content on the basis of Loar's essay insisting on the importance of the fine-grained determination that enters into phenomenal experience.

However, if it phenomenal consciousness is necessary it is not so clear where the coarse-grained level of intentionality or representation enters. Burge does much to indicate that his efforts on this topic are underdeveloped and speculative. As we will see below, Burge more recently considers this passage to be written before his more mature views on these matters. However, let us focus on whether there might be some tension underlying this passage and how it might be resolved.

2.4.3 Tension

Given the passage above, Burge holds the idea that the phenomenal character *may* be necessary, but that it certainly is not sufficient for perceptual representational content. However, Burge also holds that phenomenal consciousness is not necessary for perceptual content.[76] It might seem strange that Burge can hold that phenomenal character *may* be necessary. Can he?

Burge in his (Burge, 2003c) has two disclaimers on these points: The first and perhaps most crucial disclaimer which is that there may be perceptual representational content that disregards the specifics of phenomenally conscious discrimination. The second is the admission that perceptual representational content *commonly* will be different, on at least some level, if the phenomenal experience is different.

With these, Burge can hold that for some representational contents, some specifics of phenomenal experience may be unnecessary. This can be held while simultaneously holding that phenomenal experience is necessary to representational content. There are but some phenomenal characters, in some cases, that are not necessary for representational content. But still, phenomenal experience is necessary to perceptual representation.

[76] Strictly speaking he only said above that it was 'irrelevant' (in the context of typing perceptual content). However, as we will see as we go on, Burge is very explicit about his conjecture that phenomenal consciousness is not (even) constitutive to perception, much less necessary for perceptual content in general.

Unfortunately, even with these disclaimers, the idea that phenomenal experience is necessary to perceptual representation does not go well with the following conjunction: the idea that perception is necessarily representational *and* the idea that perception can be non-conscious. Burge accepts both of these. But they imply that non-conscious perception is representational. Thus, phenomenal experience cannot be necessary to perceptual representation, since there are non-conscious perceptual representations.

To resolve this directly, we can perhaps interpret Burge as meaning that phenomenal consciousness is constitutively associated with perceptual representation. For instance, perceptual representations lacking phenomenal experience might be parasitic on successful instances of perception that do not lack phenomenal experience.[77]

Unfortunately, this reply too is one Burge cannot accept. If Burge is to accept that phenomenal consciousness is in *any way* constitutively associated with perceptual representation he cannot hold that perceptual constancies are sufficient to demarcate perception from other kinds of reactions to stimuli. This is because Burge believes that it is plausible that constancies are dissociable from phenomenal consciousness not only in abnormal cases (that might be parasitic on successful cases) but also in cases such as that of the insect, which cannot constitutively associate to any phenomenal experience at all simply because the nature of that particular insect does not constitutively associate with phenomenal consciousness at all. The insect might not be phenomenally conscious at all.

Thus, if Burge is to have phenomenal consciousness constitutively associated with perceptual representational content, he might consider *rejecting* either of the following two statements: Either, perceptual constancies are sufficient to mark the difference between non-perceptual processes and perception. Or, there here are some animals that have perceptual constancies that do not have phenomenal consciousness.

If Burge rejects the former, then he can for instance hold that only perceptual constancies in conjunction with phenomenal consciousness can mark the difference between non-perceptual processes and perceptual representation. If he rejects the latter he will have to conjecture, for instance, that the bee has phenomenal consciousness, since we know that it has perceptual constancies.

[77] I give a more examples of how one might understand "constitutively associated with" in the next section.

2.4.4 Burge Replies to the Tension

Burge's comment that phenomenal consciousness *may* be a necessary condition on perceptual representational content was very noncommittal, exploratory and in the context of trying to find something conjecturally positive in Loar's essay; it was before the fuller development of his views on this topic (Burge in private correspondence). However, given the idea that it is something conjecturally positive, what would Burge's position be if he *did* commit to it? Which of the two statements above would he reject? Burge would reject the idea that phenomenal consciousness is dissociable from perceptual constancies, instead of rejecting the idea that perceptual constancies are not sufficient for distinguishing sensation from perception. The reason for this choice is that his idea that perceptual constancies is the mark of perception is more central to his philosophical framework than conjectures about how phenomenal consciousness constitutively fits in with it. In particular, Burge thinks of perceptual constancies as giving a scientific explanation of how objectification is possible. Objectification, recall, is what marked the difference between non-perceptual and perceptual processes. In addition, Burge thinks that proposals about consciousness are at this point mostly inevitably conjectural.[78]

Thus, *if* Burge had held that phenomenal consciousness was necessary to perception, he would have had to accept the idea that phenomenal consciousness is not dissociable from perceptual constancies. This idea has two important components that are worth highlighting. Both relate to the nature of the 'association.'

One is the component idea that unconscious perception may turn out to constitutively associate with phenomenal consciousness. This can be understood at least in three ways. First, consider abnormal cases of unconscious perception such as that occurring in blindsighted people. Blindsighted people are people who apparently do not have any phenomenal experience, but when forced to answer questions about what they see, will 'guess' correctly. Burge could hold of these people that their perceptual representational content constitutively has to associate with instances where this perceptual representational content was conscious. Second, consider that there might be types of perception in *individual subjects* (who are constitutively conscious individuals), that are never phenomenally conscious. Still, this

[78] This whole section is indebted to correspondence with Burge. In particular, he clarified his speculations about the relation between perceptual and conscious content, and the relationship between empirical conjectures and claims of constitutive conditions.

always unconscious perceptual content might constitutively have to associate with phenomenal experience. Third, consider an aspect of perception such as the early stages of visual processing. These early stages of visual processing might always be unconscious, but they can still constitutively associate with phenomenal consciousness. Burge could hold that these early stages constitutively serve, and even always serve, phenomenal experience without ever, necessarily or contingently, being part of it.[79] Burge would likely accept the first and third interpretation if he had held that phenomenal consciousness is constitutive to perceptual representational content. I think he would also accept the second, but I am not sure.

Another is the component idea that perceptual constancies simply associate with phenomenal consciousness by - at least - empirically overlapping with perceptual experiences. This is the idea that the empirical facts might turn out to reveal that every normal system that contains perceptual constancies also has phenomenally conscious perceptions. This is the idea that we might find that individual subjects and perceptual constancies correlate, while remaining open on whether there is any constitutive connection. For instance, Burge does not hold that ants, desert ants in particular, have a capacity for perception representation, because they do not have perceptual constancies (Burge, 2010, pp. 499-500). And so desert ants do not have perceptual representational content. But Burge does hold that bees have a capacity for perceptual representation, because they have perceptual constancies (Burge, 2010, pp. 508, 517). And so bees can have, and do have, perceptual representational content. The statements about perceptual constancies are securely established facts.[80] Aside from these known facts, it might turn out empirically that both or neither of the ant and the bee has phenomenal consciousness. In both these cases, perceptual constancies would not overlap empirically with phenomenal consciousness. In the case of both the ant and the bee having phenomenal consciousness the ant would have experiences without perceptual constancies. If neither the ant nor the bee has phenomenal consciousness, the bee would have perceptual constancies without perceptual experiences. However, it might also turn out that the ant does not have perceptual experiences but that the bee does, just as we would expect if we thought that perceptual constancies overlap with phenomenal content. This kind of a correlation would be evidence in favor of the idea that, in general, perceptual constancies and phenomenal content

[79] It should be apparent that this is a very complex terrain. Obviously I have to draw a limit somewhere to the amount of detail I go into on these issues. However, I think these distinctions are worth keeping in mind.

[80] I grant that the bee has perceptual constancies, even if it does not have, for instance, a neo-cortex like mammals do. I believe perceptual psychologists are well on their way in uncovering general principles.

overlap empirically (i.e. if you find one in an animal, you find the other).[81] And this, in turn, might be evidence for the idea that the correlation is not accidental, but has perhaps implicates a constitutive association.

With both of these components in mind, the "dialectic" is the following: Burge holds that, *if* phenomenal consciousness is constitutive to perceptual content, then we should think of unconscious perceptual representation, in particular, as constitutively associating (at least in the first and third ways) with phenomenally conscious perceptual representation. In particular, if unconscious perceptual representation did not constitutively associate with phenomenally conscious perceptual representations. then the early stages of vision and blindsight would be knockdown objections against their constitutive connection. Blindsight would be an objection because blindsight makes use of perceptual constancies, which is the very mark of perception. The same is true of early stages in vision. However, both are unconscious.

But if Burge holds that unconscious perception constitutively associates with phenomenally conscious perceptual representation, then Burge cannot hold that insects, insofar as they lack phenomenal experience altogether, can have perception. For instance, the bee has perceptual constancies, but we do not know whether it has phenomenal consciousness, so if perceptual representation must constitutively associate with phenomenal consciousness, we do not know if the bee has perceptual content before we know whether the bee has phenomenal consciousness.

If Burge were to accept that phenomenal consciousness was constitutively associated with to perception, the position would partly (by Burge himself) be evaluated by whether it is plausible that perceptual constancies and phenomenal consciousness overlap empirically. In particular, it comes down to such questions as if and why the bee has consciousness or not, and how the answers relate to the fact that the bee has perceptual constancies.

Burge does not think that it is empirically plausible that they overlap, but he concedes that little is known about phenomenal consciousness.[82] Thus, Burge thinks that opinions of whether insects do or do not have consciousness are fairly speculative at this point in time.

[81] After this is shown, one might think that the empirical correlation is non-accidental. That is, there might be some deep reason why they are correlated. In particular, it might shift the positive claim on the person holding that they are not constitutively related, at least, until more is known.

Given that Burge *in fact* thinks that it is implausible that they overlap empirically, Burge thinks that there are no reasons for, and some reasons against holding that unconscious perceptual representations constitutively associate with phenomenally conscious perceptual representations. [83] The reasons for would be speculative, involving theoretical conjectures about consciousness. The reasons against, by contrast, simply note that it is implausible that perceptual constancies have any deep link to phenomenal experiences. This is the most immediate reason why Burge thinks that we should not take perceptual representation as being constitutively conscious.

2.4.5 Reply to Burge

I want to note here how we can reply to Burge, given that we are analyzing Burge's position on phenomenal consciousness with an eye to how one might use phenomenal experience as a supplement to help explain how perception, and perceptual representational content in particular, is for the subject as a whole.

Immediate Reply

The two most immediate ways to convince Burge, without rejecting any of his assumptions, that perceptual representation is constitutively associated with conscious is, 1) to show that conjectures about phenomenal consciousness are not speculative, and 2) to show that phenomenal consciousness does overlap with perceptual constancies, for instance, in insects like the bee. I take these in turn, and then move on to considerations involving rejections of assumptions Burge would not reject.

First, it seems impossible to convince anyone, let alone Burge, that an account involving phenomenal consciousness is not speculative. If speculative is taken in the common sense meaning of conjectural consideration, then it seems that given our scientific progress at this point in time, any account involving consciousness will presently be quite speculative, simply

[82] For instance, in his (Burge, 2010, p. xiii) he says that "It is an open question whether or not consciousness starts, phylogenetically, before perception does."

[83] Again, this is partly from (private correspondence). Methodologically, I want to partly challenge this way of thinking about the matter in the next section. However, Burge has many other reasons for not including phenomenal consciousness. But as a proximate cause, these considerations might serve as the most immediate cause of the negative conclusion. We explore his deeper reasons in the next chapter, chapter 3.

because the study of consciousness is speculative. Thus, I agree with Burge that any account involving consciousness is, at this point in time, speculative.

Second, even if any account of phenomenal consciousness is speculative, one can still show, for instance, the likelihood of the bee having phenomenal consciousness. The science of consciousness is indeed speculative and conjectural, but it is not void of progress. Thus, it is possible to make progress on showing whether perceptual constancies empirically overlap with phenomenal consciousness or not, even though this would be a mere preliminary result.

This then, is the most immediate way to reply to Burge if we want to insist that perceptual content is constitutively associated with phenomenal consciousness. However, we can also trace our steps further back, as it were, and reject assumptions that Burge would not reject.

Rejecting assumptions

If we reject the premise that perceptual constancies are –ceteris paribus - necessary *and* sufficient, we would diverge too much from Burge's framework; I won't discuss it here. [84] Instead we might reject the premise that perceptual constancies are, ceteris paribus, necessary. In this case, there could be another mark of perception that in some way plays a similar objectifying role to that of perceptual constancies. On such an account, both perceptual constancies and *some other thing* would be markers of perceptual representation, such that we can imagine an ant having perceptual content without perceptual constancies if it has *some other thing*, and a bee having perceptual content without *some other thing* altogether since it has perceptual constancies. Alternatively, we could reject perceptual constancies as sufficient, but still hold that they are necessary. If we reject that perceptual constancies are sufficient, then we could hold, for instance, that *some other thing* would have to be present as well. So perceptual constancies would only, ceteris paribus, be partially sufficient for perceptual representation. On such an account, the constitutive presence of *some other thing* would be required for the bee to have perceptual representational content.

It is the claim that perceptual representation constitutively *also* involves *phenomenal consciousness* that I want to argue for. The claim in the context of replying to Burge is that

[84] I add the 'ceteris paribus,' because, of course, Burge does not hold perceptual constancies are sufficient all on their own, without several background conditions, such as associating to successful functions and being formed involving causal relations an external environment.

perceptual representation is constitutively associated with phenomenal consciousness. The broader context – of using phenomenal consciousness to better explain how perception is for an individual as a whole - gives us the claim that phenomenal consciousness is constitutively involved in the explanation of why perceptual representational content is attributed to a *whole* individual. Also recall that we want to better understand how to get from perceptual constancies to the most fine-grained perspectival perceptual representational content. As we will see, I think this too is related to an explanation of perception constitutively involving phenomenal consciousness.

In the next chapter I consider *why* Burge has only had a passing inclination to hold that phenomenal consciousness is constitutively associated with perceptual representation. I look at why Burge thinks constancies are sufficient, and discuss some problems with this suggestion related to attributing representational content to a whole individual.

3 Perceptual Representation and Phenomenal Consciousness

3.1 Introduction

Burge holds that in at least in a "rough, non-critical way, perceptual constancies are necessary as well as sufficient for perceptual objectification and perceptual representation" (Burge, 2010, p. 413). Burge thinks of perceptual constancies as sufficient for perception in the sense that, at least, phenomenal consciousness is not required. Against this claim, I believe that phenomenal consciousness *is* constitutively associated with perceptual representation.

To discuss this fruitfully, we should recall that we are working with three a priori conditions on having perception, and in particular, on having perceptual representational content.

First, perceptual content is constitutively *about* objects or properties of objects.[85] I believe that Burge's account of perceptual content in terms of perceptual constancies fares well in accounting for this condition. Part of the work is done by specification, the anti-individualist considerations involving historical-formative conditions, and part of the work is done by determination, the capacities of the perceptual constancies as such, and their relations to each other in the individual as a whole.

Second, only subjects *as a whole* perceive; perceptual contents are constitutively *for* a subject.[86] Burge upholds this condition by holding that perceptual representational contents, understood as "modes of representation" that "constitute the perspective of the animal or person," "mark how the world is, representationally, for an individual" (Burge, 2010, p. 37). Both the general question that opens his recent book, "what does it take for an individual to represent the physical world objectively?" and a more specific version of this question, uses '*for an individual*' in its formulation. Philosophically, we are not just focusing on perception

[85] More specifically it is of objects in the broad sense. Thus, it is – at least - about properties, relations and kinds. On any given occasion, perceptual content purports to be about particular instances of objects, properties, relations and kinds. Beyond such a priori generic statements, I agree with Burge that it is partly an empirical question what our perceptual content is about.

[86] Since we discussed Perceptual Duality in the introduction, we should now understand how the addition of 'constitutive' to these statements adds welcome nuances.

as a kind of pattern recognition happening somewhere in the subsystems of the individual; it is the individual as a whole that perceives.

Third, we want to explain perceptual representational content in the sense that it "is a perspectival way of representing at the finest explanatory grain" (Burge, 2010, p. 38). Perceptual representation not only involves representing a specific, unique representata, but involves veridicality conditions, and perceptual representational content, that type the specific perspectival way in which the perceptual representation represents what it represents. I believe that, at least if we want a strong interpretation of the second and third condition, then Burge's account can benefit from some supplementation. This was also present in the slight tension uncovered at the end of the last chapter, that when Burge wanted to account for the perspectival way of representing at the *very finest* level of explanatory grain, e.g. shades of qualia, he held both that it may be necessary for perception, and yet that it was irrelevant for typing verbal and "even discriminative similarities" between individuals.[87]

This chapter is roughly divided into four parts. In the first part we look at how Burge argues, through three quick cases, that perceptual processing involving constancies but not phenomenal consciousness can be "attributable to the individual." Burge sketches three cases that he thinks at least make it implausible that phenomenal consciousness is essentially related to perceptual representation. After each case I will comment and challenge Burge's cases. The general critical theme in my comments is that because Burge lacks an account of how the activations of the perceptual constancies become a definite perceptual representational content for a whole subject, he cannot effectively argue against the inclusion of phenomenal consciousness as a constitutive condition on perceptual representation.

[87] This passage certainly suggests that Burge's emphasis on *perspective* and *for the individual* is **also** understood in the third person in a way that is difficult to reconcile with phenomenal experience, *even in principle*. As I see it, shades of qualia – in creatures that have them - are central explananda in understanding "how the world is, representationally, for an individual." It was also clear in some passages that Burge uses the notion of 'perspective' in a "concrete, commonly spatial-directional, sometimes phenomenological" (Burge, 2009a, p. 4) way. It is of course difficult to argue for phenomenology being necessary if a priori criteria can reveal that simple concerns about typing verbal and discriminative similarities *between individuals* can render it irrelevant. This cannot be what Burge has in mind. See (Burge, 2010, pp. 411-412). As we will see, he argues in different ways against phenomenal consciousness – but all of them are *empirically* loaded, not conceptual.

To get a flavor for this, notice that in perceptual psychology, while there are some good theories of how perceptual constancies function relative to environmental stimuli, there seems to be a general consensus that we do not have a definitive account of how such constancies relate to the individual as a whole, whether 'as a whole' is meant as explaining how they can figure in an individual's active behavior, how they relate to memory, or how they relate to phenomenal consciousness. Articles with names like *Where Bottom up Meets top Down* (Mechelli, Price, Friston, & Ishai, 2004), and *Where Vision Meets Memory* (Schendan & Stern, 2008) are seen as attempts to explain aspects of this problem. Even those who, more recently, claim to have comprehensive theories of visual object recognition, as for instance culminating in a "powerful neuronal representation in the inferior temporal cortext" admit that the algorithms describing how the various subnetworks are related to a "common functional goal" is "poorly understood" (DiCarlo et al., 2012). I think phenomenal consciousness is one way of beginning to making sense of how perception can be *for the whole subject*; thus, we must look critically at Burge's argumentation and see to what extent he establishes that perceptual representation can be *for a whole individual* without phenomenal consciousness.[88]

In the second part, I try to deal with this general theme in a systematic way. I outline four thought experiments. Each comes from the consideration of trying to connect the activations of perceptual constancies with a definite perceptual content for the whole animal, which (recall) should help type the perspective of an individual as a whole at the finest level of explanatory grain. The four cases illustrate the explanatory gap that exists between activations of the perceptual constancies as such, and definite perceptual representational content for the individual as a whole..

In the third part of the chapter I further detail some conditions Burge has on perceptual representation. The idea that Burge argues for perceptual constancies as being necessary and

[88] This is also how John McDowell, for instance, can claim that the sensory representations as described by perceptual psychology are only subindividual states in a sensory system, and thus, a priori, cannot be considered perception because perceptions are states of the *whole individual* (McDowell, 2011, pp. 249-250). While I disagree with the overall point McDowell is trying to make in this context, I agree that Burge needs a better explanation of how states of a perceptual system can be states of a whole individual. In particular, this is important insofar as Burge wants to claim that phenomenal consciousness is not necessary for perceptual representation. What excites *me* is claiming that phenomenal consciousness is necessary while still sticking mostly to a third person explanation of why it is important and how it is needed.

sufficient is a simplification. Burge does not *simply* think that perceptual constancies alone determine a perceptual content. In particular, Burge has some constitutive conditions on perceptual representation meant to help make sense of how the perceptual constancies become a definite perceptual content for a subject. I look at two such conditions that might be relevant. First, Burge holds that a particular perceptual content has to be associated with *relevant representational alternatives*. Second, Burge holds that perceptual representation is constitutively associated with *temporal representation*. Both of these conditions go some distance towards making sense of how the activations of the perceptual constancies might be related to a definite perceptual content for a whole individual. However, as I will argue, they also go some distance towards involving the idea is that phenomenal consciousness is a constitutive condition on perceptual content.

Ultimately, I want to suggest that the phenomenal experience can help solve some of the problems related to closing the gap between the "content" determined by the perceptual constancies as such, and the definite perceptual representational content that helps to type a perspective for a subject. I agree with Burge that "phenomenally conscious qualities are present for the individual" (Burge, 2007, p. 406). Thus, one need only establish how the perceptual constancies, and their associated attributes, can become part of phenomenal experience, to establish a definite perceptual content that is both constitutively *about* an object, and *for* a whole individual.[89] In contrast to Burge I claim that this combination is a minimum condition on objective perceptual representation. I believe this will not be satisfactorily settled until we understand phenomenal consciousness, and how the perceptual constancies relate to the individual as a whole, in more detail.

[89] Note that I could have gone another route here, and for instance appealed to phenomenal experience and held that the explanatory appeals that science makes to perceptual constancies map on to what Charles Siewert calls a phenomenology of object constancy (Siewert, 2006a, 2006b, Forthcoming). I certainly believe that something like it is true. While this is part of the motivation for this way of explaining perceptual representation, I want to argue for it, at least partly, from the perspective of science, rather than from my perspective as a subject with access to phenomenal experiences. It is clear that "hey, this science-stuff maps on to this experience" is true in some way; I want to partly understand how this, or even something like this, is possible in a scientific way.

