The Dissertation Committee for Katherine Anne Piatti Certifies that this is the approved version of the following dissertation:

Cognitive Penetration and the Perceptual Representation of High-Level Properties

Committee:

Michael Tye, Supervisor

Michelle Montague

Ernest D Sosa

Joshua Dever

Susanna Siegel

Cognitive Penetration and the Perceptual Representation of High-Level Properties

by

Katherine Anne Piatti, B.A; M.A.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

May, 2017

Cognitive Penetration and the Perceptual Representation of High-Level Properties

Katherine Anne Piatti, Ph.D.

The University of Texas at Austin, 2016

Supervisor: Michael Tye

In recent years, two of the most popular debates in the philosophy of perception have been about whether there's cognitive penetration and whether high-level properties are represented in our perceptual experiences. My dissertation concerns both of those debates. In the first paper, I argue that the debates about cognitive penetration and the perceptual representation of high-level properties are related. Namely, there has to be cognitive penetration for it to be plausible that our perceptual experiences represent high-level properties. In the second paper, I add to the literature challenging the alleged empirical evidence for cognitive penetration. I focus, in particular, on studies purportedly finding that (direct) influence from subjects cognitive states had altered their perceptual experiences over time. I argue those studies provide empirical evidence of perceptual adaptation, not cognitive penetration. In the third paper, I argue the view that our visual experiences represent high-level properties is unmotivated and I argue the explanation for how we come to perceptually represent high-level properties is implausible.

Table of Contents

Chapter 1: No High-Level (Perceptual) Representation, Without Cognitive Penetration

Looking at a contour map, the student sees lines on paper, the cartographer a picture of a terrain. Looking at a bubble-chamber photograph, the student sees confused and broken lines, the physicist a record of familiar subnuclear events. Kuhn, 1970: 111

1.1 INTRODUCTION

Two of the hottest debates in the philosophy of perception, in recent years, have been over whether high-level properties are represented in our perceptual experiences (see e.g. Hawley and Macpherson, 2011; Siegel, 2006) and whether our thoughts can directly influence or alter our perceptual experiences (see e.g. Raftopoulos, 2009; Pylyshyn, 1999) . The first is the debate over the existence of *high-level (perceptual) representations* and the second is the debate over the existence of *cognitive penetration*.

I am skeptical of the existence of both high-level representations and cognitive penetration. However, for the purposes of this paper, I am going to set my doubts aside. My aim in this essay is to argue that the debate over high-level representations is connected to the debate over cognitive penetration. In particular, I will argue that it's implausible that there are high-level representations if there's no cognitive penetration.

I have two main arguments. My first argument is based on the premise that the perceptual representation of high-level properties plausibly requires the possession and application of high-level property concepts. For example, it's plausible that for an individual to perceptually represent the property of being a pine tree, she needs to possess the concept PINE TREE and be disposed to apply that concept to her experiences. In short, it's plausible that perceptual experiences representing high-level properties have (or would have) conceptual contents. And I argue if

perceptual experiences representing high-level properties have conceptual contents, then subjects can't come to perceptually represent high-level properties without their cognitive states directly influencing their perceptual experiences.

My second argument is based on the premise that the plausibility of there being high-level representations is contingent upon there being a plausible account of how subjects come to perceptually represent high-level properties. And I will argue that the only plausible account in the offing is in terms of cognitive penetration.

The remainder of the paper will proceed as follows. In the next section, I will summarize the debate over the existence of high-level representations, and present the main argument that's been put forward in their favor—the phenomenal contrast argument. In section 3, I will summarize the debate over cognitive penetration, and discuss how high-levelists have appealed to cognitive penetration in arguing for their view. In section 4, I will present my arguments and respond to a few potential objections. In section 5, I will offer some concluding remarks.

1.2 HIGH-LEVEL (PERCEPTUAL) REPRESENTATION

As I sit here typing these words, I am having a variety of perceptual experiences—e.g. I am feeling the pressure of the floor beneath my feet, I am hearing the tick-tick-tick of my wall clock, I am seeing a half-full mug of coffee on my desk, I am smelling coffee, and I am tasting the sweet-minty flavor of the gum I am chewing.