3.2 Burge's Cases for Constancies without Phenomenal Character

Burge thinks that phenomenal consciousness is necessarily attributed to an individual. Burge holds that "when a sensory state, perceptual or non-perceptual, is conscious, it is attributable to an individual" (Burge, 2010, p. 374). However, Burge thinks the way through consciousness is not the only way to attribute perception to individuals. In particular, Burge believes that "consciousness is not the basic factor in determining what in a perceptual system is an individual's and what is merely a subsystem's. This matter is associated with consciousness's not being a necessary condition on perception" (Burge, 2010, p. 374).[90]

Thus Burge needs an alternative way to understand why something is attributed to "whole individuals" (Burge, 2010, pp. 332-337). When Burge speaks of attribution to the individual in this context the attribution is specifically to the *whole* individual. In contrast to cases where we attribute states or actions to the *whole* individual there can be states or actions that pertain simply to the *subsystems* or *peripheral systems* of the individual. For instance, Burge's notion *primitive agency* is explicated as requiring *whole* animal behavior, where this in turn is understood as animal behavior that is a "product of coordination with central behavioral capacities," (Burge, 2010, p. 332) coordination of subsystems, and coordination of central capacities with peripheral realizations (Burge, 2010, p. 331).

One way Burge decides whether something in a perceptual system is attributed to the individual or not is through looking at whether the relevant sciences attribute perception to the whole animal or not. For instance, if perception guides the action of the animal in an appropriate way, then "standard perceptual and ethological accounts attribute perceptions to whole animals" (Burge, 2010, p. 189). On the basis of such scientific accounts, Burge holds that we should attribute perception, and thus perceptual content, to the whole animal.

Burge uses the phrase "attributable to an individual" where I would use "for a whole subject." Both phrases get at roughly the idea something is for an individual.[91] In contrast to Burge I

[90] One might think this comes close to conflicting with what I have said before. For instance, in discussing phenomenal (and rational access) consciousness, Burge says the following: "They are fundamental to what counts as non-derivatively the individual's own" (Burge, 2007, p. 395). The crucial thing to notice is of course that the statement above is relative to the *perceptual system*.

[91] As we have seen, Burge also uses the locution 'for the individual' (Burge, 2010, p. 37).

think that phenomenal consciousness is the basic factor in determining what *in a perceptual system* is an individual's and what is merely a subsystem's. Burge has three cases that he brings to bear on this issue; the case of the bee, the case of blindsight and the case of basic visual processes.

3.2.1 The Bee

Burge claims that "Bees and certain spiders visually perceive color, shape, motion, spatial location, and so on. They exhibit associated perceptual constancies. Whether bees and spiders are phenomenally conscious is unknown" (Burge, 2010, p. 375). I will shorten 'phenomenally conscious' to 'conscious' in most of what follows. From this epistemic situation Burge argues in the following way: We have lots of evidence that bees and certain spiders exhibit perceptual constancies; we do not have any evidence that they are (phenomenally) conscious; therefore, we have no reason to, or should not, believe that consciousness is necessary to perceptual representation.

Notice that Burge argues from the premise that perceptual constancies are the unfailing mark of perceptual content. However, even without this premise, he can claim that the situation regarding evidence is the following:

We have an intuition that bees have perception. There is a strong intuition that we can attribute perceptual content in the Bee to the *whole* animal. It seems that whatever the animal sees may directly figure in the actions and decisions of the whole animal. If bees are conscious, this counts in favor of a view that consciousness is constitutive to perceptual representation. If they are not conscious this counts against view that consciousness is constitutive to perceptual representation. There is no evidence for the view that bees are conscious. Burge can simply hold that the positive claim falls on the proponent of consciousness, and then argue that we have no reason to, or should not, believe that consciousness is necessary – at least until more is known.

There is also a stronger argument Burge could make. This argument is based on positively imagining the sufficiency of perceptual constancies as determining perceptual representation to an individual without appealing to consciousness. Because Burge turns to empirical evidence, I do not think Burge intends the argument to be that strong. In particular, it would establish by imagination what is in question: Whether perceptual constancies not

constitutively associated with consciousness can be sufficient for perceptual representational content that is attributed to a whole individual.

Comment on Bee

First, I agree that we have an intuition that bees have perception, and that it is attributed to the whole animal.[92] I agree with Burge that we do not yet know whether bees are conscious or not. However, I am not so sure about the idea that there is *no* evidence. The problem is that we are still in the process of finding out what should count as evidence for consciousness. We do not know enough about consciousness to know whether the evidence points this way or that. For instance, if consciousness is a gradual property, as those convinced of panpsychism believe, then there is plenty of evidence. Most importantly, we know too little to rule out the possibility that consciousness might figure constitutively in the explanation of bee cognition and behavior.[93]

Second, Burge's argument *for* perceptual constancies being necessary for perception appealed to empirical evidence. Before Burge knew that constancies where constitutively necessary for perception, he did not know whether animals had perceptual constancies or not. It does not follow from this epistemic condition that they are not constitutive to perception. Likewise, we currently do not know whether bees and spiders are conscious; in particular we do not know whether they have phenomenal experiences. So although it might give a certain skeptical conclusion to the effect that we should not speculate without empirical facts at our disposal, we need not accept this conclusion so long as we acknowledge that we are in fact philosophizing in a conjectural spirit.

[92] Insecta is a class within the kingdom Animalia.

[93] That is, there cannot be a lack of *positive* evidence, because we do not know yet what would count as positive evidence. We do not know if what the bee does requires consciousness. For instance, in the case of some lower level animals, we have *no* evidence that there is consciousness in a stronger sense. There is no evidence that consciousness might explain. In the case of the bee, by contrast, there is evidence for behavior that we do not yet know whether consciousness is required to explain or not, because our theories of the cognition of the bee and theories of consciousness are not detailed enough yet; for instance, as I will mention below, bees seemingly have a kind of working memory, and are able to index themselves as being in the 'present' as distinct from the past and future (Randolf Menzel, 2009). These kinds of abilities might be paradigmatically conscious abilities (Bernard J. Baars, 2002, p. 49). I will come back to these points.

Ultimately, whether the Bee is conscious or not only makes a view more or less intuitively plausible. If it turns out that the Bee is conscious then it becomes more plausible to insist that phenomenal experience is constitutive to perceptual representation. If it turns out that the Bee is not conscious, then my view that phenomenal experience is constitutive to perceptual representation would have to be that bee perception is deflationary. In particular I might claim that in such cases, where there are perceptual constancies but not phenomenal experience, there is no perceptual representational content attributable to the bee. This would be analogous to the claims Burge make about all those beings who do not have perceptual constancies, yet appear to navigate and act on sensory stimulations; Burge claims that these beings do not have perception as such. In addition, Burge claims about some senses, like olfaction (the sense of smell), that they usually do not count as perception since they usually do not make use of perceptual constancies (Burge, 2010, pp. 415,419fn.58).[94]

3.2.2 Blindsight

People with blindsight are people with damage to their visual cortex, who, while claiming to be totally blind, are under certain circumstances able to "guess" the correct response in a variety visual discrimination tests. For instance, a researcher might put some item in front of a blindsighted person and have the blindsighted person guess whether it is, for instance, an a dice or an apple they are confronted with. While it seems to the blindsighted person that he or she is just guessing, a number of trials reveals that the "guessing" is quite reliable. Such cases show that people with blindsight do have a range of perceptual abilities that makes use of constancy mechanisms in the way Burge requires for something to qualify as perceptual. The studies done on this assume that people with blindsight claim that they are blind because they have no phenomenal experience that connects specifically to their perceptual abilities. I agree with Burge that this is a fair assumption, and will not discuss this assumption further here.[95]

[94] A recent review on odour object recognition in both rodents and humans gives a central role to "object constancy" (Gottfried, 2010, pp. 633-635). Burge sees no principled reason why olfaction should not be perception, and holds that it is open empirically. Nothing much hinges on this.

[95] Note that taking a person with blindsight's use of constancies as the sole evidence for their having perceptual content would beg the question in the context of arguing that constancies are sufficient, just as taking the lack of consciousness as the sole evidence for their not having perceptual content would beg the question in the context of arguing that consciousness is necessary.

What blindsight can be taken to demonstrate, instead, is that blindsight is, or should, be attributed to *whole* individuals. For this we need an independent way of evaluating whether it is, or should be, attributable to the whole individual.[96] One such way of evaluating whether a sensory state should be attributed to the whole individual is the conditional that if a sensory state "can initiate action by the individual, it is attributable to the individual" (Burge, 2010, p. 373). Another criteria follows from the explication that "perception is available to the whole animal in the sense that it can guide activity or other responses by the whole animal" (Burge, 2003b, p. 158). Thus, if we use guidelines like these we might find that blindsight is attributable to the individual without even mentioning consciousness.

In particular, Burge argues that Blindsight does satisfy these criteria. Burge takes blindsight patients as showing that it is possible for nonconscious perception to be for a whole individual. In particular, he takes their ability to verbally report and act on what is before their eyes as satisfying the criteria that the perceptual state that they are in can initiate action in a way that coordinates subsystems with central capacities of the *whole* individual.

Comment on Blindsight

I think these are interesting criteria for whole individual attribution. Of course, they are not exhaustive, or complete, but they are good as tentative suggestions for how one might decide whether to attribute the perception to the whole individual as a not. [97] I have three comments.

First, note that even if perceptual content in blindsight is attributed to *whole* individuals, this does not alone show that phenomenal consciousness is not constitutively related to perceptual content in blindsight. As we have seen at various points, we can take phenomenal consciousness as constitutively associated with perception, even if it is not necessary on a given occasion, such as in blindsight. In a similar way, Burge does not claim that "*all*

[96] We might get a conflict between the blindsighted patient and the scientist studying him or her. The patient might insist that they are "not really" seeing, while the scientist will insist that "really, you are." Thus, from the third person perspective, I think it is important to be absolutely sure that we have a good way of evaluating whether to attribute the perception to the individual as a whole or not, partly just because we should respect the striking nature of the individuals under consideration.

[97] One piece of evidence for their imperfection is that they differ with respect to each other. However, I grant them as approximately valid because I think these criteria correlate highly - at least in actual instance - with more ideal criteria for whole individual attribution.

perceptions are perceptions by an individual," he merely claims that "necessarily and constitutively, some perceptions in an individual's perceptual subsystem are perceptions by the individual" (Burge, 2010, p. 369). The idea is that "all perceptions, including any that are not strictly attributable to the individual, serve perception by the individual. Fundamentally, it is the individual that perceives" (Burge, 2010, p. 369). Thus, one must really argue from the *fundamental* cases, rather than from the cases that only constitutively associate with those fundamental cases, to argue decisively either way - whether for the idea that consciousness constitutively associates with perception or for the idea that it does not.

Second, the studies on blindsight are not univocally in support of Burge's tentative criteria. The studies on blindsight actually show that these patients are passive in relation to most of what enters into the field of view were they are blindsighted. The novel experimental method used in discovering blindsight is a 'forced-choice' method where the patients are *forced* to make a choice even where no subjective inclination is present (Weiskrantz, 2009, pp. 64-65, 148-155). Aside from some engaging in certain reflex actions, a blindsighted person would not spontaneously initiate action based on their perceptions.

In more detail, where Burge spends some time developing the active-passive distinction concerning action, he mentions the following as examples of passive actions not attributed to the individual: systems muscle spasms, the firing of neurons, saccades by the eyes, shivering, coughing and sneezing (Burge, 2010, p. 333). Many processes that occur in perceptual systems, including transformations of sensory information into perception, are not attributable to the individual (Burge, 2010, p. 369). The typical reason given is that they are not "available to central coordinating agency" (Burge, 2010, p. 333). Thus, the fundamental question really is whether blindsight is "available to central coordinating agency."

In even a modest interpretation of availability to "central coordinating agency," it is questionable whether blindsight qualifies. The most successful theory of phenomenal consciousness, the global workspsace theory, holds that the main difference between conscious and unconscious processes is exactly that "unconscious processing of stimuli activates localized brain regions whereas conscious processing of the same stimuli activates widely distributed brain regions" (Andrade, 2012, p. 602).[98] Without mentioning

[98] Daniel Dennett noted already in 2001 that "Theorists are converging from quite different quarters on a version of the global neuronal workspace model of consciousness" (D. C. Dennett, 2001, p. 221). See (Dehaene &

consciousness directly, one can say that *unconscious processing only activates localized brain regions*. Thus, given that localized brain regions not activating further regions implies that the whole animal cannot act on the processing happening in these local regions, it is not clear that Burge can use his criteria successfully to attribute perception to the individual in the absence of consciousness; at least without weakening the criteria for what should be attributed to a whole individual and what should not.

One may wonder how blindsight patients can produce verbal behavior when it was just asserted that unconscious processing only activates localized brain regions. The key here is of course to realize that blindsight patients are not unconscious full stop. It is merely the visual perceptual states, or even only particular visual states, that are postulated to remain outside consciousness. But when we deal with patients that are fully conscious in most other respects, then this can skew the results. Burge would, for instance, agree that my stomach rumbling is neither an act by the whole individual or a central coordinated move by me as an agent, and yet here I am with a 'verbal report' on this distant aspect of my own being with no problem. The same can be said of a sneeze, a saccade of the eye, and other obscure tics of our own body that I think we should thank phenomenal consciousness for our introspective access to. Thus, it is clear that when something is available to verbal report or behavior *when consciousness is at all present* this alone can tell us little about its constitutive nature.[99]

Third, we should also note that Burge wants perception to "type-individuate a level of agency at which individuals can represent goals of, obstacles to, or threats to their activities, and act accordingly" (Burge, 2010, p. 370). If we want to understand the idea that perceptual states *can* initiate action with this level of agency in mind, then it seems that blindsighted patients only satisfy these criteria in some weak way. Of the little that is known in this matter, there is no evidence in general for the idea that one can unconsciously perform "an instrumental act for an incentive" (Morsella, 2005, p. 1013) without consciousness. And although unconscious "processes can yield elaborate skeletomotor actions" (Morsella, 2005, p. 1013) when the instrumental and incentive systems function independently, the "acts would be fractured and

---

Changeux, 2011) for a review. They discuss whether nonconscious stimuli can produce a "global ignition" on page 215-216. They mostly conclude in the negative, but leave the possibility open.

[99] I think this argument also highlights the power of phenomenal consciousness to integrate and bring together things that would never otherwise have been brought together, such as a tick of the eye, a blushing sensation, a rumble in the stomach and a bluish visual shade.

aimless" (Morsella, 2005, p. 1013).[100] Thus, at least in individuals that we know have phenomenal consciousness, it seems to be an open question whether unconscious perception can be coordinated with whole individual behavior in any interesting way.[101]

3.2.3 Early States in Vision

Burge holds that "certain early states in vision (states in the first micro-seconds of visual processing) may count as perception by the individual, but fail to be conscious. Again, such states exhibit perceptual constancies" (Burge, 2010, p. 375). Burge then simply states that "Perception occurs and figures directly in guiding action" (Burge, 2010, p. 375). Burge notes that it is not known "whether [these early states] are phenomenologically conscious" or not.

Thus, since these states exhibit perceptual constancies, whether they are conscious or not, Burge would hold that they count as perception by the individual.

Comment on Early States in Vision

First, we must ask ourselves whether this is a *fundamental* case of perception. If it is not, then it only serves other states that in turn have to be explained. But this would just push the problem one step further back. For the question is whether phenomenal consciousness is constitutively associated with perceptual representation. In particular, there is a *tremendous mess* of fluctuating information that flows through the early stages of vision. While it is clear that we might be able to understand the activations of particular perceptual constancies at any given time, this by itself is of course not sufficient on Burge's account to (fundamentally) individuate a perceptual content. For one thing, we need attribution to a *whole* individual. It is not enough that perception "figures directly in guiding action" on Burge's account. These early states in vision have to be available to the whole animal in the sense of engaging central

[100] This is perhaps better understood with the following example: "phenomenal states are necessary, not to express or suppress actions but, more precisely, to suppress the action tendencies of response systems … although one can unconsciously respond to harmful stimuli, one cannot unconsciously withstand any degree of tissue damage for some end. … the tissue-damage system is inflexibly concerned with avoiding physical harm … without phenomenal states, this system would cause one to avoid damage even when sustaining such damage is adaptive" (Morsella, 2005, p. 1014). See (Haggard & Clark, 2003) for a similar account.

[101] Again, much of this turns on our intuitions about insects like bees. If bees do not have phenomenal consciousness, then it seems natural that unconscious perception can be coordinated with whole animal behavior. However, this is just what is in question here.

capacities and/or be able to initiate action by the individual. If these requirements are not satisfied, we do not have a fundamental case of perceptual content attributed to the whole individual. And if early states in vision are not a fundamental case of perceptual representation attributed to the whole individual, then these early states must serve some other perceptual representational states that in turn are attributed to the individual. Thus, the problem would be pushed back to the fundamental case.

Second, various constancies active at early stages of vision often have to resolve conflicting cues in one way or another. A cue conflict happens when the two cues activate perceptual constancies that are in conflict with respect to what they "perceptually represent" at any given time. I will sometimes use quotes to indicate that the activations of perceptual constancies can be understood as having "content" in some weak sense. Cue conflict is often intermodal, e.g. between hearing and vision, but also "affects representations *within* specific perceptual modalities" (Burge, 2010, p. 442). For cue conflict to happen, there has to be different "contents" that the early states represent; however, each of these different "contents" cannot be a perceptual representational content attributed the individual, otherwise the individual would have several different representational contents that are in conflict with respect to the perspective and veridicality conditions that are attributed to the whole individual.[102]

3.2.4 Temporary Conclusion

The dialectic above has a theme concerning the relationship between the activations or "content" of constancies and the definite perceptual representational content that we are to attribute to a whole individual. In particular, the dialectic was against the proponent of phenomenal consciousness. Burge tries to account for the gap between the activations or "content" of the perceptual constancies the definite perceptual content attributed to the whole individual, by looking at what and why science tends to attribute perception to a whole individual. In particular, Burge found criteria in notions such as *coordinated subsystems related to agency*, *central versus peripheral systems* and *active versus passive action*.

Whether one is satisfied with these criteria may simply rest on whether one thinks there are better ways to make the distinction in perception between a whole individual and its various

[102] I discuss this further in section 3.3.4

subsystems. For instance, one may well hold that science uses these criteria *pending better* explanations of whole animal attribution.[103]

We have responded to Burge's cases. Let us now move more systematically across this landscape and explore the relation, and possible dissociation, that exists between the activation of the perceptual constancies and the perceptual representational content that is attributed to the individual as a whole. In what follows I simplify the perceptual psychological story that needs to be understood concerning the interconnectedness and dynamical complexity of perceptual constancies. The simplification is in the service of highlighting some gaps in our understanding of *how* perceptual constancies give rise to a definite perceptual content in a fine grained way.

3.3 The Gap Between Constancies and Content

In this section I provide a handful of scientifically inspired but philosophical thought experiments intended to show that there is a gap between the activation of perceptual constancies on the one hand and a perceptual representational content for a whole individual on the other. Given this gap, I think it is premature to hold that phenomenal consciousness is not constitutively associated with perception. The gap also signals things we need to understand in a philosophical account of perception. The section is divided into four parts.

The first part details a case that trades on the idea that there is a difference between what a perceptual constancy itself responds to, when it is maximally activated, and what is perceptually represented for a subject when the perceptual constancy is typically activated. Related to the first case the second case trades on the idea that to represent absences of a particular content, the perceptual constancies responsible for this particular content have to be inactive. Both of these arguments concern the level of activation of different constancies.

The second part details a case that trades on the idea that there is a problem of how the constancies are conjoined together into a unified perspective. Related to the third case the

[103] For instance, even if we want, in purely scientific terms, to account for free will, agency, unified perspective, and so on, this does not mean that science is mature enough to give a satisfactory explanation of them. While the study of visual systems has matured as Burge claims, the study of the way the activations of a given perceptual constancy relates to the whole visual system, and how some specific visual representational content can relate to an individual as a whole, is still in its infancy.

fourth case trades on the idea that different constancies at the same time may have conflicting representational content. Both these arguments concern the unity and disunity of the various constancies that are active in different parts of the brain.

None of these arguments are knock-down arguments against the sufficiency of constancies or for the view that consciousness is necessary if constancies are not sufficient. However, they do provide some theoretical holes that have to be filled in, or disputed, in a prospective theory of how the activations of the perceptual constancies relate to a definite perceptual representational content of the whole individual.

Note that, for our purposes, we can imagine the perceptual constancy as simply being instantiated by a neuron or a cluster of neurons in the context of being on top of a hierarchy as understood by the HTM model of Jeff Hawkins detailed in section 2.2.4. I speak of such neurons or clusters of neurons as having "content" in quotes, where this can be understood as instantiated (given some context) whenever the neuron or clusters of neurons are active.

I list the arguments in the order I presented them here.

3.3.1 Case From Cubism (Max Activation)

The following case considers the possibility of a perceptual constancy being highly active.

Case from Picasso Cubism: The activation of perceptual constancies is a gradual affair. Perceptual constancies respond most strongly to a particular spatio-temporal pattern or sequence. As neuropsychologist Vilayanur S. Ramachandran notes, this propensity of the perceptual system to respond to particular idealized patterns is exploited by artists to create a pattern that stimulates a particular perceptual constancy exceptionally well (Ramachandran, 2004, pp. 42-44). Just such exploitation of the perceptual system occurs when artists make caricatures. More exceptionally, such exploitation is used by Picasso when he draws cubist pictures that give us information about different perspectives of an object simultaneously. As philosopher Jesse Prinz points out, this is evidence for the idea that constancies can represent, for instance, a "Picasso monstrosity" that is "simultaneously representing different perspectives" (Jesse J. Prinz, 2010, p. 313). As Prinz also points out, however, reflection on perceptual experience reveals that our perceptual representational content is not of this kind. For our purposes, we might worry that if the perceptual constancies are simultaneously

representing different perspectives, then it is not clear how this cubist-like representational "content" relates to the representational content and the veridicality conditions that we typically attribute to individuals at particular times; in particular, it is not clear how such "content" could help type the perspective of the whole individual. Thus, the perceptual content presented to an individual is not straightforwardly related to the activation of the perceptual constancies that are required to have the perceptual content in question.

This line of reasoning raises at least two important questions: First, does the fact that a strongly activated constancy respond to these cubist, caricatured, idealized things, mean that even normal activation is relative to such "Picasso monstrosities"? Second, if a caricature or cubist picture strongly stimulates a particular perceptual constancy, so that it becomes maximally active, what exactly is the "content" of this constancy?

The default answer to the first question it would seem is a plain 'yes.' In some sense the study of perceptual constancies does make reference to the precise patterns that they respond to, whether these are conglomerations of different perspectives that cannot exist empirically or caricatures or whatever. Of course, Burge would, and I think should, say that the constancies are constrained by the empirical environment that formed the constancies. However, while this added *specificity* settles *what* thing, if any, is represented, it does not constrain the veridicality conditions this thing is represented under and *determined as*, or help type the perspective of the individual as a whole at the finest level of explanatory grain.

The second question also raises some issues for Burge. We typically recognize the difference between a caricature of a face and a real face, even if the caricature is more easily recognized as being as of a particular face than the particular face itself is. As the above makes clear the caricature might more strongly active the constancies responsible for responding to a particular face. And yet it does not seem sufficient to produce the perceptual representational content with the associated veridicality conditions that we typically associate with the representation of a face. Again, there seems to be a gap between the "content" of the perceptual constancies, and the typical representational content had by the individual when this perceptual constancy is active. Obviously, the solution lies in the contextual elements and so on that provides the whole individual with a definite perceptual representational content. However, it is unclear exactly how this happens.

It is a fact that the perceptual system can respond to very complex spatio-temporal sequences in the absence of consciousness. It is even true that one can produce extremely complex spatio-temporal output, such as uttering sentences or singing songs in the absence of consciousness (Morsella, 2005, pp. 1000-1001). However, there is little evidence to suggest that there is any meaning intended in these unconsciously uttered sentences; and, as we know from the case of the parrot, who speaks perfectly with no inkling of its meaning, we must be careful in our use of behavioral measures to get at striking and presumably deep phenomena such as intentionality and representation. Thus, from the fact that a perceptual constancy is active, in recognizing or even in producing some output, we cannot conclude that the perceptual representational content correlated (scientifically) with this constancy is for the individual as a whole.

3.3.2 Case From Absence (No Activation)

The following argument considers the possibility of a perceptual constancy not being active.

Case from Absence: Consider the difference between a born blind person and an individual with normal vision staring into a dark or empty room. One might say that the person born blind sees nothing. The person with normal vision might also happen to see nothing. However, they do not see *nothing* in the same way. The person with normal vision seems to positively represent an absence of visual content, or positively represent darkness (or emptiness), whereas the person born blind does not positively represent an absence of visual content, or positively represent darkness (or emptiness). One might claim that the person with normal vision represents an absence of visual content because he or she has certain visual constancies that are currently inactive. However, these very same visual constancies are also inactive when the person with normal vision sleeps in dreamless sleep, and yet in that case we would not similarly say that he or she represents the absence of visual content. Thus, one may wonder what is required for an inactive constancy to enter positively into the perceptual representational content as an absence.

What is necessary for being aware of absences in this way? When we look into a room, we represent the space between us and the various things as containing *no things* in a positive way. As Dretske points out in his (Dretske, 2010, p. 18) that we know, for instance, that there was not an elephant in our room when we woke up. The idea is that we can know something based on the absence of a particular perceptual experience if we would have noticed had that

particular perceptual experience been present. However, this kind of epistemic reasoning does not seem available to the blindsighted person. And even if it is available, we can at least say that the blindsighted person's sense of 'I would have noticed' would be weak and less fit for epistemic leverage. Thus, it might be that phenomenal consciousness plays a role in this kind of perception of absence.

It is worth noting that the problem of positively representing absences is a general unsolved problem. Some might even reject that it is a problem. It certainly appears to have a certain mystic quality to it. It would certainly be hard to make the case that Burge leaving out an account of positively representing absences would count against his view that constancies are sufficient (Burge, 2010, p. 540). Burge mentions perception of absences only briefly to note that there *may* be a perception of absences. Thus, it does not even seem to have a central role to play in Burge's account in any case.

This does not bar one from thinking that the perception of absences is a deep problem that we need to solve.[104] I think it helps to highlight the idea that the activations of perceptual constancies, and their correlated content, is in dire need of supplementation, especially if it is to help explain fine-grained veridicality conditions and perspectives of whole individuals.[105]

3.3.3 Case From Space Of Activation (Unitary Activation)

The following case considers the possible unity of the activation of different perceptual constancies.

Case from Color Space: Consider that humans can distinguish a million or so different colors.[106] Certainly, there are not millions of different perceptual constancies that classify all

[104] To see how one might begin to understand the problem of absences, consider that a bit of information can be represented by a light being on or off. An array of lights that could be on or off would give n bits and 2^n possibilities or states of the system; if we walk into the room and no light hits our eyes, *but potentially could have*, we have been informed just as much of the states of the lights as we would have been, had they all been on, or some other arbitrary combination. Thus, if we base our theory of mental representation on an idea of information (as the exclusion of possibilities) the inactivity of *relevant states* should count just as much as active ones.

[105] Recall that I am generally happy with Burge's account of how this content relates to the environment. It is the lack of detail in the explanation of how it becomes content *for a whole* individual that I investigate.

[106] I can *at least* distinguish between 16bit and 24bit colors on my screen. With 16bits the screen uses 2^16,

these different colors. In fact, as Paul Churchland notes, there are only a few perceptual constancies that govern the perception of color (P. M. Churchland, 2005, pp. 532-534). This means that at least some colors that we can distinguish do not have a corresponding perceptual constancy corresponding just to it. From just a handful of neuronal responses, with their capacity for varying their activity, we can model a space of all possible colors that we can see (P. M. Churchland, 2005, pp. 532-534). Thus, one might think that the perceptual constancies explain the fine-grained perceptual representational color content for whole individuals only insofar as they are constitutively bound up with their role as providing axes (e.g. blue-yellow, green-red, white-black) in a space of possible color representations.

This makes it plausible to suppose that perceptual constancies work in concert in the perceptual system.[107] If we take this example of color as holding more generally for perceptual representations involving constancies, then the view emerging is that of perceptual constancies being used in various ways as axes or dimensions giving coordinates in a larger space relative to which these coordinates are interpreted. This view presents a subordinate role for the perceptual constancies, as merely providing the space of possible representations, and never in themselves strictly representing anything for the subject as a whole. Importantly, on this view caricatures might reveal that our perceptual constancies respond to exaggerated, impossible or idealized patterns, not because these necessarily capture something true about objective features, but because the application of several such perceptual constancies *jointly do* capture something true about objective features.

As we saw in the last chapter, Burge entertained the thought that "maybe" phenomenal experience might be constitutive to perceptual representation. While this resolved quickly into a definite "no," it is interesting to note that he entertained this thought exactly when faced with the extremely fine-grained nature, or "shades of 'qualia'" that enter into our phenomenal experience of color. Burge was insisting that these fine-grained shades were irrelevant for "typing verbal and even discriminative similarities," (Burge, 2003d, p. 414). However, to the

---

65535 "colors." If I fix my eyes upon a particular sample color fading horizontally into another color, the limited 16bit resolution of color is clearly visible as vertical bands. (However, as in the case of sound, and brightness, I suppose one can only grasp a more limited range at any particular time.) For a critical discussion of the common claim that humans can distinguish a million or so colors, see (Papineau, 2013).

[107] This is actually an instance of Churchland's Map-Indexing Theory of Perception that I will come back to. The impatient reader can check out section 4.3.1 in chapter 4.

extent that perceptual representational content aims to explain "a perspectival way of representing at the finest explanatory grain" (Burge, 2010, p. 38) it seems we should include in our explanation even such a level of grain.[108]

3.3.4 Case From Binocular Rivalry (Non-unity of Activation)

The following argument considers the possible "conflict" between the activation of different perceptual constancies.

Case from binocular rivalry: Binocular rivalry is a visual effect that can occur when you present different stimuli to each eye: Under certain circumstances, one can present a feature X to one eye and a feature Y to the other eye, and instead of seeing both X and Y or a combination of X and Y, one sees either the one or the other but never both. For instance, if X are a set of vertical lines, and Y are a set of horizontal lines, presented each to each eye, there are circumstances under which we would not perceive both or a mixture of them, but only the vertical or only the horizontal lines at any one time. Much like one can see the Necker cube or the Rubin vase under two different interpretations, people undergoing binocular rivalry have perceptual experiences that alternate or shift between a stable perceptual experience of X, say vertical lines, and a stable perceptual experience of Y, say horizontal lines. During binocular rivalry, we know from neurological studies that the non-consciously perceived stimuli is suppressed to a certain extent even at the early stages of visual processing (Tong, Meng, & Blake, 2006, p. 509). Now suppose we see the horizontal lines, and thus we see the vertical lines unconsciously; when we are conscious of the horizontal lines, the other vertical lines are to some extent suppressed; but it is still there, it can still have some influence on the dominant stimuli. As (Fang & He, 2005) put it, "substantial information in the suppressed eye can escape ... suppression." For instance, it can have further effects in the dorsal cortex (Fang & He, 2005, p. 1383).[109] Since the suppressed vertical lines makes use of perceptual constancies, and can influence dominant stimuli and have further effects in the dorsal cortex, it should – ceteris paribus - qualify as perception on Burge's account, and thus as a perceptual

[108] Whether such a level of explanatory grain is necessary to perception is of course another matter. That depends, in this context, on the relation between perceptual constancies and phenomenal consciousness.

[109] Specifically (Fang & He, 2005, p. 1383) note that "the information from the suppressed eye could be represented 'unconsciously' at the input levels of V1." The details of the empirical situation are too complicated to get into.

representational content. Of course, the same argument applies to the dominant horizontal lines. But then we end up with two perceptual representational contents that are in conflict.

Should we worry about contents conflicting in this way? One might think that the Conditional Principle of Determination (CPD) is satisfied. CDP in this context means that *if* perceptual constancies alone are insufficient to specify the *representata* of perceptual states, they are insufficient to individuate their representational content. Thus, since the perceptual constancies are individuated anti-individualistically, each of the low level perceptual constancies seem to specify their *representata* just fine, X represents *X* and Y represents *Y*. Thus, one might think that there is no problem of conflicting contents.

However, we wanted X to represent *X for a whole Individual*. And we also wanted Y to represent *Y* to the whole individual. But now the individual as a whole has two representational contents, X and Y. The whole individual has veridicality conditions as of vertical lines at some position and as of horizontal lines in exactly the same position at the same time. Thus, the individual will have two contents at the same time that type different perspectives and veridicality conditions.

This is problematic if we want to understand the perspective and function of whole individuals. In particular, if we want a fine-grained determination of what content is attributed to the individual as a whole. As (Morsella, 2005) notes, a whole individual is limited by only producing one unified action-movement at any particular time. An action orchestrated by a whole individual requires a coordinated sequence of muscle-activations, any one of which may be sabotaged by a conflicting muscle-activation. Thus, at least in the service of coherent, initiated action, the whole individual must resolve, before committing to an action, conflicting representational contents; and during the action, suppress conflicting intentions, goals, and desires to avoid conflicting muscle-activations that sabotage whole individual action.[110]

In addition, there is some evidence from research on macaques that, at least the difference between perception in anesthetized and alert macaques is that when perception is unconsciously registered, contradictory "content" in separate perceptual constancies can go unresolved, whereas when the macaques were alert and conscious, the "content" of these "ambiguous local motions features" quickly resolve (within approximately 150ms) into an

[110] The implication is that the fundamental cases of perception do resolve such conflicting contents, otherwise they would not fundamentally be able to coordinate or initiate action in a successful way.

"unambiguous global representation" (Pack, Berezovskii, & Born, 2001, p. 907). [111] Thus, it seems that if we want a perceptual content "available to central coordinating agency" and all the rest, the absence of phenomenal consciousness may present a problem.

Burge (in private correspondence) thinks the case of binocular rivalry is interesting. In particular he thinks there are interesting issues about how perceptual constancies might conflict. Fundamentally, Burge thinks we need to know more about how to localize and specify relevant states at different states of processing. While Burge agrees that consciousness might be a factor in *some* cases, he does not think it is necessary. So even in this case we are short of a knockdown case against the sufficiency of perceptual constancies, or for the necessity of phenomena consciousness. [112]

3.3.5 Discussion

One way Burge might respond to the arguments above is to distinguish between perceptions in those creatures that are phenomenally conscious from perception in those creatures that are not. Recall that Burge called individuals that are phenomenally conscious *individual subjects*. Burge might just agree that consciousness is constitutive to perception in individual subjects, but not constitutive to individuals. This possible response suggests a complaint with a number of the arguments above.

In the specific examples I have used, the contrast is between conscious and unconscious perception *in creatures we know are conscious*. If there is some animal without consciousness but with perceptual constancies, we might get different results. For instance, such an animal

[111] This is also at least *consistent* with the idea that what is measured is the time it takes for this "content" to enter into consciousness, as the minimum estimate of neural processing time required for stimuli reaching consciousness is on the order of 50-80 milliseconds, with the maximum time for more complex stimuli being estimated as high as 500ms (Dainton, 2010).

[112] In general one might think that Burge *could* claim that conflicting and ambiguous contents help serve successful representational contents that are not contradictory or ambiguous for the whole individual. Perhaps that would seem a bit ad hoc in this context, since he used early states in vision as an argument for the sufficiency of perceptual constancies against the idea of phenomenal consciousness as a necessary condition. If it turns out that the perceptual constancies that the blindsight people use are simply serving some successful representational contents, then we are no wiser as to whether these more successful representational contents require consciousness, which would really be at issue.

could conceivably have a perceptual system that does not "allow" contradictory low level representations. In essence, the presence of consciousness seems to ensure coherence to the whole individual and its actions; but if an individual does not have consciousness, there might be other mechanisms to ensure this coherence.

As it stands, however, it seems that Burge has an unsatisfactory account of how the activations and "content" of perceptual constancies gives rise to a definite perceptual representational content for a whole individual. We have seen this theme recurring both in my answers to Burge's cases - meant to illustrate that phenomenal consciousness was not constitutively associated with perceptual representation - and in the cases I just gave to systematically investigate the relation between the activations of the perceptual constancies and the definite perceptual representational content attributed to a whole individual.

I said in the introduction that I would simplify the perceptual psychological story that needs to be understood concerning the interconnectedness and dynamical complexity of perceptual constancies. I further said that the simplification was in the service of highlighting deep gaps in our understanding. I consider all four cases - high activation, no activation, unitary activation, and non-unitary activation - as giving force to the problems that we must face up to. One might think that Burge goes some distance towards satisfying these problems in the further story that he tells about how the perceptual constancies relate to the individual as a whole.

In particular, Burge has two conditions on perceptual representation that helps make sense of the qualification in his statement that "perceptual constancies are necessary as well as sufficient" for perceptual representation only in a "rough, non-critical way" (Burge, 2010, p. 413). Both conditions help give more substance to the explanation on his account of how the activations and "content" of the perceptual constancies become a definite perceptual representational content attributed to the whole individual. However, to the extent that these conditions do give more substance, I argue that they also push Burge closer to the edge of accepting that phenomenal consciousness is constitutively associated with perceptual content.

Next, I say a bit about Burge's further conditions on perceptual representation and then discuss the extent to which they might also (implicitly) go some distance towards establishing the view that phenomenal consciousness is constitutive to perceptual representation.

3.3.6 Two Conditions Burge does Accept

One necessary condition that is part of the larger story for Burge is that perceptual representation must be associated with temporal representation (Burge, 2010, p. 527).[113] As an empirical example, Burge holds that the "perceptual and actional systems of the bee are egocentrically tensed. Present perceptions are present-tensed; some temporal order is retained in perceptual memory, which guides actional representation" (Burge, 2010, p. 522).[114]

This new evidence that Burge mentions comes from (R. Menzel et al., 2005) where they conclude that "the map-like memory in bees is rich and can be used in a flexible way" and further that the general strategy employed in the paper could "further suggests that spatial relations between environmental features appear to be coherently represented in a maplike memory in insects, as is the case in other animals and humans" (R. Menzel et al., 2005, p. 3045).[115]

In a more recent study (Randolf Menzel, 2009) Menzel points out that such navigation with map-like perceptual content requires something akin to working memory.[116] A textbook definition of working memory is that it "refers to our ability to co-ordinate mental operations with transiently stored information during cognitive activities such as planning a shopping trip or reading a newspaper" (Hitch, 2005, p. 307). Menzel gives several reasons for why map-like perceptual content requires something like working memory. One is that the animal needs to "localize itself by recruiting remote spatial memory and integrate it with the current views and directions;" a second is that it needs to "select a goal out of several potential goals and decides to aim for one of these goals;" and third, it needs to "continuously compare the actual conditions with the expected ones."[117] Menzel notes that *bees* in particular (and bees have

[113] He says explicitly that "past perception must be coordinated through memory with present or anticipated representational use." Burge then says that "The coordination requires some sensitivity, at least to temporal order…"(Burge, 2010, p. 527).

[114] Burge also holds that "Egocentric frameworks of some kind are a necessary feature of any perception" (Burge, 2010, p. 201). This relates back to the discussion about what constitutes a perspective at the finest level of explanatory grain in the introduction to this chapter.

[115] Actually, Burge points to (Randolf Menzel, Brandt, Gumbert, Komischke, & Kunze, 2000) in this passage. Burge refers to Menzel's 2005 study on page 206 of *Origins of Objectivity*. The interested reader might also take note of what Burge says, about a 2006 paper by Menzel, on page 517 in *Origins of Objectivity*.

[116] Related comments are made in (Randolf Menzel, Fuchs, Kirbach, Lehmann, & Greggers, 2012).

[117] More direct evidence is offered for this hypothesis: "Active memory processing within expectation theory

perceptual constancies) need something like a working memory to make use of their perceptual representational capacities.[118]

Another necessary condition that is part of the larger story for Burge is that representation must be able to "distinguish instances of that [representational] attribute from instances of other attributes that the individual can discriminate…" (Burge, 2010, p. 466). This is the demand that an individual can only "perceptually indicate and attribute an attribute (kind, property, relation)" if at least something in the individual's psychology is capable of distinguishing instances of the particular attribute from "*relevant representational alternatives"* (Burge, 2010, p. 466). This is what Burge calls the *Relevant Representational Alternatives* (RRA) principle (Burge, 2010, p. 466).[119] This principle holds that one must be able to represent "both different attributes and the same attribute differently" to have a capacity for perceptual attribution, which is necessary to have perceptual representational content (Burge, 2010, p. 446). Burge says, "*given* that (RRA) is met, having perceptual constancies with respect to an attribute suffices to have a capacity to represent the attribute perceptually [my emphasis]" (Burge, 2010, p. 466).

If we keep the 'relevant' in 'Relevant Representational Alternatives' constant, and hold that several different objective features of the environment are all relevant, then for some animal to represent one of them, it has to have a capacity to distinguish or discriminate it from the others.[120] The idea behind this principle is that we do not want to grant that some animal is representing say, a square, or a face, without also respectively being able to represent a circle, or a head. For instance, there are artificial programs running on computers that are capable of sophisticated face recognition; and while there is a one to one relation between the presence

---

predicts that animals will pause at decision points (e.g., in a maze) as they mentally explore available possibilities (Johnson et al., 2007). We believe that our analysis of the search flights shows just such behavior."

[118] See (Neuser, Triphan, Mronz, Poeck, & Strauss, 2008) for a similar result for flies (Drosophila). Flies, too, have perceptual constancies. For instance, see (Tang, Wolf, Xu, & Heisenberg, 2004).

[119] Burge also makes this point in (Burge, 2009a, p. 28) where he notes that "The ability to discriminate a particular must be marked by some ability general attributive representation that under certain normal, standard-making, conditions would be successful in helping the perceptual system discriminate the perceived particular from discernible particulars of other types that are in the same environment."

[120] Whether something is relevant is specified by broadly externalist ways; for instance, by etiological, ethological and teleological concerns.

of a face and this recognition software being active, it is not clear that they represent faces in any meaningful way, when their only *alternative* response, to all other things, reduces to "not face." Even if one did not bar a particular artificial program from being *perceptual* on grounds that they were not for a whole individual (or, for that matter, were individuated anti-individualistically), one could still bar them by their lack of representational alternatives, even if they made use of something like perceptual constancies.

The positive account Burge gives of how the representational alternatives enter into perceptual representational content partly appeals to the counter-factual nature of constancies. For instance that "the capacity to perceptually represent a surface as being a specific orientation and slant is associated with capacities to represent the same surface at different slants" (Burge, 2010, p. 445). In addition he appeals to the idea that the constancies are 'associated' with each other in *some* way (left largely unexplained.)[121]

Temporary Conclusion

I think these further conditions get us closer to the idea that phenomenal consciousness is related constitutively to perception.

First, the capacity for temporal representation that the bee has, its ability to index itself in time, as (Randolf Menzel, 2009) also noted, leverages capacities that look startlingly like working memory, one of the features most commonly correlated with consciousness (B. J. Baars & Franklin, 2003). We might think of this temporal *indexing* as a primitive indexical component in the mental language of the bee. Baars notes in his (Bernard J. Baars, 2002, p. 49) that "working memory depends on conscious elements, including conscious perception, inner speech and visual imagery, each mobilizing widespread functions." If one cannot have working memory without consciousness and one cannot have temporal representation without working memory, and one cannot have perceptual representation without temporal representation, then we cannot have perceptual representation without consciousness.

[121] As Burge himself notes, "neither philosophy nor psychology should be satisfied with an exhaustive contrast between associative connections and representational states" (Burge, 2010, p. 529). (Associationist explanations without quantitative ones are crude, quantitative explanations without postulations of representational states are blind, postulations of representational states without either are empty.)

Burge (in private correspondence) thinks that this is a stretched conclusion. Given the enormous neural differences between bees and humans, we should be careful in thinking that the bee has a kind, 'working memory,' in common with humans and other individuals that we know are phenomenally conscious. And even if it did have such a capacity, we should not draw the conclusion that is associated with phenomenal consciousness just because it is in higher mammals, for instance, which is what Baars likely has in mind. Thus, Burge does not view this kind of inference as giving us any good reason for thinking that bees are conscious. He thinks the question of whether bees are conscious or not is an *open* question.