Many contemporary philosophers think our perceptual experiences, like our beliefs and other attitudes, are mental states that represent the world as being some way or other, and thus are either accurate or inaccurate (our experiences are accurate if they represent the world a way it actually is, and inaccurate if they represent it some way it is not). The *content* of a state is a

specification of the way it represents the world as being. So, for example, the belief that the Milky Way is a galaxy represents and has the content that the Milky Way is a galaxy.

While many contemporary philosophers may agree that our perceptual experiences are representational and thus have contents, there is a lot of disagreement among them over the nature of those contents (see Siegel, 2016). Are they structured? Unstructured? Propositional? Non-propositional? Sets of Possible Worlds? Fregean? Russellian? Existential? Singular? Gappy? The debate over high-level representations concerns a disagreement about whether certain properties—namely, high-level properties—can figure in the contents of our experiences (i.e. whether our perceptual experiences represent high-level properties).

Take the perceptual experiences I was having as I typed the first few lines of this section. To keep things manageable, let's focus on my visual experience, which I described as "seeing a half-full mug of coffee sitting on my desk". There are lots of things you might want to know about that experience: was it veridical (and accurate) or hallucinatory (and inaccurate), what (if anything) was it an experience of (i.e. what were its objects), did the fact that I was focused on what I was writing and not my experience have an effect on my experience? Something else you might want to know is what *properties* did my experience represent?

There is widespread agreement that the experience represented sensible qualities or low-level properties like the color of the coffee, the shape of the mug, and the spatial relationship between the mug and the desk. But what about so-called high-level properties like being a mug, being a desk, and being coffee? Does my visual experience represent those properties too? A somewhat indirect way of getting at the same question is to ask whether my experience would have been *inaccurate* or *illusory* if, instead of coffee, the mug had actually contained cola.

The orthodox view (if there is one) is that our perceptual experiences can only represent "sensible qualities" like color, shape, pitch, tone, pressure, and texture properties. So on the

orthodox view my experience could only have represented properties like the color of the coffee and the shape of the mug. However, a handful of philosophers have rejected the orthodoxy, and contend that our perceptual experiences can also represent (some) high-level properties (e.g. McClelland, 2016; Siegel, 2010). On this view, my experience could also have represented the properties of being coffee and being a mug. Since I'm calling perceptual experiences that represent high-level properties *high-level representations*, this can be understood as a debate over the existence of high-level representations. I'll call those who believe in the existence of high-level representations, *high-levelists*, and those who deny their existence (i.e. believe that our perceptual experiences only represent low-level properties) *low-levelists*.

You might be tempted to think high-levelists are right since I obviously *see* more than colors and shapes (and other low-level properties), I see a mug of coffee. That's true as far as it goes. But it doesn't go far enough. I *see* a mug of coffee in the sense that a mug of coffee is the object of my experience. But no one—least of all low-levelists—contends that the objects of our experiences are just colors and shapes. But seeing a mug of coffee doesn't require visually representing the high-level properties of being a mug and being coffee. So, the fact that I see a mug of coffee doesn't support the claim that my visual experience represented those high-level properties. I also *know* that I'm seeing a mug of coffee (I filled the mug, and set it there myself). But I can know I'm seeing a mug of coffee without visually representing the properties of being a mug and being coffee.

Some high-levelists contend that common sense is on their side. They claim it's common sense that high-level properties are represented in our perceptual experiences. But that's just false. It's not false because it's common sense that only low-level properties are represented in our perceptual experiences. That's not common sense either. It's false because there is no common sense view about which properties are visually represented.

1.1.1 Phenomenal Contrast Arguments for High-level Representation

High-levelists have tried to establish that high-level properties are represented in our perceptual experiences (i.e. the existence of high-level representations) using *phenomenal contrast arguments* (e.g. Bayne, 2011; Siegel, 2007; Siewert, 1998).[1] These arguments start by describing a pair of situations in which it seems like subjects have (or would have) perceptual experiences that differ, but where it doesn't seem like those perceptual experiences could be representing different low-level properties. High-levelists then argue that the best (or perhaps only) way to explain/account for subjects (purportedly) having different perceptual experiences in these pairs of situations is that their perceptual experiences were representing different high-level properties. But if the explanation is that different high-level properties were represented in subjects' perceptual experiences, it must be the case that our perceptual experiences can represent high-level properties, and thus that there are high-level representations.