Second, if we think that there are some deep explanation of how the *actual* perceptual content at any given time relates to the counterfactual possibilities that the animal has capacities to represent, one might want a theory to explain it; in particular, instead of an "associationist" explantion one might want a theory giving a rigorous, mathematical theory of how the actual neuronal activations relate to the various possible activations. As we will see in the next chapter, such an account is given in (Tononi, 2008). As it happens, Tononi's theory is also a theory of phenomenal consciousness. In addition, the problem of how the activation of perceptual constancies can relate to the animal as a whole is partly the sort of problem that the most prominent and successful theories of consciousness focuses on and attempts to solve (G. M. Edelman et al., 2011). Consciousness also seems to figure in drawing a distinction between passive and active *initialization of action* (Morsella, 2005).

Burge (in private correspondence) thinks that each modality of perception comes with both perceptual memory and perceptual anticipation. Based on this, Burge speculates that he might be able to account for perception of absence. I also think the same argument works towards giving a more detailed account of his RRA principle. If Burge can explain perception as involving perceptual memory and perceptual anticipation without involving phenomenal consciousness, then he might be able to give an explanation of RRA and perceptual absence, insofar as the memory integrates relevant representations, and anticipations can be present even in the absence of stimuli.

Thus, even if it turns out that consciousness can help explain how various criteria on perceptual content are satisfied, there may be something else, or a bunch of separate things, that can also explain or satisfy these criteria. Note however that, if criteria are behaviorist at the outset, then it seems always possible, in principle, at least for typical animal behavior, that even if they are produced in part by consciousness, they could have been produced by some

sufficiently complex processes that are not conscious.[122] Insofar as we are out to explain one of the most striking features of mind, however, we should accept that some results can be ruled out by common sense. For instance, it is easy to recognize (although small children might not) that even if a parrot can reproduce any possible snippet of speech, it does not follow that it speak or understands speech. The same applies to people who speak unconsciously. I think the same applies to unconscious perception. Instead of uprooting Burge's views any further, I want to begin supplementing his account.

3.3.7 Why Supplement Burge

As we have seen Burge may find another solution to the problems above other than introducing phenomenal consciousness as a constitutive condition on perceptual content. However, Burge has not yet given a detailed account of exactly how this is possible, and our goal is to establish the conditions under which Perceptual Duality is possible. Thus, I focus on the supplement to his account that *I* think holds the most promise in understanding how the activations of the perceptual constancies become a definite perceptual content for a subject.

While some may consider this attempt to include phenomenal consciousness a daunting move, we should also recognize that there is a high potential payoff in arguing that phenomenal consciousness is constitutively associated with perceptual representational content. There are two main reasons for this.

First, phenomenal experience is already for a whole subject. Phenomenally, we encounter it as a unified perspective. I more or less assume these points, but it is worth noting that these assumptions are shared by Burge.[123] If phenomenal experience, which is essentially *for* a subject, can be constitutive in the explanation of perceptual representational content, which is essentially *about* an object, one will satisfy (PD), the idea that perception is constitutively for a subject and that perception is constitutively about objects. The idea that phenomenal experience is available to the whole subject can also give us a certain epistemic foothold that

[122] Compare: According to (Bourget & Chalmers, Forthcoming) about one in four philosophers think that zombies, creatures physically identical to humans but lacking phenomenal consciousness, are metaphysically possible. By sacrificing some identity (to rough behavioral identity) and narrowing it down to empirical possibility, I think it is still more plausible than zombies.

[123] Maybe he would not agree with the formulation "encounter it as a unified perspective." As we will see in the next chapter, I give empirical point to such formulations.

can be important philosophically, especially for solving traditional philosophical problems related to perception.[124]

Second, phenomenal consciousness is intrinsic to the being it occurs in.[125] If we believe phenomenal consciousness arises from the physical events in the brain, then the phenomenal character has to be explained in accordance with the intrinsic limitations of these physical events going on in the brain.

Thus, there is some philosophical motivation for taking phenomenal character to have perceptual content. Similarly motivated, but taking a different approach, Susanna Siegel (Siegel, 2011) argues directly that one can find perceptual content in phenomenal experience. Instead of arguing directly from phenomenal experience I continue the methodological path set by the explanatory and implementation constraints, and Burge's exemplary account of perception. I will do this by arguing in a scientifically informed way that phenomenal consciousness involves something like what David Lewis called 'phenomenal information' (Lewis, 1990); with phenomenal information, and a model of perception due to Churchland (P. M. Churchland, 2012), I will sketch a supplementary account of perception involving phenomenal information that helps resolve some of the problems encountered in this chapter.

3.3.8 Chapter Conclusion

We have seen through the three quick cases why Burge thinks that phenomenal consciousness is not constitutively associated with perceptual representation; in particular as it relates to helping explain the subsystem/individual distinction. In outlining why Burge thinks that phenomenal consciousness is not constitutively associated with perceptual representation, I commented that the cases he appealed to did little to support the idea that phenomenal consciousness is not constitutively associated with perceptual representation.

In the four thought experiments I have gave reasons to think that it is not clear how the activations and "content" of the perceptual constancies become a definite perceptual representational content attributed to a whole individual. Thus, there is still a lot of work to do

[124] Certainly many philosophers assume this to be the case, and quarrel with those that do not. I want rather, in a scientifically informed way, to explore *how* phenomenal consciousness can have epistemic import.

[125] One might dispute this. Burge would not (Burge, 2003b, p. 159). I want rather, in a scientifically informed way, to explore *how* phenomenal consciousness is intrinsic, if it is.

to understand exactly how perceptual constancies can give rise to a definite perceptual representational content that is for the individual as a whole, particularly as this is understood to include the finest level of explanatory grain.

While Burge perhaps tentatively thinks that further associative supplements like temporal representation and the RRA principle can begin to explain the gap between perceptual constancies and the indvidiual as a whole, I think that these only go some distance towards the view that we need phenomenal consciousness in this further explanation of how perception fundementally can be *for* a whole individual and *about* objects. In particular, I think these associative supplements cries out for quantitative explanations that can more rigorously explain how perceptual constancies, relevant representational alternatives, anticipations, memory, and temporal representation can come together in a perceptual representational content for an individual as a whole, and an individual's whole perspective in particular, particularly at the finest level of explanatory grain.

In the next chapter, chapter 4, I sketch a view of phenomenal experience as being a kind of phenomenal information. I develop a strong a weak interpretation of this. I then outline Paul Churchland's Map-Indexing Theory of Perception. This theory shows, in fine-grained way, how perceptual constancies might relate to a definite perceptual representational content. It also makes phenomenal predictions. Finally, I show that Churchland's theory, in combination with both strong and weak phenomenal information, helps solves some of the problems encountered in this chapter. Most speculatively, and spectacularly, on the strong interpretation, phenomenal consciousness *just is* Integrated Information, which I argue is required to fully make sense of Churchland's theory, in particular, how his notion of '*indexing*' is, fundamentally, possible.

4 Phenomenal Information and The Map-Indexing Model of Perception

4.1 Introduction

In this chapter I sketch a way to supplement Burge's account. As the title suggests, there are two key ideas that will figure prominently in this sketch. First there is the hypothesis of 'Phenomenal Information.' Second there is the 'Map-Indexing Model of Perception.' I give a brief sketch of each here and discuss them in detail as the chapter progresses.

The hypothesis of phenomenal information is the idea that phenomenal consciousness is informative. What do we mean by informative? I follow David Lewis in holding that the "genuine hypothesis of phenomenal information …treats information in terms of elimination of possibilities" (Lewis, 1990, p. 94). On this notion of information, Lewis explains that "When we lack information, several alternative possibilities are open, when we get the information some of the alternatives are excluded" (Lewis, 1990, p. 94).

The Map-Indexing Model of Perception is a model due to Churchland (P. M. Churchland, 2012) that tackles the problem of how perceptual constancies (what he calls 'feature detectors') can give rise to perceptual representational content. In short, Churchland sees feature detectors as being axes in a multidimensional space of possible representations.[126] The 'map' is the relevant space of possible representations, and the 'indexing' is the act of determining a particular place ("coordinate") in such a space.

Why do these two ideas help supplement Burge's account? In the previous chapter we basically wanted two things. On the one hand we wanted to understand how perception could be *for* a whole subject in a stronger sense than what Burge had resources to explain in detail. On the other hand we wanted a fine-grained account of what the perceptual representational content was in any given case, and how this relates to the activations of perceptual constancies.[127]

[126] To see how this might work in the nitty-gritty of theoretical neuropsychology *in humans*, I urge the reader to examine the work of (DiCarlo et al., 2012).

[127] I should perhaps say activations *and operations* of the perceptual constancies, given the latter part of the last

The Map-Indexing Model of Perception gives a fine-grained account of how the activations of perceptual constancies can give rise to a definite perceptual representational content. The hypothesis of phenomenal information is the idea that the phenomenal consciousness is informative. I conjecture that when a space of possible representations, the 'map,' is *indexed* by the subject *as a whole*, then perception constitutively involves phenomenal information.

I begin by outlining a weak and a strong version of phenomenal information. To give a 'weak' version of phenomenal information I appeal to various philosophers who have made use of notions such as Quality Space, Information Space, and Activation Space; I label these '*Structural Information*,' because they all involve the idea that the individual excludes a structured set of possibilities. It is also clear that while phenomenal experience may be constitutively characterized in terms of structural information, the opposite is not true. To give the 'strong' version of phenomenal information I appeal to Tononi's theory of *Integrated Information*, which is an intrinsic measure of information conjectured to explain the very nature of phenomenal consciousness. By intrinsic I mean that it aims to be immune to worries that any measure of information is always relative to a standpoint or interpreter.[128]

After elucidating the 'strong' (integrated information) and 'weak' (structural information) version of phenomenal information, I detail Churchland's Map-Indexing Model of Perception (MIMP). I also show that MIMP entails some phenomenal predictions. Given Churchland's MIMP we have a 'weak' and 'strong' interpretation of it, according to whether we take it to involve structural or integrated information. While I argue that even the 'weak' interpretation goes some distance towards settling the gaps between the activation of perceptual constancies and a perceptual representational content for a whole individual, I argue that the 'strong' interpretation goes an even greater distance in bridging those gaps.

---

chapter as getting involved with the RRA principle, temporal representation, and so on.

[128] In this sense Tononi's theory aims to mathematically explain the intrinsic nature of consciousness much like our best physical theories aim to explain the intrinsic nature of physical forces by equations. I simply note the analogy. I take no detailed stance on how such theories are to be understood (as in philosophy of science).

4.2 Two Cases: Phenomenal Character as Informative

If we now focus on phenomenal consciousness, our aim is to show that one can find *information* in phenomenal characters considered as such. According to David Lewis, there is an important hypothesis in this direction, which he calls "The genuine hypothesis of phenomenal information…" which "treats information in terms of elimination of possibilities" (Lewis, 1990, p. 94). Note that while Lewis defines this genuine hypothesis of phenomenal information as necessarily excluding materialism, one need not do so.[129] For instance, in a reply to Lewis, William Lyncan defends "a perfectly good sense of 'information', more finely grained than that of 'fact,' in which there is after all phenomenal information, indeed (!) phenomenal information inaccessible to objective, third-person science" (Lycan, 1995, p. 254; 1998).

The idea of information as tied to a *range of excluded possibilities* is a general idea of information, for instance, that Jaakko Hintikka attributes approvingly to Quine. More specifically, the idea is that "the more alternative possibilities a proposition excludes, the more narrowly can we restrict our attention, which is the manifestation of having more information" (Hintikka, 2007, p. 191). For our purposes, we can substitute 'proposition' with 'phenomenal character.' According to Hintikka, Quine says that "such use of the notion of information would be feasible only if we could specify all the different possible states of the physical world" (Hintikka, 2007, p. 191). Against this Hintikka notes that we can simply speak in terms of "systems" and specify the possible states of such systems; in this sense, the (actually) excluded possibilities are always relative to the (counterfactual) possibilities of the system.[130] I will argue that this notion of information is constitutive (on a weak and strong interpretation of constitutive) to the explanation of phenomenal character, considered as being relative to an intrinsic "system," e.g. the brain.

I do this in two steps. Both center on the idea that phenomenal character constitutively involves a range of excluded possibilities, and for this reason count as informative.

[129] For a discussion, see (Perry, 2001, pp. 159-163)

[130] Note that whether such a notion is feasible might decide the fate of Burge's Relevant Representational Alternatives (RRA) principle. In particular, it will decide whether it is possible to precisely quantify the Relevant Representations that a system has.

First, I give some reasons to believe that each phenomenal character constitutively excludes a certain range of possibilities, and that the possibilities that each phenomenal character excludes is structured according to difference and similarity relations. This first step is meant to show that the phenomenal character of experience can be at least constitutively associated with a kind of highly structured, quality space of possibilities. Second, I show that there is a good - although speculative - scientific theory (the Integrated Information Theory due to Giulio Tononi) that explains phenomenal character in terms of the structured set of possibilities that the state ocurrently excludes.[131] Again, like Burge shows that perceptual constancies are constitutive to perception by appealing to psychological explanations, I show that structural information has a central explanatory role in the explanation of phenomenal experience by appealing to quite a few philosophical frameworks. Then I show that it is possible, in a scientifically grounded way, to hold that Integrated Information has a central role to play in the explanation of the nature of phenomenal consciousness.[132]

4.2.1 Quality Space (Structural Information)

Quine introduced his crucial notion of a 'quality space' in (Quine, 1960, pp. 82-84) to solve the problem of how a subject could distinguish some things as more similar to each other than others. For instance, a child might distinguish cyan as being more similar to green than to red. As Quine notes, it follows naturally from considerations on such similarity and difference relations to "credit the child with a sort of prelinguistic quality space" (Quine, 1960, p. 82).[133] Notice that a quality space is not just a set of discriminations, but a structured set of discriminations. The structured set of information gives rise to many dimensions along which particular discriminations are related according to similarity and difference relations.

---

[131] One might think that the comment from Lycan above is inconsistent with a science of such phenomenal information. However, it is just the accessibility of such information that is 'intrinsic.' Just like a mathematical formula of gravity is not itself gravity, the scientific explanation of phenomenal consciousness in terms of integrated information is not itself integrated information.

[132] Specifically, Tononi identifies phenomenal consciousness with a *complex* of Integrated Information; the space and timescale on which a *single* system reaches the maximum quantity of integrated information. He formalizes all this in (Balduzzi & Tononi, 2009; Tononi, 2004; Tononi & Sporns, 2003). Responses by philosophers to the theory have been mixed. David Chalmers, perhaps most neutrally, has commented that "It's the sort of proposal that I think people should be generating at this point: a simple and powerful hypothesis about the relationship between brain processing and conscious experience," (Zimmer, 2012).

[133] See (Broughton, 1981) for an indepth discussion of Quine's use of 'Quality Space.'

Chalmers comes to a similar notion partly from reflection on phenomenology. Chalmers simply notes that we "find information realized in our phenomenology. States of experience fall directly into information spaces in a natural way" (Chalmers, 1996, p. 284). Chalmers justifies this by appealing to the "natural patterns of similarity and difference between phenomenal states, and these patterns yield the difference structure of an information space" (Chalmers, 1996, p. 284). While Quine wanted to rigorously rely on behavioral measures, Chalmers finds it natural to appeal directly to the difference relations in our phenomenology to establish that it involves (what I call) structural information.

Shoemaker agrees with Chalmers point that, if we assume physicalism there "must be a similarity ordering on the physical states that realize experiences having complex contents; there must also be a similarity ordering on the physical properties of these states that contribute to their similarity or difference along particular dimensions" (Shoemaker, 1996, p. 63). Shoemaker comes to such conclusions based on philosophical considerations akin to the argument from color-inversion.

Shoemaker agrees that two *phenomenal characters*, green in individual x and green in individual y, can have different representational content in virtue of individual x and y's differing relation to the external environment. However, he also believes that there is a striking similarity between them, insofar as they have the same phenomenal character, green. In particular, "given a background of externalist constraints on content, the qualitative similarities and differences can be said to determine the intentional similarities and differences; and, with the same qualification, the qualitative character of an experience can be said to determine its representational content" (Shoemaker, 1996, p. 63). Thus, while it *may* be possible to invert our quality color space, this is only possible insofar as we keep most or all of the similarity and difference relations intact. In particular, Shoemaker holds that the inversion thought experiments, that he focused on already in (Shoemaker, 1982), are not to be taken as arguments against using phenomenal characters as a resource in content-individuation. In fact, Shoemaker argues "that a satisfactory account of perceptual experience requires qualia" (Shoemaker, 1996, p. 61). Shoemakers thinks that phenomenal characters are constitutive to an account of perceptual experience. He argues this by holding that perceptual

experience is constitutively related to a qualitative space, and that the "notion of a quale is implicit in the very notion of a quality space" (Shoemaker, 1996, p. 62). [134]

To make sense of perceptual representational content, Shoemaker appeals to two different things. On the one hand, Shoemaker holds that "A color experience represents an object as having a 'phenomenal' property that is constituted by a relation to sense-experience" (Shoemaker, 1994, p. 35). This 'phenomenal' property is what according to Shoemaker "gives the experience its phenomenal character" (Shoemaker, 1994, p. 35). On the other hand, "the experience also represents the object as having a certain color" (Shoemaker, 1994, p. 35).That is, the representation of this property also constitutively depends on its external relation to a specific, say, reflectance property or photon wavelength.

Churchland comes to a similar conclusion based on philosophical and scientific reflection on psychological theories of perception. Churchland holds similar to Shoemaker that the representational content is given by two things. For Churchland, this means that the content of a given position in an information-space ("a point in activation space") is given by two different things. It is given both by "(1) its spatial position relative to all of the other contentful points within that space" and by "(2) its causal relations to stable and objective macrofeatures of the external environment" (P. M. Churchland, 1998, p. 8).

While Chalmers, Shoemaker and Churchland may hold that phenomenal character is always associated with or even constitutively explained in terms of structural information, only Shoemaker holds that a notion of a quality space is constitutively linked to phenomenal character.[135] All three, with the very slightly possible exception of Chalmers, hold that structural information is constitutive of phenomenal experience.

Shoemaker partly argues from philosophical inversion thought experiments, Chalmers partly argues from reflection on phenomenology and Churchland partly argues from perceptual psychological considerations. That so many different thinkers and methodologies converge on

[134] Note that this removes one reason for supposing that phenomenal consciousness is not necessary, namely, that different phenomenal characters can represent the same thing, and the same phenomenal character can represent different things. For now we have some definite structural information invariantly attached to a particular phenomenal character, although it does not directly pertain to what the representation is *of*.

[135] He does not simply mean this in trivial terminological sense, that quality space consists of the word 'quality' and for this reason somehow constitutively entails 'qualia.'

the explanation of phenomenal experience in terms of such a similar notion of structural information should be taken seriously. In particular, that these thinkers give an account of phenomenal characters explained partly in terms of a quality or an informational space suggests that structural information is constitutive of phenomenal characters.[136]

Related to the view I want to espouse here, Robert Stalnaker in his book *Our Knowledge of the Internal World* (Stalnaker, 2008) outlines a view which *makes use of* this feature of phenomenal characters that they are informative. In a recent interview, Stalkner explains that "while the information expressed by saying 'seeing red is like this' is not detachable from the context in which it is expressed or thought, it is (I want to say) a piece of information about what the world is like, which means that we understand the content of the thought by the way it distinguishes between alternative possibilities" (Marshall, 2013). Being on a track similar to the one I am on, Stalnaker wants to understand how such information can be "a way of connecting the content of a thought to the thinker who is thinking it" (Marshall, 2013).

Next, we look at Tononi's theory of the phenomenal consciousness. Tononi gives a scientific, formal, testable account of a kind of information that according to the theory is both necessary and sufficient for experience. To simplify, Tononi conjectures that whenever there is a complex of Integrated Information there is consciousness.[137]

[136] Note that this way of talking of phenomenal characters squares with research on the phenomenal consciousness. For instance, when (G. M. Edelman et al., 2011, p. 4) notes that "the experience of qualia occurs in each individual as a set of discriminations: 'heat' is not 'green,' 'green' is not 'touch,' etc." Thus, even researchers in the field think of the very nature of these phenomenal characters as *somehow* inherently, intrinsically excluding these other possible phenomenal characters (or neurological states).

[137] Less simplified, Tononi formalizes a measure of integrated information, and holds that whenever integrated information reaches a maximum, there is what he calls a complex of integrated information, which is he conjectures is the formalized essence of phenomenal consciousness (Tononi, 2008, p. 216). (In the sense that a mathematical theory say, of gravity may be conjectured to be the formalized essence of gravity.)

4.2.2 Integrated Information Theory (Integrated Information)

According to Tononi's theory, there are two aspects that make the phenomenal experiences what they are. First, phenomenal experiences are rich in information. They are *informative*. Second, this information is *integrated*. Phenomenal experiences consist of *integrated information*. To argue for and make clear these two points, Tononi makes use of two philosophical thought experiments.

To characterize *informativeness*, Tononi uses the *The Photodiode Thought Experiment:* Imagine a photodiode responding only to light. The photodiode has two options, to remain off or to turn on. Each time there is light causally interacting with the photodiode in a certain way, it will turn on. Does the photodiode see light?

According to Tononi, a crucial difference between us and the photodiode with respect to phenomenal experience is the space of possible discriminations it has. When a photodiode makes a "discrimination between two alternatives, the detector in the photodiode generates log2(2) = 1 bit of information" (Tononi, 2008, p. 217).[138] When we humans, or other animals, see 'light,' by contrast, we discriminate between vast numbers of alternative things that we could have seen. In Tononi's words, "we are implicitly being much more specific: we simultaneously specify that things are this way rather than that way (light as opposed to dark), that whatever we are discriminating is not colored (in any particular color), does not have a shape (any particular one), is visual as opposed to auditory or olfactory, sensory as opposed to thought-like, and so on" (Tononi, 2008, p. 218). In contrast to the photodiode, we produce much more information when we see light than simply 1 bit of information. We exclude all the trillions of possible states that we have in our repertoire.[139]

[138] Recall that log2 of 2 is simply the power we have to raise two to, to get two. We have to raise two to the power of one (2^1) to get two. Thus, if there were three photodiodes, there would be discrimination among eight alternatives; and since we need 2x2x2 (2^3) to get eight, we would now have 3 bits of information – that is, a situation we can characterize by three binary digits.

[139] To get a feel for the problem, take some time to think about the quantity of distinct possible conscious states that a single human might be able to experience. Every frame in every movie we have ever seen only scratches the surface of the different conscious states we might be in. One can sit alone in a movie theatre and perceive one frame after the other, each corresponding to a different conscious state, … Then consider that each and every one of these possible states have a correspondingly specific physical substrate.

To understand why we also need *integration*, Tononi uses *The Camera Thought Experiment*: imagine a "digital camera, say one whose sensor chip is a collection of a million of binary photodiodes, each sporting a sensor and a detector" (Tononi, 2008, p. 218). With the camera, in contrast to the photodiode, it seems we get a huge amount of information. If we consider the camera "as a single system with a repertoire of 2^1,000,000 states" (Tononi, 2008, p. 218) then we have a million bits of information. When the camera takes a picture, it settles into one of 2^1,000,000 states. If phenomenal experience was just a matter of discriminative information, then it seems that the camera should be *very* conscious.

What is lacking? Tononi starts with the intuition that this system, "the camera chip is not an integrated entity: since its 1 million photodiodes have no way to interact" (Tononi, 2008, p. 219). The problem is of course that there is no special vantage point from which we might consider the camera as a single system or entity. What is needed is ultimately "an intrinsic point of view associated with the camera chip as a whole" (Tononi, 2008, p. 219). This is needed because, as it stands, there is no objective measure that uniquely settles which states should be grouped together; that is, there is no intrinsic measure of which alternative representations are relevant, and which are not.[140]

How can we measure and quantify what should count as a single system? According to Tononi, we begin to understand such an objective measure if we cut the camera sensor chip in two, such that it consists of two chunks of 500 thousand photodiodes each. By doing this, we find that the information contained in the whole does not change at all. In fact, if we cut the camera sensor chip into "1 million pieces each holding its individual photodiode, the performance of the camera would not change at all," (Tononi, 2008, p. 219). The components in this system work independently of the whole.

In contrast to this, Tononi argues that conscious beings only "discriminate among a vast repertoire of states *as an integrated system* [my emphasis]" (Tononi, 2008, p. 219). For instance, in popularizing his work (Tononi, 2012, pp. 160-162), Tononi uses the example of an Italian speaking person watching the text 'SONO' on a screen. The person would read and think of this as "I am." However if we split this person's brain in two along the corpus

[140] This is a general problem, as Hilary Putnam and others have argued; one can interpret almost any system to instantiate any computation. For instance, see Putnam in his discussion of David Lewis (Putnam, 1991, pp. 94-100). Computation naturally involves information as we typically think of it, and so the argument implies that any information too can be instantiated by almost any system.

callosum for instance, the left hemisphere might end up reading 'SO' which means "I know," while the right hemisphere ends up reading 'NO,' which means "no." The intuitive idea is that we cannot separate and reduce the various discriminations we make into separate component parts, because they work in a way that is explicable only holistically, unlike the camera.

If this were simply a thought experiment showing that our brain processes information holistically, nothing much would follow. However, Tononi follows it up with a quantitative measure of integrated information. This measure, that he calls phi (Φ), is an exact measure of how much information is generated above and beyond the information generated by the parts of the system (Tononi & Sporns, 2003). [141] I won't go into how this quantity is measured here. There are other efforts in the same direction.[142]

The upshot is that the theory can use this measure in explaining phenomenal experience; has testable predictions; and retrodicts that we should find experience where we do in fact find it. While specific predictions are forthcoming due to computational limitations, Tononi holds that it can at least account for why experience occurs on such and such a time-scale and spatial scale. In addition, Tononi holds that it can account for why specific parts of the brain such as the cerebellum, or specific states, such as dreamless sleep, anesthetics, and epileptic seizures are not associated with consciousness, even though they often involve just as much, or more neuronal activity. Further, Tononi argues in his (Balduzzi & Tononi, 2009) that the phenomenal characters of experience can partly be accounted for by a 'geometry of integrated information.'[143]

[141] It can really be interpreted as an objective measure of anti-reductionism in terms of information theory. It is a measure worth its salt even divorced from its claim to explain experience.

[142] An overview and introduction to the notion of integrated information and how it relates to complexity and networks of the brain, see (Sporns, 2010, pp. 293-298).

[143] While theories of consciousness are necessarily speculative, this is a simple, quantitative theory with testable predictions. It has also created quite a buzz. Chalmers notes about this theory that 'it's the sort of proposal that I think people should be generating at this point: a simple and powerful hypothesis about the relationship between the brain processing an conscious experience" (Zimmer, 2012) .Cristoph Koch, chief scientific officer at the Allen Institute for Brain Science in Seattle, who with Francis Crick arguable made the scientific study of 'consciousness' respectable, have recently come out as a strong advocate of the theory. Among other things, he has stated that it is "the only really promising fundamental theory of consciousness" (Zimmer, 2012).

Although highly intriguing, it is outside the scope of this essay to explain in detail how Tononi's theory works out. In general it is also outside the scope of a philosophical essay to decide whether a scientific theory is correct or not. What this case shows is that there at least is a scientific theory, consistent with the known facts, that phenomenal experience is created by a certain kind of information integration in the brain.

4.3 The Map-Indexing Model of Perception

We uncovered two kinds of information related to phenomenal consciousness. The first makes the case that there is a kind of structural information instantiated whenever there is a phenomenal character instantiated. The second makes in addition the stronger case that whenever there is integrated information there is phenomenal character.

Now we must see whether structural or integrated information can help make sense of how the "content" of perceptual constancies can be for a *whole subject*. In this section I argue that there is a theory of perception that –at least - makes use of structural information to relate the activations of the perceptual constancies to a definite perceptual content. In particular, Paul Churchland's Map-Indexing Model of Perception (MIMP) relates the activations of the perceptual constancies to a coherent, precise perceptual content. In addition, Paul Churchland argues that this way of relating the perceptual constancies to perceptual content also helps give predictions about phenomenal character. In doing this, the theory also helps explain how it is possible that the phenomenal character of experience *can possibly* relate to information concerning the activations of the perceptual constancies, effectively going some distance towards showing how phenomenal information is possible.[144]

[144] Note that Paul Churchland has traditionally been an eliminativist about phenomenal character. Given the prospects of empirical theories, he is not shy of giving predictions about the phenomenal character of colors on the basis of the MIMP. In (P. M. Churchland, 1986, p. 104) he holds that the "visual sensation" of colors are "literally identical" with a certain set of specific activation of neurons. To a certain extent, this seems clearly untrue, stemming from what Burge would call a reductionist, materialist ideology. It simply does not make sense to say that only a few neurons can give rise to such a specific sensation of color, as there are bound to be a similar, physically identical set of neurons somewhere else in the brain that activate in the same specific fashion. On the other hand, it does, and indeed must have some truth to it, given the successful predictions of such theories. (The obvious solution is that the context matters, but *how*, exactly?)

The basic problem is of course that structural information, as such, is not sufficient to solve the kinds of problems that Tononi's Integrated Information is developed to solve. In particular, structural information is neutral on the topic of how it is possible, in a deep sense, to objectively measure, from the intrinsic perspective of the system, the relevant space of possibility for the neuronal activations doing the representational work. In some obvious sense, it seems that the context that these specific set of neurons Churchland appeals to must enter into the very nature of phenomenal experience. And the quantification of this is partly what Tononi's theory is motivated to formalize.[145] Ultimately, I will speculate that Tononi's theory of integrated information might help complete the Map-Indexing Model of Perception (MIMP). However, let us first look at how Churchland's MIMP works.

4.3.1 Map-Indexing Model of Perception

In the previous chapter I used the case of a color space to argue that an account of perception focusing on perceptual constancies themselves as giving rise to content, will lose sight of the richness of perceptual content. In particular, I noted that humans can discriminate around 1 million colors, in spite of there likely being only a few constancies responsible for color perception. Each of these shades can be about the environment in a particular way. For this reason we might suppose perceptual constancies to be constitutively associated with their role in some representational space.

However, this exclusive focus on color might be too narrow to establish something general about perception. One might worry that this feature of color - that it has a straightforward representational space associated with it - is unique to color representation and a few other kinds of perceptual representation only, such as perhaps spatial and geometrical representation.

In this section I introduce another example that is sufficiently distant from color (and these other suggested cases) that we might consider the model used, if not universal to perceptual representation, then at least very commonplace.[146] The example is from (P. M. Churchland, 2012). It highlights in an intuitive way how perceptual constancies might give rise to a

[145] And if it does not succeed, it seems that some other, similar theory must take its place, if phenomenal character is to relate essentially to some kind of structured information.

[146] Of course, if one does not agree with the assumption that the model captures how perceptual (face) representation works, then the objections are multitude. I cannot discuss this here.

definite perceptual representational content. In particular, the example focuses on how perceptual representation of faces might work. The upshot will be that face representation, just like color representation, seems to rely on constancies only insofar as the perceptual constancies give the dimensions and thus "activation space" from which a perceptual content has to be determined ("indexed").

4.3.2 Face Space[147]

Churchland details a specific face-discrimination network developed from Cottrell (Cottrell & Metcalfe, 1990). The network is a simplified model of how the visual system might work. It has "4096 grayscale sensitive 'retinal' cells, arranged in a 64 x 64 grid, on which various images can be 'projected'" (P. M. Churchland, 2012, p. 64). From this 64x64 grid, this layer of 'retinal' cells, there are 80 different ways that this layer is processed by constancy mechanisms. Thus, the 4096 cells taking input end up compressed into 80 invariant features, or constancies. Here, each of the 80 invariant features are features taking into account all of the 4096 input cells in a way that is similar to Jeff Hawkins' way of conceiving of constancy mechanisms that I detailed in section 2.2.4.

The idea is that these 80 invariant features are features that faces, as a whole, might have. For instance, one of the invariant features might respond exclusively to the femaleness of the face. Another might recognize maleness. Another recognizes that it is a face; that is has eyes, nose, mouth etc. in the appropriate place. And yet others might recognize the properties of faces that make them seem young, old, smiling, and so on (P. M. Churchland, 2012, p. 62).

For instance, the constancy detecting the masculinity of a face might respond to features involving the eye-brows, that they are furry and close to the eyes, and features involving the jaw, that it is strong and edgy. Meanwhile, the constancy detecting the femininity of a face might respond to features of the mouth, that the lower lip is voluminous, or features of the eyes, that they are more open and exposed.

[147] Pun somewhat intended. 'Phase space' is a space where all possible states of some system are represented. In the color space case, a few neurons might be thought of as axis in a multidimensional space, with each point in the space representing a unique state of these neurons taken collectively, with points close by representing similar collective states. Thus Daniel Dennett's philosophical lexicon entry on 'Churchland,' as "n., (1)Two-ring traveling circus …at which philosophers are giving entertaining religious instruction in Science and nothing to eat but "phase space sandwiches" (D. Dennett).

The idea is that the constancies capture the characteristics of femininity or masculinity that deviates from the average face. We assume that there are various invariant features that might be detected by the constancy mechanisms. On this face recognition model, there are eighty or so other dimensions that a face will typically vary in. Thus, when we actually see a face, there will be activation on several of these constancy dimensions. And these dimensions will activate in different degrees according to the strength that the actual face perceived is detected as an instance of femininity, masculinity, oldness, youthfulness, etc.[148]

Let's say that the constancies that single out features of faces are labled C1, C2, C3… C80. Further suppose that each of these can be activated to a certain extent, such that they can become activated less or more strongly. Then we can imagine perception of a particular face as activating each of the constancies C1, C2, C3… C80 to a certain extent. For instance, C44 might respond to femininity and respond at 90% of its activation level, while C12 that responds to smiling might respond at 80% of its activation level, and so on. When Churchland speaks of *indexing* he means the precise way in which several such constancies are combined in a definite perceptual representational content.

Notice that this exactly mirrors the empirical model of how we perceive colors. Edelman and Tononi uses a similar example; they say, "let us now imagine that the mean firing rate of these three sets of neuronal groups can vary from 0 (completely inhibited) to 10 Hz (spontaneous firing range) to 100 Hz (maximum firing)" and further, "that differences in the firing rates of as little as 5 Hz make a difference to the firing of the neurons to which they project and connect" (G. Edelman & Tononi, 2001, p. 161).[149] With these constraints, they go on to note, we have a three dimensional color space.[150] Chuchland makes the same point, that the color space can be seen as a three dimensional color space (P. M. Churchland, 2005).[151]

---

[148] Many of these dimensions will not be recognizable as an attribute the whole individual would use in characterizing the face, but might simply reflect deviances that are not noticeable but collectively add to the power of facial discrimination and "modeling" for the whole subject. Also note here that a face might activate both male and female to a high degree. (In our perceptual system this might be true, or we may think of there being inhibitory connections between the perceptual constancy for femininity and masculinity.)

[149] Note that this simplified model would not account for upwards of a million colors, but only eight thousand or so discriminations.

[150] See (G. Edelman & Tononi, 2001, p. 162). They also note that a model implementing this "can account for our ability to discriminate colors and reproduces well various phenomena in the psychophysics of color

Generally, Churchland's account relies on two factors; a map, and an index. The 'map' part in the above example is the structured space of possible representations. The 'index' part is what determines the actual perceptual representational content in any given instance.

The Map-Indexing Model of Perception makes sense of how the "content" of the constancies become a definite representational content. In particular, it gives an explanation of the way in which different lone constancies combine into a definite, unitary content. In general, it is a more fine-grained attempt at explaining how the activations of the perceptual constancies become a perceptual representational content that might be for the subject as a whole.[152]

4.3.3 Phenomenal Predictions

The MIMP also gives us a way to make *predictions* about the phenomenal character of perceptual experiences. The following example comes from (P. M. Churchland, 2005) where he argues that a Map-Indexing Model of color vision predicts a color that he calls "styan blue," a shade of blue as dark as black and yet distinctly blue. This is a color that makes sense given a specific theory of human color vision, but is typically prevented from occurring due to the workings of our eye. In particular, our eyes cannot register something as distinctly blue without also activating the registration of brightness, so to speak, thus typically excluding the experience of a completely dark blue hue (P. M. Churchland, 2005, p. 532; 2012, pp. 134-135;plate 6). However, Churchland *shows* in his article that it is possible, by inviting the reader stare at various patches of particular colors for an appropriate amount of time, to induce - by fatiguing the eyes in the right way - the phenomenal character "styan blue" and other exotic hues that the theory predicts *should* be possible to experience, given that the theory is correct and phenomenal characters maps on to the color space that the theory postulates.[153]

---

perception, including color constancy" (G. Edelman & Tononi, 2001, p. 161).

[151] Churchland bases this whole article, *Chimerical Colors: Some Predictions from Cognitive Neurscience*, on The Hurvich–Jameson (H–J) opponent-process network; see (P. M. Churchland, 2005, pp. 529-540).

[152] See (DiCarlo et al., 2012) for a state of the art approach. Also see (Nere, Olcese, Balduzzi, & Tononi, 2012) for a neural network inspired computational architecture.

[153] There are other 'exotic' colors too. I urge the interested reader to check it out (P. M. Churchland, 2005, pp. 548-549).

What this shows is that there are some exotic colors (phenomenal characters!) that figure in our color space that we are not even aware of, or worse yet, might consider impossible – such as 'impossibly dark yellow' and 'hyperbolic orange' (P. M. Churchland, 2005, pp. 548, 552). Unless one fatigues the eyes in a particular way while looking at these colors patches as illustrated in the article, one might never phenomenally experience those particular hues. What is special about this "styan blue" after-image is that it does not seem to track anything we might possibly see through our eyes in the normal fashion. It is, as Burge might put it, psychological noise.

However, the theory that predicts it - along with all the other colors we see - is mathematically quite beautiful. It is noteworthy that the quality space is partly constrained by mathematical symmetry. Also noteworthy is the fact that the theory takes into consideration possibilities that we might not even realize are part of our visual system.[154]

4.4 Two Interpretations

Given that we accept The Map-Indexing Model of Perception (MIMP) as being supplementary to Burge's account of perception, then we now have two ways to interpret the constitutive relation to phenomenal consciousness. On a weak interpretation, the phenomenal predictions above only reveal a correlation between phenomenal characters and MIMP. The correlation goes through what we called structural information. Notice that the relation between 'map' and 'indexing' inherently appeals to structural information because information is understood here as the exclusion of possibilities; the map is understood as a structured set of possibilities, and the indexing is understood as the determination of one among many possibilities (i.e. excluding all the other possibilities). As we have seen, various thinkers have conjectured that there is a constitutive link between phenomenal consciousness and structural information; whenever there is phenomenal consciousness there is structural information; whenever there is phenomenal consciousness there is an exclusion of a structured set of possibilities. The weak interpretation is too weak insofar as we want phenomenal consciousness to be constitutively associated with perception, however. For all it

[154] The idea that our perceptual system is constrained by mathematical principles of symmetry might also do some work in differentiating between parochial, pragmatic, evolutionary instincts of the animals on the one side, and its objectively aiming to be accurate in an act of perception on the other. (Initially, I thought about writing my whole thesis on this issue.)

says perception as structured information does not need a constitutive link to phenomenal consciousness, even if the opposite is true.

On a strong interpretation the phenomenal predictions above are taken to reveal something deeper about the connection between phenomenal consciousness and MIMP. In particular, one might conjecture that it reveals phenomenal information to be constitutive to any account like The Map-Indexing Model of Perception insofar as it involves the idea of 'indexing' in a similar fashion and aspires for completeness. The strong interpretation, by taking Tononi's Integrated Information Theory as its cue, solves additional problems that MIMP does not have resources to address. In particular, the nature of the *indexing* remains mysterious and incomplete on MIMP, and the predictions the model yields about phenomenal characters such as "styan blue" remains puzzling and unexplained.

4.4.1 Weak interpretation

Structural information, as we found it in phenomenal experiences, was just the idea that there is a kind of rich structure of possible phenomenal characters relative to which a particular phenomenal character gets its particular nature. On the weak interpretation we have a third-person formulation of the very same thing: There is a kind of rich structure of possible neuronal states relative to which a particular state gets its particular nature. Unlike Burge's RRA principle, which is framed in terms of associations between representational alternatives, the MIMP gives a very precise, quantitative way to determine a perceptual representational content in any given instance, and given a set of representational alternatives. Thus, the MIMP, even on a weak interpretation, helps give us solutions to some of the fundamental problems uncovered in the last chapter.[155]

[155] The four problems were as follows: Max Activation, No Activation, Unitary Activation, Conflicting activation. The Max Activation case highlighted the strange nature of the "content" of the constancies. Taken in isolation, the neurons or clusters of neurons underlying these constancy mechanisms responded more strongly to caricatures and "Picasso monstrosities," than what one might have taken them to typically have been formed in response to. The no activation case highlighted the mysterious problem of how one might begin to understand the perception of absences in terms of the activation of the perceptual constancies. Only by being inactive relative to some context, it seemed, could absences be represented positively. It was not clear how to flesh out what this context was in any given case. The unity activation highlighted the problem of binding together different constancies into a unified content. It also questioned whether the successful

Helps Solve Max Activation and Unitary Activation Case

In the last chapter I outlined four cases meant to show the theoretical difficulty of bridging the gap between the activation of the perceptual constancies and the perceptual content that is presented for an individual. On the weak interpretation, the MIMP begins to make sense of two of these cases.

First, it begins to solve the problem of explaining why the activation of the perceptual constancies only indirectly help type a typical perceptual content for a whole individual. It makes sense of this by noting that the perceptual constancies are formed with respect to serving a global representational content that in turn is for a whole individual. For instance, caricatures reveal that our perceptual constancies respond to exaggerated, impossible or idealized patterns, not because these *necessarily* capture something true about objective features, but because the application of several such perceptual constancies *jointly do* capture something true about objective features. In particular, they jointly capture something true about objective features of interest to the whole individual. Thus, it makes sense of why, and the way in which, some perceptual constancies activate in response to what they do, and why this is indirectly related to a typical representational content had by a whole individual.

Second, the MIMP helps solve the problem of how the activations of different perceptual constancies are combined into a definite perceptual representational content. It is by means of *indexing* that the different activations-levels of different perceptual constancies come together into a definite perceptual content for the individual as a whole. As Churchland notes, "what is important, for any map to be taken seriously as a representation of reality, is that somehow or other, however indirectly, it is possible to *index* it" (P. M. Churchland, 2012, p. 250). Thus, it is by means of some measure of *indexing* that the system intrinsically specifies a definite representational content.

The idea is that we simply look at the activation levels of different neurons and take their individual activation-levels to reflect dimensions in some representational space. This can be

---

development of such a theory might give a more subordinate role to the perceptual constancies. Finally, the un-unitary activation highlighted the problem of lacking an account of exactly how different constancies bind together to form a unitary content. Without such an account, one may worry that multiple representational contents are instantiated in and attributed to a single individual at any given time, such that the various veridicality conditions of the individual as a whole at any given time conflict with each other.

done in the third person by correlating the levels of activations of different neurons to the responses of the animal under experimental conditions (DiCarlo et al., 2012). Thus, one simply operates with the idea that the MIMP framework that we impose from the outside *corresponds* to some sort of representational space that the animal has intrinsically.

As we saw, however, this representational space not only correlated, but could make phenomenal predictions of phenomenal characters like "styan blue." When a lawlike description of some physical system (e.g a biological visual system) not only correlates with the system but is able to give a wide range of counterfactual predictions, one might begin to suspect that it captures something about the very intrinsic nature of the system itself. However, it seems extremely difficult to understand *how* the MIMP can get at the nature of the system itself when the nature of the *indexing* remains entirely mysterious. In particular, the problem of No Activation (e.g. perception of absence) is *more acute* on this model. For if the activation of three clusters of neurons can give rise to a representational space, then necessarily, the absence of activation of one of these neurons corresponds to a distinct place in this representational space – according to the MIMP itself. The *indexing* is always relative to the whole space of activations that the three clusters of neurons are understood in relation to. That the phenomenal predictions of the theory go outside the normal, empirical, evolved function of vision, in predicting "styan blue," only helps to confirm that the mathematical framework of the theory form a complete set of counterfactual predictions about what phenomenal characters would occur given such and such levels of activations. However, the troubling question arises of how phenomenal characters could vary as a function of the counter-factual space of possible neuronal activations.