For the phenomenal contrast arguments to have any chance of success, the audience has to find it plausible both that subjects' perceptual experiences would be different in the situations in question, *and* that the difference(s) between subjects' perceptual experiences can't be explained in terms of differences in what low-level properties those perceptual experiences represent. High-levelists have described a number of different cases in arguing for (the existence of) high-level representations. But the most prominent examples either have to do with, what we might call, *expert perception* or with our perceptual experiences of ambiguous figures.

Expert perception is the phenomenon whereby gaining knowledge/expertise in a certain domain alters (improves) one's perceptual experiences of things within that domain—e.g. when becoming a wine expert alters or improves how wines taste to you (i.e. your taste experiences of

[1] Phenomenal contrast arguments have also been used in other areas of philosophy (e.g. Chudnoff, 2015). For criticism of phenomenal contrast arguments see Koksvik, 2015.

wine), when becoming a radiologist alters how x-rays look to you (i.e. your visual experiences of x-rays), and when becoming a musician alters how music sounds to you (i.e. your auditory experiences of music). Siegel describes a pair of cases where becoming a pine tree expert (of a sort) alters how pine trees look to you (i.e. your visual experiences of pine trees).

> Suppose you have never seen a pine tree before and are hired to cut down all the pine trees in a grove containing trees of many different sorts. Someone points out to you which trees are pine trees. Some weeks pass, and your disposition to distinguish the pine trees from the others improves. Eventually, you can spot the pine trees immediately: they become visually salient to you. Like the recognitional disposition you gain, the salience of the trees emerges gradually. Gaming this recognitional disposition is reflected in a phenomenological difference between the visual experiences had before and those had after the recognitional disposition was fully developed (Siegel, 2010: 100).

Compare your visual experiences of the trees you had *before* and *after* you learned to recognize pine trees (i.e. as a novice, and then as an expert). According to Siegel, those visual experiences are different, and she contends that's because pine trees *look* different to you as an expert from how they did to you as a novice.

There are lots of banal reasons why pine trees might look different to you, after you learned to recognize pine trees – e.g. perhaps the weather conditions were different, or maybe your eyesight changed, or very plausibly, after learning to recognize pine trees, you attended to them in a very different way. However, those kinds of differences don't exist in the cases Siegel has in mind. In other words, we are supposed to be comparing the visual experiences you would have in situations that are *identical*, except for the difference in your ability to recognize pine trees by sight.

If we suppose that abilities to recognize things (e.g. pine trees) can be cashed out in terms of some kind of cognitive state(s) (e.g. know how), then the circumstances in question apparently only differ with respect to your cognitive states. And if the circumstances only differ with respect to your

product in the states, it doesn't seem like your visual experiences could be representing different low-level properties.

So, high-levelists contend that the difference between your visual experience before and after learning to recognize pine trees is the result of a difference in the high-level properties your visual experiences are representing. More specifically, the claim is that your visual experiences are different because, after learning to recognize pine trees, you came to visually represent the property of *being a pine tree*. Supposedly, learning to recognize pine trees gave you the ability to visually represent a property—namely, being a pine tree—that you could not *perceptually* represent before .[2]

Some of the other cases high-levelists have discussed are drawn from our perceptual experiences of ambiguous figures like Wittgenstein's duck-rabbit (Figure 1).[3] Compare your visual experience of the figure looking like a duck to your visual experience of the figure looking like a rabbit. According to high-levelists, those visual experiences are different, i.e. the figure looks different when it looks like a duck versus when it looks like a rabbit.