In particular, the problem occurs when we ask how *indexing* is possible. The reply that the "visual sensation" of colors are "literally identical" with a certain set of specific activation of neurons (P. M. Churchland, 1986, p. 104) is clearly untenable. There is no way that a simple three dimensional abstract space of color representations, given by some set of neurons is responsible for the richness of colors. If that were true, then there would be color spaces everywhere in our brain, as there are undoubtedly many similar clusters of neurons elsewhere in the brain with exactly the same activation profiles. The difference seems to lie in how the neurons take into account their context. As Chalmers notes, "NCCs [Neural Norrelates of Consciousness] are often supposed to be relatively limited systems, such as the inferior temporal cortex or the intralaminar nucleus, but nobody (or almost nobody) holds that if one

excises the entire inferior temporal cortex or intralaminar nucleus and puts it in a jar and puts the system into a relevant state, it will be accompanied by the corresponding state of consciousness" (Chalmers, 2010, p. 73). However, the problem of context is also present with just three neurons side by side. Why just those three neurons, understood in that particular way? Why are the counterfactual set of activations, and the absence of activations, relevant for the occurrent representational state of the individual? How do we individuate the set of possible activations that are relevant to the intrinsic nature of the representational state?

For those that are apt to worry about such problems, I think Tononi's Integrated Information theory supplies a much needed stronger interpretation of the relation between the 'map' and the 'indexing' on the MIMP. Tononi's theory aims to solve such problems. That it also happens to give a formal theory of the nature of phenomenal consciousness can be considered a happy accident.

4.4.2 Strong Interpretation

As we have seen, the most unstable feature of the MIMP is that the explanation of how *indexing* is possible remains mysterious. As we will see, the same problem is evident in the literature on phenomenal consciousness. The Strong Interpretation I advocate takes the problem of *indexing* to be at least partly solved by Tononi's measure of Integrated Information (IIT). Thus, Tononi's theory coupled with the MIMP theory gives us a very solid theory of perception. That Tononi takes Integrated Information to be an intrinsic measure of phenomenal consciousness is of course interesting for all sorts of philosophical reasons, but it also makes phenomenal consciousness constitutive to perception and perceptual representation, on the assumption that we accept MIMP as giving us perceptual representational content in the fundamental case, and IIT as giving us an explanation (in terms of phenomenal consciousness) of how this is possible.

It is unclear on Churchland's map-indexing model of perception how *indexing* is possible. However, this is not an isolated problem. For instance, in an article by Dehaene et al. they hold that a phenomenal experience is possible only if what is experienced is "represented by small groups of neurons whose firing provides an unambiguous index of the relevant attribute, and which would be amplified by top-down attention" (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006, p. 209). As in Churchland's work, the 'unambiguous index' should really be part of what such a theory aims to explain, and should not figure as a primitive in the

explanation itself. By contrast, Edelman and Tononi, after discussing an "accurate neural model of color vision," (similar to the one given by Churchland) wonder how it is possible that clusters of neurons could provide an unambiguous index relating to a *particular* qualitative space (color); they ask "what else is needed?" (G. Edelman & Tononi, 2001, pp. 161-162). Tononi, Edelman and Baars (G. M. Edelman et al., 2011) have great overlap in their respective theories of how to begin to solve this problem, and they all relate it to consciousness. Another instance of the *indexing* problem can be seen in Jesse Prinz's recent work on perception and 'attention' (Jesse J Prinz, 2012). See Appendix A.

Solving the Problem of Indexing

To solve the problem of indexing is simultaneously to solve the problem of No Activation we encountered in section 3.3.2. In contrast to leaving the *indexing* unexplained, Tononi's IIT helps to provide a formalized, mathematical measure for the integration of information. The measure of Integrated Information is essentially a measure of why and how the actual states of a system can exclude, in a specific intrinsic way, some structured set of possible states that the system has available to it. Integrated information both helps us understand how activations of neurons can determine, from the intrinsic perspective of the system, a particular measurable quantity of information, as well as a particular structure of information (Balduzzi & Tononi, 2009). It is outside the scope of this essay to detail the way in which this is possible. I urge the interested reader to check out the theory for more details.

The important point is seeing that it might be possible to explain *indexing* in a quantitative, fundamental way. The major point is that *if* Tononi's theory, or a similar theory, is true, *then* Churchland's model, or a similar model, of perception would fruitfully be related constitutively to this kind of information. It also begins to explain why individuals with perceptual constancies have phenomenal consciousness. Note that Burge explains perceptual constancies as distinct from mere sensory registration partly as requiring that the filtering out of information (or screening) be "structurally specific to an attribute, and cannot be a weighting or averaging of sensory registrations" (Burge, 2010, p. 408, fn.41) and that "these capacities cannot be explained simply as generalized weighting of registration of proximal stimulation" (Burge, 2010, p. 408).[156] However, it is exactly in such mechanisms that Integrated Information flourishes. In particular, the HTM-like structure will need to match the

[156] Also see (Burge, 2010, p. 424).

structure of the scene we are facing on many levels of abstraction that do not reduce to any simple function of the sum of their parts.[157]

While I prefer the stronger interpretation, I do think Tononi's theory is very speculative at this point in time. Even Churchland's theory is quite speculative. However, both of these theories are formulated as answers to problems that I think we should recognize. I also think it is interesting to see the kind of picture that emerges when both of these answers are combined. Due to limitations of space I include in Appendix B some broader historical comments that place these answers in the very human terrain of 20th century answers to the striking feature of mind that is intentionality and representation. Next, I outline two puzzles (historically salient – see Appendix B) that Integrated Information helps solve.

Two Puzzles Integrated Information Helps Solve

First puzzle: Justin Fisher in his (Fisher, 2007) has argued that *any* theory of *phenomenal experience* will face counter-intuitive consequences that suggests solutions similar to the one thinkers like Millikan (Millikan, 2009), Dretske and (Fish, 2013) have proposed in solving the problem of intentionality.[158] The force of Fisher's article stems from deep intuitions we have about the world that jointly cannot be true. On the one hand, it is intuitive that the phenomenal character of green, for instance, cannot be isolated to a single, say, neuron, or a single specific point in time, but must involve some spatio-temporal context. On the other hand, consciousness seems unitary and irreducible. Thus, there is a gap between a process taking some physical time, involving many elements and reaching from the past into the future, and a unitary singular phenomenal character of green. According to Fisher, we are left with accepting 1) wide functionalism, that take into account the context and surroundings of the individual neurons (at an extreme the very environment we are in). Or to accept 2) a kind of history-oriented version of teleo-functionalism that takes into account what has happened recently (at an extreme the evolutionary history of the system). Fisher thinks the history-

[157] This line of reasoning could also question Burge's comment that "the psychological realization of a quantitative map-like geometrical structure could be phylogenetically prior to perception." At least for all but trivial instantiations of such structures, perceptual constancies, as HTM-like, might be required.

[158] Dretske, in particular also uses his theory to explain phenomenal consciousness. See (Bourget & Mendelovici, Forthcoming) for a critical discussion of how Dretske and others use a kind of externaist "tracking" representationalism to explain phenomenal consciousness.

oriented version of teleo-functionalism sounds most plausible, landing him in terrain that is quite familiar to thinkers like Fodor, Dretske, and Millikan.[159]

IIT solves this puzzle since Integrated Information is maximized at a certain spatio-temporal scale, that roughly corresponds to the spatio-temporal scale most researchers in the field (and Burge seems to) expect to find phenomenal consciousness at, namely on the order of 100ms (milliseconds) across a rapidly shifting neuronal (global) workspace (Dehaene & Changeux, 2011, p. 215). The integrated information theory is compatible with some of the *logical structure* of the philosophical arguments that leads theorists like Fisher, Dretske and others (e.g. (Byrne & Tye, 2006)) to accept externalism about phenomenal characters. However, at the same time, it vindicates our intuition that, *in fact*, phenomenal consciousness does not depend on things that are radically far away in spatio-temporal terms. Thus, the IIT effectively helps resolve this tension.

Second puzzle: There is a perennial philosophical idea that phenomenal consciousness has a certain epistemic, and often metaphysical priority, that has been important philosophically, such as in Russells 'aquaintance' (Russell, 1997, pp. 46-59). And yet how can phenomenal consciousness, or anything else for that matter, give us such an epistemic grip?

IIT begins to solve this puzzle by holding that phenomenal consciousness *just is* a type of phenomenal information, namely integrated information; this means that whenever we phenomenally experience something, the very nature of the experience excludes a great range of other possible experiences that we could have had at that very moment. In this sense it gives us an epistemic grip on the nature of the relevant state we are in, as contrasted with several other possible states that we might have been in at that very moment. In addition, IIT vindicates the idea that phenomenal consciousness has a certain metaphysical priority. According to Tononi's theory, Integrated Information is a fundamental property of nature, more primordial than any parochial stance that we can take in giving an explanation of it. In sum, if phenomenal consciousness is essential to intentionality then there might be an important and irreducible sense - contra Dennett and Fodor (see Appendix B) - in which intentionality is fundamental, and not deflatable.

[159] In light of this we can perhaps better see why for instance Ned Block is so careful to assume only "a very weak form of physicalism" where "the phenomenal character of experience supervenes on physical states of the body and … the physical environment surrounding the body and even the past physical states of that environment" (Block, 2007, p. 10, fn.12).

This is a fascinating but limitless and speculative topic. Although there is a lot more to be said, I hope to have provided at least some reasons for thinking that MIMP and IIT (and thus phenomenal consciousness) are constitutively associated with perceptual representation, particularly insofar as it is a striking feature of mind. In the context of exploring a scientifically grounded possibility of Perceptual Duality, I think this is sufficient.

4.4.3 Chapter Conclusion

I have argued that, at least on the map-indexing account of perception, structural information is essential to perception. I also argued that the map-indexing account made sense of some puzzling features uncovered in the last chapter about how the activations of the perceptual constancies could relate to a definite perceptual representational content. In particular, the Map-Indexing Model of Perception takes the activation-levels of perceptual constancies as giving dimensions in a representational space. The *indexing* of the activation-levels of perceptual constancies thus means a determination (of position) in a space of possible representations. I further argued that there is a plausible case to be made for a further supplement to the Map-Indexing Model of Perception.

In particular, the Map-Indexing Model of Perception did not quite make sense of all the puzzling features uncovered in the last chapter. For instance, the failure to explain the problem of absences follows from the theory's lack of explanation of how *indexing* is supposed to be possible. For those who are apt to worry about such things, I made the case that one should look for an additional supplement. I speculated that Tononi's Integrated Information theory is a good candidate. Tononi's theory unites a theory of experience with a determinate, formalized measure of information that could help us to objectively determine the perceptual content of any given individual at any given time.

In the next chapter I answer some objections, and provide a final conclusion.

5 Objections, Replies and Conclusions

5.1 Introduction

In this (short) chapter I discuss a select few objections. I will not discuss Burge's assumption that the nature of perception is representational. I will also not discuss the rejoinder that Churchland's MIMP or Tononi's IIT might be false. I include in Appendix C some more advanced discussion about how supplementing Burge's account with Churchland's MIMP and Tononi's IIT might change the focus in explaining the nature of veridicality conditions, and, in particular, whether there are two differing conceptions of veridicality conditions in place and whether these veridicality conditions are, in some sense, too abstract to explain PD.

5.2 Two Kinds of Perspective?

In relation to supplementing Burge there seems to be two notions of 'perspective' in play. On the one hand, there is the perspective of an individual as being individuated by its biological, evolutionary, ethnological function relative to its environment. On the other hand, there is the perspective of an individual as being individuated by its intrinsic perspective on its environment. In my discussion, I have largely presupposed that Burge wants to make sense of the perspective of the individual in both the first and the second sense.

Thus, one might object that Burge only aims to answer the first. One might even object that they are incompatible. I leave it open as to whether they are incompatible or not. I think it is an interesting question. There are four reasons why I think that Burge should answer the second as well as the first: First, Burge says in his introduction that he wants to explain one of the most *striking* features of mind. If Burge only means perspective in the first sense, the charge of deflation might be relevant. Second, Burge often speaks of explaining perceptual representational content to the finest level of explanatory grain. If Churchland's MIMP or Tononi's IIT are constitutive to perceptual representation, it seems clear that the finest level of grain involves the second sense. Third, Churchland's MIMP has implications to the fine-grained understanding of whole-animal behavior, which was explicitly assumed by Burge to be relevant (P. M. Churchland, 2012, pp. 139-157). Fourth, Burge speaks of perspective in

connection with the difference between perceiving the same particular from this or that spatial perspective, or with this or that specific phenomenal character – implying the second sense.

5.3 Two Kinds of Perceptual Constancies?

In relation to supplementing Burge there seems to be two notions of 'perceptual constancies' in play. On the one hand, there are the perceptual constancies that center on capturing law-like regularities in the environment, such as those underlying at least some of our geometrical, spatial understanding of space. On the other hand, there are the perceptual constancies that center on simply capturing invariant patterns that recur in the environment of the individual, with no special emphasis on them reflecting anything resembling laws in the physical world.

In my discussions I have presupposed that the latter understanding of perceptual constancies is more general. There are several reasons for this: First, from what I know of perceptual psychology, I do not think we understand yet how to restrict and explanation of perceptual constancies to apply only to the second conception. Second, this essay was concerned primarily with the relation between perceptual constancies and the individual's perspective as a whole; many of the paradigmatic examples of perceptual constancies capturing law-like regularities, like convergence, are simple mathematical constancies that primarily figure in early stages of visual processing. Like the simple primitives in the visual system they only indirectly type the perceptual representational content of the animal as a whole.[160] Third, according to Burge whole-individual behavior is usually centered on perceptually attributing things such *flies* and *body* that while parasitic on primitive perceptual constancies, involve the particular nature of and biological concerns of the whole animal to a much higher degree.[161]

5.4 Phenomenal Consciousness – Part/Whole?

One might suppose that phenomenal consciousness does not need to be for the whole individual. If this is true, then it undercuts using phenomenal consciousness as glue for the perceptual representational constancies in the subsystem of the whole individual. Consider people who have had their corpus callosum cut for instance to prevent seizures in one

[160] That is, it only makes sense to call them 'representations' by appealing to their constitutive association with the fundamental cases of perceptual representational content that are for a whole individual.

[161] Personally, I would use the word 'fundementally' instead of 'usually' here. I do not because Burge might not.

hemisphere from affecting the whole brain; if their phenomenal consciousness is also split in two, one might still want to hold that they are a single whole individual. In that case, the presence of phenomenal consciousness, in conjunction with perceptual constancies, is not a sufficient condition for the whole individual to have perceptual content.[162]

In such cases I would hold that there are actually two perceptual representational contents, attributed to two whole individuals. There are two whole perspectives. As one would expect, these people often express conflicting interests, points of view, motivations, desires and behavioral expressions. However, it is also clear that many functions of such individuals are defective, in the sense that the very nature of many of these functions are constitutively associated with the other half's body, brain, memory, environment, etc. More subtly, one might argue that phenomenal consciousness can be isolated in different parts of the brain. Or one might argue that phenomenal consciousness is a gradual phenomenon. Against this, a consequence of Tononi's IIT theory is that phenomenal consciousness only exists in spatio-temporal pockets within which the would-be conscious parts melt together like droplets of water forming a biggest droplet. According to IIT, it is at least unlikely that there will be distinct pockets of consciousness in a brain.

Integrated information clearly achieves a kind of global unity among different subsystems in the brain. However, although many theories of consciousness postulate that consciousness has something to do with global availability of information, the science has not yet progressed far enough to answer exactly *how* consciousness might be *instrumental*, as opposed to for instance epiphenomenal, in such global access.[163] We do not yet fully understand how phenomenal consciousness can be instrumental, for instance, in suppressing non-dominant stimuli. Nevertheless, progress is being made, the future looks bright, and I am optimistic.

[162] I do not actually need to hold that they are sufficient, only that they are necessary. Nevertheless I should say that in a rough and noncritical way they are necessary and sufficient, for this example clearly goes against the spirit of my approach (to use phenomenal consciousness as the criteria that makes perception *for the whole* individual.)

[163] I agree with Burge epiphenomenalism should be regarded as a skeptical worry only to be taken seriously as an abstract possibility (Burge, 1993). A position on which phenomenal consciousness does not take *part in* the causal structure of the world is *surely* a skeptical position only to be taken seriously as an abstract possibility.

5.5 Final Conclusion

5.5.1 How I discussed it

The overarching goal was to find out how Perceptual Duality might be possible. In particular I wanted to understand the idea that perceptual content is constitutively for a whole subject and that perceptual content is constitutively about objects. I assumed that perceptual representational content was central to the explanation of the nature of perception, and I used Burge's framework and account of perception to give substance to the idea that perceptual content could partly be explained in terms of perceptual constancies. In particular, the idea was that perceptual constancies are constitutively about objects. In agreeing with Burge on these points, I went on to investigate how perceptual content, partly explained in terms of perceptual constancies, could be attributed to the subject as a whole. Anticipating my positive account, I investigated how Burge saw the relation between phenomenal consciousness and perceptual content. As Burge did not think phenomenal consciousness was even constitutively associated with perceptual content, I became interested in his reasons and worried about how to explain the subject's perspective in perceiving. My worry developed into growing sense that Burge's account did not have sufficient resources to make sense of exactly how, in scientifically grounded way, the perceptual constancies become a perceptual representational content for the whole subject.

In outlining four fundamental problems of how to connect the perceptual constancies to a subject as a whole (full-, no-, unitary- and non-unitary activation) I found a need to supplement Burge. With phenomenal consciousness I wanted to give a positive explanation of how one might explain the connection between the activation of the perceptual constancies and the perspective of the subject as a whole.

I drew on a wide range of thinkers to show that phenomenal character was associated with a kind of structural information. Using the notion of structural information I showed that Paul Churchland's MIMP could relate perceptual constancies to a definite perceptual representational content. In addition, I noted that the definite representational content we get from MIMP, which is based on perceptual constancies, actually has phenomenal predictions.

More speculatively, I argued that phenomenal experience might be constitutively associated with Churchland's MIMP. In particular, I conjectured that Tononi's integrated information

theory is constitutively associated to a theory of perception along the lines that Paul Churchland has given. In particular, Tononi's theory is unique in having, or at least attempting to develop, an intrinsic measure of information integration. Even if one does not accept Tononi's measure, one can recognize that such a measure, that is both intrinsic to the system itself and can determine a subjective perspective, is needed to *index* perceptual constancies in the way Paul Churchland's theory proposes, and to make sense of the predictions that the map-indexing model of perception makes about phenomenal experience, such as in case of an "impossibly dark blue."

If we assume that Burge has given a good account of how the perceptual constancies are constitutively *about* objects then the condition under which Perceptual Duality is possible is given by the precise way in which the *whole subject* perceives by making use of such perceptual constancies. Paul Churchland has given a decent account of how it is possible that perceptual constancies can give rise to definite perceptual representational content. If we supplement Churchland's MIMP with Tononi's IIT, we get a deeper explanation of how *indexing* is possible on Churchland's account. This leads to an even stronger account of how Perceptual Duality is possible. And since IIT is a theory of phenomenal consciousness, this stronger account of perception has phenomenal consciousness as a constitutive condition.

5.5.2 Why it is important

A big part of this essay consists of explicating Burge's views on the relation between perceptual constancies, perceptual representation and phenomenal consciousness. To understand Burge's views on these matters is interesting in and of itself. Burge gives both perceptual constancies and phenomenal consciousness important positions in his philosophical framework, and is tireless in seeking their scientific ground. In exploring the relation between these two aspects of his philosophical framework, I hope to have shed some light on Burge's position on these matters.

My outline of the four problems in chapter 3 highlights, in a concrete way, some abstract problems that need to be resolved in understanding how components in the subsystem of an individual can give rise to a perspective of an individual as a whole. I think the resolution of these problems is not specific to a debate centering on perceptual constancies and phenomenal consciousness. Perceptual constancies and phenomenal consciousness, however, turn out to be pieces can instantiate a resolution to this puzzle in a concrete sense.

In addition to Burge exegesis and the four problems, one might say that my positive account consists of clashing together Burge's philosophical account of perception with Churchland's account of the map-indexing theory of perception, and Tononi's Integrated Information theory of consciousness. There is some truth to this; however, I think this combination is highly interesting, particularly in the sense that it opens avenues for further investigation.

My positive account makes it clear the idea that there are three theories out there. Burge has a scientifically grounded philosophical account of how perceptual constancies are about objects and features of objects in the world. Churchland's theory is about how perceptual constancies become a definite perceptual representational content. And Tononi's theory is about how phenomenal experience, with an intrinsic perspective, *just is* a complex of integrated information. Thus, in part I think this essay shows that they are at three different levels that need to come together to solve the *philosophical* problem of how a *subject*, in the most fine-grained sense, can perceive objects in the world. My thesis also implicitly shows that these three levels can be, and actually are, separated, and carried out by different people with different agendas and motivations.[164]

While in fact all of Burge's, Churchland's and Tononi's theories might be wrong in detail, they still illustrate the possibility of Perceptual Duality. In addition, they suggest a research program. In particular, they each try to explain phenomena that taken together seems to be at least one solution to the problem of Perceptual Duality. Let me take each project in turn.

Burge tries to explain the origins of representation. Fundamentally, Burge wants to explain what the *minimum* conditions for objective perceptual representation are. One can give a very good and nuanced account of the minimum conditions of objective representation without saying much about how different parts of the brain coordinate with each other and relate to a subject *as a whole*. Understanding the nature of perceptual constancies can in part be seen as a computational problem, or a problem of generalized pattern recognition. In part, Burge's talk of perception being for the whole individual is motivated by anti-individualism.[165]

[164] On top of that we may consider philosophical projects such as epistemology, metaphysics, and philosophy of mind as independent projects in their own right.

[165] This is of course because he needs that it is the *individual as a whole* that perceives, insofar as he wants particular perceptual constancies to have the added specificity that they have, namely, not to indicate a certain spatio-temporal pattern, but to indicate exactly what is biologically important to *the individual as a whole*.