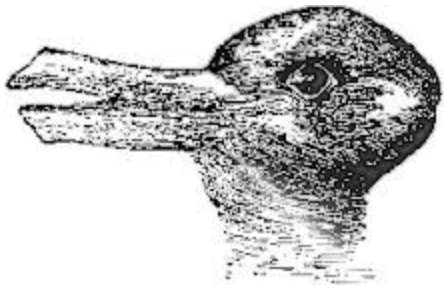

Figure 1: Wittgenstein's Duck-Rabbit

We find the fact that ambiguous figures can take on different looks surprising and delightful because nothing about the figure has changed (and, in a sense, it looks exactly the same). But what makes our

[2] Of course you could *conceptually* represent the property.
[3] See, e.g., Price, 2009.

experiences of ambiguous figures surprising and delightful is also what makes them fodder for high-levelists. Like cases involving novice and expert perception, cases involving disambiguated experiences of ambiguous figures seem to be cases where subjects have perceptual experiences that are different, and it doesn't seem like there are differences in the low-level properties their perceptual experiences are representing.

As with (purported) experiential differences associated with expert perception, there are banal reasons why disambiguated experiences of ambiguous figures might be different. However, since we can flip back and forth between the disambiguate experiences – e.g. between the experience of the figure looking like a duck and the experience of the figure looking like a rabbit – banal explanations are not as forthcoming. Since we can switch back and forth in real time, we can ensure there are no differences in the lighting conditions, our eyesight, and, perhaps even, our attention. High-levelists contend that your visual experience of seeing the figure as a duck is different from your visual experience of seeing the figure as a rabbit because your experiences are representing different high-level properties—the property of *being a duck*, in the one case, and the property of *being a rabbit*, in the other. In other words, when the figure looks to you like a duck, that's because your visual experience is representing the property of being a duck, and when the figure looks to you like a rabbit, that's because your visual experience is representing the property of being a rabbit.

If phenomenal contrast arguments are successful, there are differences between our perceptual experiences that can best (or perhaps only) be explained/accounted for in terms of differences in high-level properties those experiences represent. And if that's the case, there are high-level (perceptual) representations.[4]

[4] There is a real question as to whether the phenomenal contrast arguments are successful. Low-levelists have objected to these arguments either by denying that subjects would have different perceptual experiences in

1.3 COGNITIVE PENETRATION

The debate over high-level representations has to do with the contents of our perceptual experiences – whether high-level properties can figure in those contents. The debate over cognitive penetration also has to do with the contents of our perceptual experiences – whether those contents can be (directly) influenced/altered by our cognitive states (see e.g. Stokes, 2013; Wu, 2013; Zeimbekis, 2013).

It's uncontroversial that the contents of our perceptual experiences are directly influenced/altered by the objects of our experiences (e.g. the properties they have), our perspective on those objects (e.g. how far away they are), the condition/state of our perceptual systems and sensory organs (e.g. the acuity of our eyes), the way we distribute and focus our attention (e.g. whether we are distracted), and the surrounding environment (e.g. the lighting conditions). But proponents of cognitive penetration contend that what we *think* (e.g. our beliefs, and desires, intentions, and expectations) can also directly influence/alter the contents of our perceptual experiences.[5]

It's uncontroversial that our thoughts can *indirectly* influence/alter the contents of our perceptual experiences. Suppose you look in the cupboard because you want a cookie and believe there are cookies in the cupboard. The visual experience you have while looking in the cupboard has to do with your thoughts in two obvious respects: (i) it is because of your thoughts (namely, that you want a cookie and believe there are cookies in the cupboard) that you are looking in the cupboard—and the visual experience you have depends on where you are looking—and (ii) presumably what you are thinking also influences what you're paying attention to and what you're ignoring—and the

---

the contrast cases, or by accepting there are differences but denying those differences would be best explained in terms of differences in high-level properties being represented in their perceptual experiences.

[5] The debate over cognitive penetration can also be understood as a debate about our *mental architecture*, specifically, whether our perceptual systems are *modular* or *informationally encapsulated*. For the classic defense of the thesis that our perceptual systems are modular see Fodor, 1983.

contents of your visual experience also depends on what you're paying attention to. But opponents of cognitive penetration maintain that once you are looking in the cupboard, it doesn't matter what you're thinking—the content of your visual experience (i.e. the way things look to you at that moment) is fully determined by the things in the cupboard, your viewing perspective, your perceptual systems and sensory organs, the way you focus and distribute your attention, the surrounding environment, and perhaps how you are related to your environment (e.g. your evolutionary history). To put it differently, once you are looking in the cupboard and attending, say, to the jar of cookies, it makes no difference what you think.