Churchland tries to give a conceptual framework for understanding how different parts of the brain give rise to a perceptual representational content; one can explain this in quite some detail without either talking about how this can give rise to the phenomenal character of experience, or talk about how this content relates to objective things in the world. The problem is one of giving a unified account of how the various perceptual constancies can be transformed into a definite, precise perceptual representational content on any given occasion. However, it is worth noting that the specificity inherent in Churchland's account also gives a fine-grained understanding of how an animal can initiate and execute coherent behavior (P. M. Churchland, 2012, pp. 45-50).[166] And, as we have seen, it might also contribute or help determine the fine-grained accuracy-conditions on perceptual content.

Tononi is motivated by asking how, in a fundamental sense; information in different parts of the brain can be integrated into a unified perspective. As it happens, and I think amazingly, he associates this measure with phenomenal consciousness, and formulates mathematically precise, scientific, predictions based on it. Obviously, one can explain much about how information in different parts of the brain can be integrated into a unified perspective without talking about how the origins of representational capacities is possible, or how this unified perspective relates to the external world. However, as we have seen, IIT might also contribute to a very fine-grained understanding of how the perceptual representational content 'matches' the environment in a dynamical, mathematically precise way.

Another reason why I think the essay might be of interest is because it outlines a possible relation between phenomenal consciousness and perception. Aside from Searle mentioned earlier, there is a recent resurgence of philosophers who claim that intentionality is constitutively conscious. For instance, (Horgan & Tienson, 2002; Loar, 2003), and more recently, (Farkas, 2010; Kriegel, 2007; Masrour, 2013; Pautz, 2013).[167] If we regard perception as essentially intentional, and intentionality as constitutively conscious, then it is

---

This also counts in favor of regarding the early stages of vision as individuated representationally only with respect to understanding what they are used in representing at later stages of vision.

[166] Churchland intimately links the solution to the problem of sensorimotor coordination and perceptual representation (P. M. Churchland, 1986, pp. 280-291). However, if (Morsella, 2005) is right, then this explanation too must take into account phenomenal consciousness as a crucial feature in global whole-individual coordination of action.

[167] See (Siewert, 2006a) and (Menary, 2009) for overview, and (Kriegel, 2013) for a collection of papers.

highly relevant to understand whether, and to what extent, perception too is constitutively conscious.

This essay provides a way to understand the possibility of how one might find scientifically grounded reasons for thinking that perception is constitutively conscious. In particular, it gives us scientifically grounded reasons both to suppose that phenomenal consciousness has an important role to play both in determining the various "representations" floating around in the subsystems of a subject (indexing), and to suppose that phenomenal consciousness has a crucial role in integrating information that is available to the epistemic perspective of a subject as a whole (it *just is* a complex of integrated information). It is also important to note that this way of understanding follows from the effort in this essay to understand the relation between perceptual constancies and phenomenal consciousness, a problem that is at the forefront of scientifically grounded approaches to the intersection between philosophy of mind and philosophy of perception.[168] The further investigation of the relation between perceptual constancies and phenomenal consciousness is likely to remain fruitful.

Thus, given that one is motivated to understand the possibility of Perceptual Duality, with the relevant assumptions, I think the foregoing should be of interest.

However, in conceding the force of some of the objections above (and in Appendix C below), and in recognizing the limitations of the discussion on many points, it is clear that much work still needs to be done to see how (PD) is possible. In particular, I think two things stand out. First, I believe that - under philosophical analysis - 'perceptual constancies' as I have used it, and as I suspect most perceptual psychologists use it, will be seen to consist of at least a few distinct notions. Second, it needs to be investigated how to unite the difference between Burge's notion of 'a whole individual,' as figuring in anti-individualistic individuations of a state, and a notion of 'a whole subject,' as figuring in typing the perspective and epistemic situation intrinsic to the subject at the finest level of grain.[169]

[168] See (Hatfield & Allred, 2012) for a recent collection of papers on visual experience and constancies.

[169] I note here a third point that follows from Appendix C. We should investigate to what extent the correspondence between the perceptual representation and the environment consists of the workings of independent components (i.e. isolated perceptual constancies, simple computations) and to what extent it is an achievement of a collective, complex set of dynamic interactions in the brain. In particular, it is crucial to establish whether there can be a quantitative measure of the extent to which and how the internal structure of

a perceptual representation can mirror (say, be functionally isomorphic to) abstract properties of the objective world, for instance in an embodied, dynamic relation to the environment.

Bibliography

Andrade, J. (2012). Consciousness. In A. G. Nick Braisby (Ed.), *Cognitive Psychology* (Vol. Second Edition, pp. 577-606). Oxford: Oxford University Press.

Balduzzi, D., & Tononi, G. (2009). Qualia: The Geometry of Integrated Information. *PLoS Comput Biol, 5*(8), e1000462. doi: 10.1371/journal.pcbi.1000462

Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron, 76*(4), 695-711. doi: 10.1016/j.neuron.2012.10.038

Block, N. (1990a). Inverted earth. *Philosophical Perspectives, 4*, 53-79.

Block, N. (1990b). Inverted Earth. *Philosophical Perspectives, 4*(ArticleType: research-article / Issue Title: Action Theory and Philosophy of Mind / Full publication date: 1990 / Copyright © 1990 Ridgeview Publishing Company), 53-79.

Block, N. (2007). Wittgenstein and Qualia. *Philosophical Perspectives, 21*(1), 73-115.

Borge, S. (2007). Some remarks on Reid on primary and secondary qualities. *Acta Analytica, 22*(1), 74-84.

Bourget, D., & Chalmers, D. J. (Forthcoming). What do philosophers believe? *Philosophical Studies*.

Bourget, D., & Mendelovici, A. (Forthcoming). Tracking Representationalism. In A. Bailey (Ed.), *Philosophy of Mind: The Key Thinkers*: Continuum.

Bowers, J. S. (2009). On the biological plausibility of grandmother cells: Implications for neural network theories in psychology and neuroscience. *Psychological Review, 116*(1), 220-251. doi: 10.1037/a0014462

Broughton, L. M. (1981). Quine's 'quality space'. *Dialectica, 35*, 291-302.

Burge, T. (1986). Individualism and psychology. *Philosophical Review, 95*(January), 3-45.

Burge, T. (1993). Mind-body causation and explanatory practice. In J. Heil & A. R. Mele (Eds.), *Mental Causation*: Oxford University Press.

Burge, T. (2003a). Descartes, Bare Concepts, and Anti-Individualism: Reply to Normore. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 291--335): Mit Press.

Burge, T. (2003b). Perception. *International Journal of Psychoanalysis, 84*(1), 157-167.

Burge, T. (2003c). Phenomenality and Reference: Reply to Loar. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 435--451): Mit Press.

Burge, T. (2003d). Qualia and intentional content: Reply to Block. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 405--415): Mit Press.

Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics, 33*(1), 1-78.

Burge, T. (2007). Reflections on Two Kinds of ConsciousnessFoundations of MInd (Vol. 2, pp. 392-420): Oxford University Press.

Burge, T. (2009a). Five Theses on De Re States and Attitudes. In J. Almog & P. Leonardi (Eds.), *The Philosophy of David Kaplan*: Oxford University Press.

Burge, T. (2009b). Perceptual objectivity. *Philosophical Review, 118*(3), 285-324.

Burge, T. (2010). *Origins Of Objectivity*. Oxford New York: Oxford University Press.

Byrne, A., & Tye, M. (2006). Qualia ain't in the head. *Noûs, 40*(2), 241-255.

Baars, B. J. (2002). The conscious access hypothesis: Origins and recent evidence. *Trends in Cognitive Sciences, 6*(1), 47-52.

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends Cogn Sci, 7*(4), 166-172.
Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*: Oxford University Press.
Chalmers, D. J. (2010). *The Character of Consciousness*: Oxford University Press.
Chomsky, N. (1995). *The Minimalist Program*: The Mit Press.
Churchland, P. (2007). On the reality (and diversity) of objective colors: How color-qualia space is a map of reflectance-profile space. *Philosophy of Science, 74*(2), 119-149.
Churchland, P. M. (1986). Some reductive strategies in cognitive neurobiology. *Mind, 95*(July), 279-309.
Churchland, P. M. (1998). Conceptual similarity across sensory and neural diversity: The fodor/lepore challenge answered. *Journal of Philosophy, 95*(1), 5-32.
Churchland, P. M. (2005). Chimerical colors: Some phenomenological predictions from cognitive neuroscience. *Philosophical Psychology, 18*(5), 527-560.
Churchland, P. M. (2012). *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, Massachusetts

London, England: The MIT Press.
Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*: Oxford University Press.
Cottrell, G. W., & Metcalfe, J. (1990). *EMPATH: face, emotion, and gender recognition using holons*. Paper presented at the Proceedings of the 1990 conference on Advances in neural information processing systems 3, Denver, Colorado, United States.
Cummins, R. (1996). *Representations, Targets, and Attitudes*: Mit Press.
Cummins, R. C., & Poirier, P. (2004). Representation and indication. In H. Clapin (Ed.), *Representation in Mind*: Elsevier.
Dainton, B. (2010). Temporal Consciousness.
Davidson, D. (1987). Knowing One's Own Mind. *Proceedings and Addresses of the American Philosophical Association, 60*(3), 441-458. doi: 10.2307/3131782
Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron, 70*(2), 200-227.
Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences, 10*(5), 204-211. doi: 10.1016/j.tics.2006.03.007
Dennett, D. Daniel Dennett's Philosophical Lexicon.
Dennett, D. C. (1987). *The Intentional Stance*: MIT Press.
Dennett, D. C. (2001). Are we explaining consciousness yet? *Cognition, 79*(1), 221-237.
DiCarlo, James J., Zoccolan, D., & Rust, Nicole C. (2012). How Does the Brain Solve Visual Object Recognition? *Neuron, 73*(3), 415-434.
Dretske, F. (1995a). Dretske's awful answer. *Philosophia, 24*(3-4), 459-464.
Dretske, F. (1995b). *Naturalizing the Mind*: MIT Press.
Dretske, F. (1998). An Interview with Fred Dretske. *The Dualist (Stanford's Undergraduate Journal of Philosophy)*. Retrieved from http://philosophy.stanford.edu/apps/stanfordphilosophy/files/wysiwyg_images/dretske.pdf
Dretske, F. (2003). Burge on mentalistic explanations, or why I am still epiphobic. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge*: Mit Press.
Dretske, F. (2010). What we see : the texture of conscious experience. In B. Nanay (Ed.), *Perceiving the World*: Oxford University Press.

Edelman, G., & Tononi, G. (2001). *A Universe of Consciousness: How Matter Becomes Imagination*: Basic Books.
Edelman, G. M., Gally, J. A., & Baars, B. J. (2011). Biology of Consciousness. [Perspective]. *Frontiers in Psychology, 2*. doi: 10.3389/fpsyg.2011.00004
Edgar, G. P. a. G. (2005). Perception. In N. B. a. A. Gellatly (Ed.), *Cognitive Psychology* (pp. 71-112). Oxford: Oxford University Press.
Fang, F., & He, S. (2005). Cortical responses to invisible objects in the human dorsal and ventral pathways. *Nat Neurosci, 8*(10), 1380-1385. doi: 10.1038/nn1537
Farkas, K. (2010). Independent intentional objects. In T. Czarnecki, K. Kijanija-Placek, O. Poller & J. Wolenski (Eds.), *The Analytical Way*: College Publications.
Fish, W. (2010). *Philosophy of Perception: A Contemporary Introduction*: Routledge.
Fish, W. (2013). Perception, hallucination, and illusion: reply to my critics. *Philosophical Studies, 163*(1), 57-66. doi: 10.1007/s11098-012-0072-8
Fisher, J. C. (2007). Why nothing mental is just in the head. *Noûs, 41*(2), 318-334.
Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*: MIT Press.
Fodor, J. A. (1990). *A Theory of Content and Other Essays*: MIT Press.
Fodor, J. A., & LePore, E. (1993). Precis of holism: A shopper's guide. *Philosophy and Phenomenological Research, 53*(3), 637-682.
Francescotti, R. (1991). Externalism and Marr's theory of vision. *British Journal for the Philosophy of Science, 42*(June), 227-238.
Friston, K. (2010). The free-energy principle: a unified brain theory? [10.1038/nrn2787]. *Nat Rev Neurosci, 11*(2), 127-138. doi: http://www.nature.com/nrn/journal/v11/n2/suppinfo/nrn2787_S1.html
Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics, 36*(4), 193-202.
Fukushima, K. (1988). Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks, 1*(2), 119-130. doi: http://dx.doi.org/10.1016/0893-6080(88)90014-7
Gallagher, S. (2005). *How the Body Shapes the Mind*: Oxford: Clarendon Press.
Ganson, T., Bronner, B. E. N., & Kerr, A. (2012). Burge's Defense of Perceptual Content. *Philosophy and Phenomenological Research*, no-no. doi: 10.1111/j.1933-1592.2012.00632.x
Geisler, W. S. (2003). Ideal observer analysis. *The visual neurosciences*, 825-837.
George, D., & Hawkins, J. (2009). Towards a Mathematical Theory of Cortical Micro-circuits. *PLoS Comput Biol, 5*(10), e1000532. doi: 10.1371/journal.pcbi.1000532
Goldman-Rakic, P. S., & Rakic, P. (1991). Preface: Cerebral Cortex Has Come of Age. *Cerebral Cortex, 1*(1), 1. doi: 10.1093/cercor/1.1.1-a
Goldstein, E. B. (2010). *Sensation and Perception*: Wadsworth Cengage Learning.
Gottfried, J. A. (2010). Central mechanisms of odour object perception. *Nature Reviews Neuroscience, 11*(9), 628-641.
Gross, C. G. (2002). Genealogy of the "Grandmother Cell". *The Neuroscientist, 8*(5), 512-518. doi: 10.1177/107385802237175
Guigon, G. (Forthcoming). Overall Similarity, Natural Properties, and Paraphrases. *Philosophical Studies*.
Haggard, P., & Clark, S. (2003). Intentional action: Conscious experience and neural prediction. *Consciousness and Cognition, 12*(4), 695-707.
Hatfield, G., & Allred, S. (2012). *Visual Experience:Sensation, Cognition, and Constancy*: OUP Oxford.

Hawkins, J. (2004). *On Intelligence* (S. Blakeslee, Trans.). New York: Owl Books.
Hawkins, J., George, D., & Niemasik, J. (2009). Sequence memory for prediction, inference and behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences, 364*(1521), 1203-1209. doi: 10.1098/rstb.2008.0322
Hawrylycz, M. J., Lein, E. S., Guillozet-Bongaarts, A. L., Shen, E. H., Ng, L., Miller, J. A., . . . Jones, A. R. (2012). An anatomically comprehensive atlas of the adult human brain transcriptome. [10.1038/nature11405]. *Nature, 489*(7416), 391-399. doi: http://www.nature.com/nature/journal/v489/n7416/abs/nature11405.html#supplementary-information
Hintikka, J. (2007). *Socratic Epistemology: Explorations of Knowledge-Seeking by Questioning*: Cambridge University Press.
Hitch, G. J. (2005). Working Memory. In N. B. a. A. Gellatly (Ed.), Cognitive Psychology (pp. 307-341). Oxford: Oxford Press.
Horgan, T. E., & Tienson, J. L. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings* (pp. 520--533): Oxford University Press.
Hung, C. P., Kreiman, G., Poggio, T., & DiCarlo, J. J. (2005). Fast Readout of Object Identity from Macaque Inferior Temporal Cortex. *Science, 310*(5749), 863-866. doi: 10.1126/science.1117593
Kelly, S. D. (2008). Content and constancy: Phenomenology, psychology, and the content of perception. *Philosophy and Phenomenological Research, 76*(3), 682–690.
Kriegel, U. (2007). Intentional inexistence and phenomenal intentionality. *Philosophical Perspectives, 21*(1), 307-340.
Kriegel, U. (2013). *Phenomenal Intentionality*: Oxford University Press.
Kroustallis, B. (2006). Content individuation in Marr's theory of vision. *Journal of Mind and Behavior, 27*(1), 57-71.
Kurzweil, R. (2012). *How to Create a Mind: The Secret of Human Thought Revealed*: Penguin Group US.
LePore, E. (1991). *John Searle and His Critics*: Cambridge: Blackwell.
Lewis, D. (1980). Veridical hallucination and prosthetic vision. *Australasian Journal of Philosophy, 58*(September), 239-249.
Lewis, D. (1990). What experience teaches. In W. G. Lycan (Ed.), *Mind and Cognition* (Vol. 13, pp. 29--57): Blackwell.
Llinás, R. R. (2002). *I of the Vortex: From Neurons to Self*. Cambridge, Massachusetts

London, England: The MIT press.
Loar, B. (2003). Phenomenal intentionality as the basis of mental content. In M. Hahn & B. Ramberg (Eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 229--258): Mit Press.
Lycan, W. G. (1995). A limited defense of phenomenal information. In T. Metzinger (Ed.), *Conscious Experience*: Imprint Academic.
Lycan, W. G. (1998). Phenomenal information again: It is both real and intrinsically perspectival. *Philosophical Psychology, 11*(2), 239-242.
Laakso, A., & Cottrell, G. (2000). Content and cluster analysis: assessing representational similarity in neural systems. *Philosophical Psychology, 13*(1), 47-76.
Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Fransisco: Freeman and Company.
Marr, D. (1983). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*: Henry Holt and Company.

Marshall, R. (Producer). (2013). The Possible Worlds Hedgehog. Retrieved from http://www.3ammagazine.com/3am/the-possible-worlds-hedgehog/
Masrour, F. (Ed.). (2013). *Phenomenal Objectivity and Phenomenal Intentionality: In Defense of a Kantian Account*.
McDowell, J. (2011). Tyler Burge on disjunctivism. *Philosophical Explorations, 13*(3), 243-255.
Mechelli, A., Price, C. J., Friston, K. J., & Ishai, A. (2004). Where Bottom-up Meets Top-down: Neuronal Interactions during Perception and Imagery. *Cerebral Cortex, 14*(11), 1256-1265. doi: 10.1093/cercor/bhh087
Menary, R. (2009). Intentionality and Consciousness. In W. Banks (Ed.), *Encyclopaedia of Consciousness*: Elsevier.
Menzel, R. (2009). *Working memory in bees: also in flies?* (Vol. 23).
Menzel, R., Brandt, R., Gumbert, A., Komischke, B., & Kunze, J. (2000). Two spatial memories for honeybee navigation. *Proceedings of the Royal Society of London. Series B: Biological Sciences, 267*(1447), 961-968.
Menzel, R., Fuchs, J., Kirbach, A., Lehmann, K., & Greggers, U. (2012). Navigation and Communication in Honey Bees

Honeybee Neurobiology and Behavior. In C. G. Galizia, D. Eisenhardt & M. Giurfa (Eds.), (pp. 103-116): Springer Netherlands.
Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., . . . Watzl, S. (2005). Honey bees navigate according to a map-like spatial memory. *Proc Natl Acad Sci U S A, 102*(8), 3040-3045. doi: 10.1073/pnas.0408550102
Millikan, R. G. (2009). Biosemantics. In B. P. McLaughlin & A. Beckerman (Eds.), *The Oxford Handbook of Philosophy of Mind* (Vol. 86, pp. 281--297): Oxford University Press.
Morreau, M. (2010). It simply does not add up: Trouble with overall similarity. *Journal of Philosophy, 107*(9), 469-490.
Morsella, E. (2005). The function of phenomenal states: supramodular interaction theory. *Psychol Rev, 112*(4), 1000-1021. doi: 10.1037/0033-295x.112.4.1000
Mountcastle, V. B. (2003). Introduction. *Cerebral Cortex, 13*(1), 2-4. doi: 10.1093/cercor/13.1.2
Nere, A., Olcese, U., Balduzzi, D., & Tononi, G. (2012). A Neuromorphic Architecture for Object Recognition and Motion Anticipation Using Burst-STDP. *PLoS ONE, 7*(5), e36958. doi: 10.1371/journal.pone.0036958
Neuser, K., Triphan, T., Mronz, M., Poeck, B., & Strauss, R. (2008). Analysis of a spatial orientation memory in Drosophila. [10.1038/nature07003]. *Nature, 453*(7199), 1244-1247. doi: http://www.nature.com/nature/journal/v453/n7199/suppinfo/nature07003_S1.html
Noë, A. (2006). Experience without the head. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual Experience*: Oxford University Press.
Pack, C. C., Berezovskii, V. K., & Born, R. T. (2001). Dynamic properties of neurons in cortical area MT in alert and anaesthetized macaque monkeys. *Nature, 414*(6866), 905-908. doi: 10.1038/414905a
Palmisano, S., Gillam, B., Govan, D. G., Allison, R. S., & Harris, J. M. (2010). Stereoscopic perception of real depths at large distances. *Journal of Vision, 10*(6). doi: 10.1167/10.6.19
Papineau, D. (2013). *Can We Really See a Million Colours?* Lecture. Retrieved from http://www.kcl.ac.uk/artshums/depts/philosophy/people/staff/academic/papineau/files/articles/colours.pdf

Pautz, A. (2013). Is Phenomenology the Ground of Intentionality? In U. Kriegel (Ed.), *Phenomenal Intentionality*: Oxford.
Perry, J. (2001). *Knowledge, Possibility, and Consciousness*: MIT Press.
Piccinini, G. (2004). Functionalism, computationalism, and mental contents. *Canadian Journal of Philosophy, 34*(3), 375-410.
Plaut, D. C., & McClelland, J. L. (2010). Locating object knowledge in the brain: Comment on Bowers's (2009) attempt to revive the grandmother cell hypothesis. *Psychological Review, 117*(1), 284-288. doi: 10.1037/a0017101
Poggio, T. (1981). Marr's computational approach to vision. *Trends in neurosciences, 4*, 258-262.
Prinz, J. (2007). Mental pointing: Phenomenal knowledge without concepts. *Journal of Consciousness Studies, 14*(9-10), 184-211.
Prinz, J. (2011). Is Attention Necessary and Sufficient for Consciousness? In C. Mole, D. Smithies & W. Wu (Eds.), *Attention: Philosophical and Psychological Essays*: Oxford University Press.
Prinz, J. J. (2010). When is perception conscious? In B. Nanay (Ed.), *Perceiving the World*: Oxford University Press.
Prinz, J. J. (2012). *The conscious brain*: Oxford University Press, USA.
Putnam, H. (1991). *Representation and Reality*: A Bradford Book.
Quine, W. V. O. (1960). *Word & Object*: The Mit Press.
Raizada, R. D. S., & Grossberg, S. (2003). Towards a Theory of the Laminar Architecture of Cerebral Cortex: Computational Clues from the Visual System. *Cerebral Cortex, 13*(1), 100-113. doi: 10.1093/cercor/13.1.100
Ramachandran, V. S. (2004). *A Brief Tour of Human Consciousness: From Impostor Poodles to Purple Numbers*: Pearson Professional.
Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex.
Russell, B. (1997). *The problems of philosophy*: Oxford University Press, Incorporated.
Schellenberg, S. (2011). Perceptual Content Defended. *Noûs, 45*(4), 714-750. doi: 10.1111/j.1468-0068.2010.00791.x
Schendan, H. E., & Stern, C. E. (2008). Where Vision Meets Memory: Prefrontal–Posterior Networks for Visual Object Constancy during Categorization and Recognition. *Cerebral Cortex, 18*(7), 1695-1711. doi: 10.1093/cercor/bhm197
Schmidhuber, J. (2012). Computer Vision Retrieved 23/9, 2012, from http://www.idsia.ch/~juergen/vision.html
Schrödinger, E. (1992). *What Is Life?: with "Mind and Matter" and "Autobiographical Sketches"*: Cambridge University Press.
Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*: Cambridge University Press.
Searle, J. R. (1989). Consciousness, unconsciousness, and intentionality. *Philosophical Topics, 17*(1), 193-209.
Serre, T., & Riesenhuber, M. (2004). Realistic modeling of simple and complex cell tuning in the HMAX model, and implications for invariant object recognition in cortex: DTIC Document.
Shagrir, O. (2001). Content, computation and externalism. *Mind, 110*(438), 369-400.
Shagrir, O. (2010). Marr on computational-level theories. *Philosophy of Science, 77*(4), 477-500.
Shapiro, L. A. (1993). Content, kinds, and individualism in Marr's theory of vision. *Philosophical Review, 102*(4), 489-513.
Shoemaker, S. (1982). The inverted spectrum. *Journal of Philosophy, 79*(July), 357-381.
Shoemaker, S. (1994). Phenomenal character. *Noûs, 28*(1), 21-38.