But proponents of cognitive penetration contend that our thoughts can *directly* influence/alter the contents of our perceptual experiences. Suppose you spot a plain sugar cookie while looking in the cupboard. According to proponents of cognitive penetration, how that cookie looks to you, can be directly influenced/altered by what you think, and thus how it looks to you depends on what you think. If, for example, you really want a peanut butter cookie, your desire for a peanut butter cookie could make the sugar cookie you see look slightly bigger and browner than it actually is (and presumably also bigger and browner than it would have looked to you if you hadn't really wanted a peanut butter cookie).

The primary support for believing there's cognitive penetration is empirical studies purportedly finding that influence from subjects' cognitive states had somehow altered their perceptual experiences. For example, a study purportedly finding that subjects' expectations about the color bananas had altered their visual experiences of black and white images of bananas—i.e. made them look slightly yellow (when they we objectively achromatic).[6] And a study reported the finding that children's desires for money had altered their visual experiences of coins – making the coins look slightly bigger than they actually were (Bruner and Goodman, 1947).

[6] Study discussed in Deroy, 2013.

These studies, and dozens of others like them, have convinced some philosophers that the contents of our perceptual experiences are (at times) directly influenced or altered by our cognitive states – and thus that there is cognitive penetration. Of course, whether any of those alleged empirical findings are, in fact, convincing evidence of cognitive penetration is up for debate.[7] And for what it's worth, I'm unconvinced. But might aim in this essay is not to argue against the existence of cognitive penetration. It's to argue that there can't be high-level representations unless there's cognitive penetration.

1.4 Cognitive Penetration & High-level Representation

There are ongoing debates in the philosophy of perception about whether high-level properties figure in the contents of our perceptual experiences (i.e. whether there are high-level representations), and whether our cognitive states can directly influence/alter the contents of our perceptual experiences (i.e. whether there's cognitive penetration). Those debates are far from settled, but I don't intend to weigh in on them directly. Instead, my aim is to argue that they are related. More specifically, that the existence of high-level representations is dependent on the existence of cognitive penetration – if there is no cognitive penetration, there are no high-level representations.

1.4.1 Argument from Conceptual Content

Imagine that you are looking at a vegetable garden on a clear sunny day and a ripe green zucchini catches your eye. At that moment, even if you have no idea what you're looking at – you have never

[7] For criticism of the empirical studies, see Firestone, 2013 & 2014, and Firestone and Scholl, 2016.

heard of or seen a zucchini before – you will be having visual experience of a zucchini. However, presumably, you won't be thinking or believing that you are having a visual experience of a zucchini. To think or believe that you are having a visual experience of a zucchini is to have a thought or belief that has the content that you are having a visual experience of a zucchini. But you can't have a thought or belief with that content unless you have at least some idea that you're looking at a zucchini. And in order to have some idea that you're looking at a zucchini, you have to possess the concept ZUCCHINI. That's because our thoughts and beliefs are mental states that represent the world in a conceptual way, and thus are mental states that have conceptual contents. A mental state M has conceptual contents if and only if (i) M has accuracy conditions and (ii) in order to have or be in M, subjects need to possess and apply the concepts that would be used in a canonical specification of M's accuracy conditions.

My first argument begins with the premise that perceptual experiences representing high-level properties (i.e. high-level representations) are the same as thoughts and beliefs in that they have (or would have, if they existed) conceptual contents. In order for you to have a visual experience that represents the property of being a zucchini, you would need to possess the concept ZUCCHINI and apply that concept to your visual experience. But I argue if high-level representations must have conceptual contents, then there can't be high-level representations without cognition directly influencing perception. And thus there can't be high-level representations without cognitive penetration. Here's the argument in full.

1) High-level representations have conceptual contents.
2) If high-level representations have conceptual contents, then subjects need to possess high-level property concepts and be disposed to apply those concepts when undergoing the relevant perceptual experiences, for those perceptual experiences to represent high-level properties.
3) So, subjects need to possess high-level property concepts and be disposed to apply those concepts when undergoing the relevant perceptual experiences, for those perceptual experiences to represent high-level properties.