Shoemaker, S. (1996). Colors, Subjective Reactions, and Qualia. *Philosophical Issues, 7*(ArticleType: research-article / Issue Title: Perception / Full publication date: 1996 / Copyright © 1996 Dr. Enrique Villanueva), 55-66. doi: 10.2307/1522891
Siegel, S. (2010). *The Contents of Visual Experience*: Oxford.
Siegel, S. (2011). The Contents of Perception. *Standford Encyclopedia of Philosophy*. Retrieved from http://plato.stanford.edu/archives/win2011/entries/perception-contents
Siewert, C. (2006a). Consciousness and Intentionality. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*: Stanford: Metaphysics Research Lab.
Siewert, C. (2006b). Is the appearance of shape protean? *Psyche, 12*(3), 1-16.
Siewert, C. (Forthcoming). On the Phenomenology of Introspection. *Introspection and Consciousness*.
Sporns, O. (2010). *Networks of the Brain*: MIT Press.
Stalnaker, R. (2008). *Our Knowledge of the Internal World*: Oxford University Press.
Tang, S., Wolf, R., Xu, S., & Heisenberg, M. (2004). Visual Pattern Recognition in Drosophila Is Invariant for Retinal Position. *Science, 305*(5686), 1020-1022. doi: 10.1126/science.1099839
Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science, 331*(6022), 1279-1285. doi: 10.1126/science.1192788
Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends Cogn Sci, 10*(11), 502-511. doi: 10.1016/j.tics.2006.09.003
Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience, 5*(1), 42.
Tononi, G. (2008). Consciousness as Integrated Information: a Provisional Manifesto. *The Biological Bulletin, 215*(3), 216-242.
Tononi, G. (2012). *Phi: A Voyage from the Brain to the Soul*: Knopf Doubleday Publishing Group.
Tononi, G., & Edelman, G. M. (1998). Consciousness and Complexity. *Science, 282*(5395), 1846-1851. doi: 10.1126/science.282.5395.1846
Tononi, G., & Sporns, O. (2003). Measuring information integration. *BMC Neuroscience, 4*(1), 31.
Tononi, G., Sporns, O., & Edelman, G. M. (1996). A complexity measure for selective matching of signals by the brain. *Proc Natl Acad Sci U S A, 93*(8), 3422-3427.
Varela, F., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*: MIT Press.
Wade, N. J., & Bruce, V. (2001). Surveying the seen: 100 years of British vision. *British Journal of Psychology, 92*(1), 79-112. doi: 10.1348/000712601162112
Weiskrantz, L. (2009). *Blindsight: a case study spanning 35 years and new developments*: OUP Oxford.
Zimmer, C. (2012). Sizing Up Consciousness by Its Bits, *The New York Times*. Retrieved from http://www.nytimes.com/2010/09/21/science/21consciousness.html?pagewanted=all&_r=0

Appendix A: Prinz on Attention

Jesse Prinz is a philosopher working intimately with psychology, and Prinz too is wondering what else might be needed. Prinz holds a perceptual theory of consciousness, according to which experience is essentially perceptual. Because of this, Prinz - like us - is deeply concerned with what is doing the work of picking out a definite content from among all the "contents" that are given at various stages of the perceptual constancy hierarchy, and all the numerous activations of perceptual constancies going on at any given time (Jesse J. Prinz, 2010, p. 313).

Prinz holds that *attention* is the special ingredient that is doing the work of *indexing* perceptual constancies (he calls it "mental pointing"). This 'pointing' confers on the individual non-conceptual, phenomenal knowledge (J. Prinz, 2007). While Prinz does not hold that attention (indexing) is necessary to perceive as such, attention figures essentially in picking out a particular coherent content from among all the "contents" that are given at various stages of the perceptual constancy hierarchy at any given time. The upshot of Prinz account is that attention is taken to be the relevant 'pointing' or 'indexing' feature that makes a perception conscious (J. Prinz, 2011). In this attempt at explaining how something like *indexing* is possible, Prinze lets the explanatory work fall on the term attention. According to Prinz, attention makes perception conscious because it "makes information *available* to working memory" (Jesse J. Prinz, 2010, p. 322).

However, in defining attention as making information available to working memory, Prinz has just defined attention in terms that are generally recognized to be the best factual correlate of consciousness, namely, making information available to working memory. Explaining how something can make information available to working memory and what this means, is exactly what a theory of consciousness should, and seems to be trying to, explain. It is the aspect that the two giants in the field, Baars and Edelman, are focusing on (G. M. Edelman et al., 2011). Likewise, Chalmers holds that "when information is *directly available for global control* in a cognitive system, then it is conscious" (Chalmers, 2010, p. 93).

Thus, in being motivated to solve the problem of how the activation of the perceptual constancies can become a definite perceptual content for an individual, Prinz has come very close to invoking consciousness, since he explains *indexing* (my term) in terms of attention (his term), which he defines extremely close to phenomenal experience, i.e. as making

information *available* to working memory. In invoking the term 'attention' in this way, Prinz highlights the connection between determining a particular perceptual content and phenomenal consciousness. However, it does not fundamentally progress our understanding of how *indexing* is possible, or what the nature of phenomenal consciousness is.

Appendix B: Broader Discussion

We set out to explain, and give a positive theory, of one of the most striking features of mind, representation/intentionality, partly by appealing to scientific theories. We set out to explain how some *whole* individual, understood in the light of science, can have mental states that are *about* things in the physical world. To understand our discussion in a broader light, let us briefly consider some history.

First, notice that there have been philosophers who have taken an elimiativist, or deflationist approach to the problem of intentionality. Daniel Dennett illustrates the way in which one can *explain away* instead of explaining: Dennett has held that the problem of intentionality reduces to a kind of intentional stance, an attitude, that we can chose to take when we explain events (D. C. Dennett, 1987).[170] At the opposite end of this way of explaining intentionality are philosophers like John Searle, who one might think *inflates* intentionality by holding that it is constitutively associated with phenomenal experience (Searle, 1983, 1989). Searle holds that the basic or fundamental form of intentionality stems from phenomenal experience; that is, all other kinds of intentionality – that phenomenal experience is not essential to – are derivative or less basic. The explanatory *ideal* was then, as it is now, to account for intentionality in broadly naturalistic terms.

For instance, (Fodor, 1987) starts his Psychosemantics with wanting to explain how his cat can have representational states such as beliefs and desires. Fodor draws a quick analogy to symbols, which are also capable of representation. However, when it finally comes cashing out the very intentionality of the symbols appealed to in the explanation, one must, according to Fodor, "fix a context for the tokenings of certain symbols" (Fodor, 1987, p. 100). Well then what does it mean to fix a context in this way? According to Fodor, we have to say "what it is for a primitive symbol of Mentalese to have a certain interpretation in a certain context" (Fodor, 1987, p. 101). Fodor goes on to suggest various causal accounts of how to fix the meaning of the primitive symbols of the (mental) language (Mentalese).

[170] Burge, in contrast, works with "a distinction between cases in which the notion of representation applies to a subject matter and cases in which it is simply imposed on a subject matter" (Burge, 2010) and assumes the former in his discussions of perceptual representation. I do the same.

There are a whole host of different broadly naturalistic attempts to explain what it is for something to intentionally represent something else. However, it is easy to *suspect* that, in the last instance, when we try to cash out intentionality in terms of causal accounts of reference, regular co-variation, teleology or function (or a great mix of them all) – explanations that have occupied the likes of Fred Dretske (Dretske, 1995b) and Millikan (Millikan, 2009) – one substitutes the striking feature of mind we want to explain with an explanation of something different and less striking. Those few, like Dretske, who are still persisting with this naturalization project, and are at the same time insisting on "closing (what has come to be called) the explanatory gap in a naturalistically acceptable way," are left with, in Dretske's own words, "strongly counterintuitive consequences" (Dretske, 1998).[171] When even the person who, in Burge's words, seems to regard materialism (or naturalism) "as a doctrine to die for" (namely, Dretske) insists that the very phenomenal character of experience depends on things that happened in our past and evolutionary history, we might be justified in speculating about alternative approaches to intentionality and phenomenal experience.[172]

In an illuminating exchange, Burge seems exasperated by Dretske's materialism (Dretske, 1998). However, Dretske in his reply to Burge seems exasperated by the very depth of the question he is out to resolve, namely "how various states of gray matter can make the person in whom that gray matter resides experience (be aware of) the color orange, movement, or middle C when there need be no sounds, colors, or movements inside (or even – during hallucination – outside) the head of the person undergoing these experiences?" (Dretske, 1998).[173] Contrast this statement to the one by Dennett above, or even Fodor's statement that "if aboutness is real, it must be really something else" (Fodor, 1987, p. 97).[174] Thus, there is a

[171] "This is all terribly counterintuitive. But then most people—on first hearing it—also find the idea that physically identical creatures could have different beliefs counterintuitive" (Dretske, 1998, p. 2). Also see, the article "Dreske's Awful Answer" (Dretske, 1995a).

[172] The commitment to ontology seems to be one of the large differences between Burge and Dretske. In this sense, I am more in line with Dretske. Like Dretske, I worried that there might be nomological danglers if we rely too heavily on appeals to explanatory practice in accounting for philosophically fundemental phenomena such as intentionality and experience. For Dretske's take on Burge concerning this issue see (Dretske, 2003).

[173] Dretske then goes on to note that he thinks "they must, it seems be properties the experience in some way represents things as having, properties the experience stands for or is about. They must, in other words, be intentional properties of the experience" (Dretske, 1998).

[174] "I suppose that sooner or later the physicists will complete the catalogue they've been compiling of the ultimate and irreducible properties of things. When they do, the likes of spin, charm, and charge will perhaps

big chasm between the ways in which different philosophers have wanted to account for this most striking feature of mind.

appear on their list. But aboutness surely won't; intentionality simply doesn't go that deep..." (Fodor, 1987, p. 97)

Appendix C: Veridicality Conditions

In this appendix I discuss three further objections relating specifically to how the two supplements, Churchland's MIMP, and Tononi's IIT might shift the focus in explaining the nature of accuracy conditions.

Two Accuracy Conditions?

Our eyes have cones in our retinas that respond to - roughly - three wavelengths. For instance, the light that is coming from a typical computer screen only consists of – roughly – photons with three distinct wavelengths, and yet it seems to us that there is a vast number of different colors on the screen. When we look at a computer screen, the colors are, in some sense, illusions. While this is true, our color representations are accurate for most colors most of the time. For instance, cyan (500nm) looks like green (550nm), cyan looks like blue (450nm), orange (650nm) looks like yellow (600nm), orange looks like red (700nm), and so on. Thus, colors that we generally consider to be close in physicalist terms, e.g. only differing by 50nm, are also typically close in terms of our phenomenal characters.

However, consider the color magenta (pink). Magenta does not, strictly speaking, have an associated wavelength. In general, our eyes can 'see' the wavelengths from roughly ~350 to ~700nm (Goldstein, 2010, p. 47). Humans organize this range into a color-wheel, so that a mixture of ~700nm and ~350nm light produces a distinct color.[175] This color is, approximately, magenta. Our subjective feel for which colors are more similar and different correspond to this color-wheel. We see deep violet (~350nm) as being pretty similar to deep red (~700nm). And we see their mixture as being a color in between, magenta, which is subjectively close to both. But if we actually could average the wavelength of deep violet and deep red, we would instead get (350nm+700nm)/2 = 525nm, a color between cyan and green.[176] Thus, magenta is a color in our color space that in a way fails to match on to the objective feature domain.[177]However, it is still clear that magenta is about some feature of

[175] In fact, this can happen to most colors. Erwin Schrödinger notes in his book *What Is Life* that "if waves of 760μμ, which by themselves produce the sensation of red, are mixed in a definite proportion with waves of 535μμ, which by themselves produce the sensation of green, this mixture produces a yellow that is indistinguishable from the one produced by 590μμ" (Schrödinger, 1992, p. 154). μμ is nanometers.

[176] Of course, our eyes do not average wavelengths in that way in any case. However, it gets the point across.

[177] If we were discussing sounds, we could perhaps have said that deep red was one octave above deep violet,

reality, namely, the conglomeration of wavelengths from opposite sides of the visible spectrum of light. In the larger scheme of things, the fact that deep red looks like deep purple, but is not objectively like deep purple is an anomaly of the way our perceptual system somehow decided to organize incoming input. For most colors, it gets things right. The perceptual system has organized color input in a systematic, mathematically beautiful way, which often happens to correspond to the regularities inherent in nature. [178]

In light of this, we should begin to suspect dissociation between the color space as a whole corresponding to some objective feature domain, and a particular phenomenal character corresponding to some objective features in the environment. We might distinguish two senses in which the perceptual representational content might be right or wrong.

It seems clear that *either* we get it right and our perceptual system presupposes that a "missing shade of blue" will be just like the other colors and vary along certain dimensions, *or* we get it wrong, as in the case of Kripke's 'killer yellow,' a shade of yellow that whenever observed kills the observer (Borge, 2007, pp. 75, 83 fn. 71). Magenta does get it *wrong* insofar as it does not correspond to a unique wavelength that exists between deep violet (~350nm) and deep red (~700nm). The perceptual representation gets it wrong in the sense that the color space seems to "assume" a kind of continuity that does not align with the objective domain of photon-wavelengths. In another sense it seems clear that magenta gets it *right* insofar as it consistently corresponds to the conglomeration of deep violet (~350nm) and deep red (~700nm), where both of these refer to unique properties of the external world. In this sense, magenta gets it right because it is grounded in the right sort of relations to the external world.

Thus, there are two ways in which one might say that the perceptual representational content is accurate with respect to its representata. In one sense one might say it is accurate because it assumes the right sort of thing, e.g. that colors are roughly on a continuum, with similarity and difference relations among them. In another sense one might say it is accurate because it

---

and perhaps that magenta was their harmonic union, and on account of this have made it seem somewhat reasonable relative to the objective features of the world. But of course, we cannot perceive the wavelength of light in the same way we can perceive the wavelength of sound. (In the latter case we can actually hear say, a harmony, when two particular wavelengths resonate in a particular way.)

[178] A good recent article explores by means of statistical and computational methods how it is that we can "make inferences that appear to go far beyond the data available" (Tenenbaum, Kemp, Griffiths, & Goodman, 2011).

is grounded in the right sort of way, e.g. associated or in the right way causally related to, what it is about.

Matching an External Feature Domain?

Churchland is more anxious than Burge to hold that the information space associated with face-recognition and color-recognition accurately matches the objective (external) feature space of how faces typically vary, and how colors typically vary (P. M. Churchland, 1998, 2012).[179] Ideally, Churchland holds that we want to globally match the "acquired internal structure on the one hand, and the independent structure of the external feature-space or property-domain that it more or less successfully portrays, on the other" (P. M. Churchland, 2012, p. 105).[180] Burge, by contrast, seems to think that there is "nothing *per se* within the limits of an individual's body from which one could recover anything relevant to specific properties in the environment that perception is as of" (Burge, 2010, p. 85). To adopt Churchland's accuracy conditions, however, one needs some way to determine whether or not the map as a whole can match an external feature domain. The idea that some informational space can be similar or that it can match the external world in some way is not only problematic for epistemological reasons. There have been challenges to the very idea that it is possible to define such measures (Fodor & LePore, 1993; Guigon, Forthcoming; Morreau, 2010). It is worth noting that as Churchland accepts that perceptual content consists of causal relations to stable and objective macrofeatures of the external environment, one does not need a perfect measure of the matching relation. As Burge puts it, "the point is that the general elements cannot do all the referential work. Some of the work is done by the perceptions' being caused by particulars that are referred to" (Burge, 2010, p. 83). Even so, Churchland aims to answer what he calls the 'Fodor/Lepore challenge' (P. M. Churchland, 1998), a challenge to the very idea that it is possible to unambiguously define a similarity or matching relation between an internal map-like structure, and an external environment. Since the idea is

[179] See (P. Churchland, 2007) for his treatment of how "color-qualia space is a map of reflectance-profile space." However, Burge also thinks that perceptual color representation is about something like surface reflectance (Burge, 2010, p. 48). Burge also says that "The formation principles describe law-like regularities, in the perceptual system, that reflect or mirror law-like regularities in the distal environment" (Burge, 2010, p. 346).

[180] Churchland at one place writes that a map can have "its internal portrayal of the external structures that it purports to, *or is deployed to*, represent.[my emphasis]" (P. M. Churchland, 2012, p. 123)

to match the internal, map-like structure with aspects of the external environment that are relevant representationally, Churchland speaks of matching an 'external feature domain.'

By the method of matching an internal, map-like structure with aspects of an external feature domain, one should perhaps only hope to exclude certain representationally relevant possibilities, just as the perceptual constancies, as it were, by themselves exclude certain representational alternatives. Churchland in his (P. M. Churchland, 1998) and (P. M. Churchland, 2012) limits himself to solving the problem of how neuronally instantiated, internal map-like structures, can be compared for similarity across individuals that differ and differ in the way they instantiate such structures. The idea is that this more limited solution generalizes to solving the problem of matching an external feature domain.[181] Churchland takes (Laakso & Cottrell, 2000) to have given a good account of representational similarity across neural systems. I will not detail this approach here.

Abstract Accuracy Conditions?

If the subject perceives through some kind of abstract structural information, should we not be just as anxious about our epistemic access to the world as we were before? This is perhaps the most serious threat given the situation this essay has led the reader into. It is no accident that Churchland's new book is called Plato's Camera (P. M. Churchland, 2012).

However, I think we should remind ourselves that we are just now beginning to see the potential inherent in phenomenal experience to help solve the problem of perceptual duality. We have all reason to suppose that the crude kind of 'matching' Churchland invokes is going to get strengthened as the science progresses. I think this is especially true if we supplement his theory with something like Tononi's IIT. The manner in which perceptual constancies can be accurate in relation to what they are about only gives us a hint of how to proceed in this direction. For instance, early work by Tononi, Edelman and Olaf Sporns show that there might be a change in what they call 'matching complexity' in a system, such as the brain when it "encounters extrinsic stimuli reflects the degree to which intrinsic statistical relationships between its constituent neuronal groups match statistical relationships present in the environment" (Tononi & Edelman, 1998, p. 481; Tononi, Sporns, & Edelman, 1996). This

[181] One reason for solving this more limited problem is the idea that we can know an objective, external feature domain is open to doubt by the radical skeptic. Initially, it is most important to establish that abstract, map-like representational structures *can* be compared mathematically according to their objective similarity.

might partly vindicate philosophical notions such as ‘structural coupling’ that embodied cognition philosophers such as (Varela et al., 1991) have invoked in explaining the embodied relation between perceptual experience and the environment. For instance, Tononi sketches the view that “through natural selection, epigenesist, and learning, informational relationships in the world mold informational relationships within the main complex that ‘resonate’ best on a commensurate spatial and temporal scale … In this way, qualia—the shapes of experience—come to be molded, sculpted, and refined by the information structure of events in the world” (Tononi, 2008, p. 240). Again the idea is to do this in a precise, formally rigid, quantitative way.

Thus, if one is antecedently inclined towards the idea of perceptual representational content, then the way through integrated information might turn out highly illuminating, and might, in the end, answer questions that now appear obscure. Indeed, almost ironically, the Intergated Information Theory of consciousness might provide reason to suppose that our phenomenally conscious experience when confronted with the actual world is of a different kind than that that which occurs in dreaming or hallucinations.

For instance, Tononi states “a working hypothesis that the quantity of ‘matching’ between the informational relationships inside a complex and the informational structure of the world can be evaluated, at least in principle, by comparing the value of Φ when a complex is exposed to the environment, to the value of Φ when the complex is isolated or ‘dreaming’” (Tononi, 2008, p. 240). In this way, we can have a testable, empirical prediction about whether these states are different or not, just like Tononi’s theory retrodicted the common sense view that phenomenal consciousness only occurs at a certain limited spatio-temporal scale in biological creatures.