4) But possessing high-level property concepts and being disposed to apply those concepts to your perceptual experiences are matters of having and/or being in certain cognitive states.
5) So, subjects need to have and/or be in certain cognitive states for their perceptual experiences to represent high-level properties.
6) But if subjects need to have and/or be in certain cognitive states for their perceptual experiences to represent high-level properties, there can't be high-level representations without one's cognitive states directly influencing or altering one's perceptual experiences.
7) But cognitive penetration just is cognitive states directly influencing or altering perceptual experiences.
8) So, there can't be high-level representations without cognitive penetration.

The only substantive premises of this argument are 1, 4, and 6. And, of those, I take it that only premises 1 and 6 are debatable. Let me say a bit in defense of each.

Again, premise (1) is the claim that perceptual experiences representing high-level properties have (or would have, if they existed) conceptual contents. And that's because perceptual experiences representing high-level properties are, I claim, like thoughts and beliefs in that in order for subjects to have or be in those mental states, they would need to possess certain concepts and apply those concepts to their perceptual experiences. For premise (1) to be false, it would have to be the case that subjects' perceptual experiences sometimes represent high-level properties for which subjects lack the relevant high-level property concepts and/or fail to apply the relevant concepts to their perceptual experiences.

I find the claim that high-level representations have (or would have) conceptual contents exceedingly plausible—at least as plausible as the claim that thoughts and beliefs have conceptual contents. Consider again the visual experience you would have on a clear sunny day while looking at a ripe green zucchini dangling from a vine. What properties would that experience represent? Assuming your vision is good and your view is unobstructed (i.e. you see the zucchini clearly), then no matter what you believe or what concepts you possess or apply to your experience, we can

assume that your visual experience would represent the color and shape properties, as well some other low-level properties, of the things you see (e.g. the property of being green).

But what about the high-level property of being a zucchini, would your visual experience represent that? According to high-levelists, it might. However, I take it that's only plausible on the assumption that you at least possess the concept ZUCCHINI and are inclined to, in some sense, apply that concept to your experience. But if we suppose that you haven't the faintest idea what a zucchini is – have never heard of or seen one before in your life – it's not even remotely plausible that your visual experience would represent the property of being the zucchini. How the heck could it do that?

Presumably, our mental states only represent those properties they have somehow glommed onto or detected. If you know what the zucchini is (and thus possess the concept ZUCCHINI), then perhaps your visual system could use that knowledge to parse or group incoming sensory information in such a way that allows it to glom onto or detect the property of being as zucchini. However, I don't see how your visual experience (or the underlying visual state of which it is an experience of) could have glommed onto or detected the property of being a zucchini is you have no idea what a zucchini is. So I conclude that your visual experience could not have (and thus would not have) represented the property of being a zucchini unless you had possessed and applied your concept ZUCCHINI.

Another reason to think that high-level representations have (or what have) conceptual contents is that all of the cases that high-levelists have appealed to in arguing for their view – i.e. all the cases in which subjects' perceptual experiences purportedly represent high-level properties – are cases in which subjects both possess high-level property concepts and apply those concepts to their relevant perceptual experiences. So, for example, Siegel argues that the experts visual experience of pine trees is different from the novice's visual experience of pine trees because the expert's visual

experience represents, whereas the novice's does not, the property of being a pine tree. But she contends that the expert's perceptual experience came to represent the property of being a pine tree as a result of her learning to recognize pine trees by sight. But if the expert can recognize pine trees by sight, she possesses the concept PINE TREE and she applies that concept to her visual experience of pine trees.

This also applies to cases involving ambiguous figures. High-levelists contend that the visual experience you have when Wittgenstein's duck rabbit looks to you like a duck is different from the visual experience you have when the figure looks to you like a rabbit, because in the first case your visual experience is representing the property of being a duck, whereas in the second case your visual experience is representing the figure as being a rabbit. But their claim is that when your visual experience represents those properties it's because you're thinking of the figure as a duck or you're thinking of the figure as a rabbit. What if you're thinking of the figure in those ways, you possess the concepts DUCK and RABBIT and are applying those concepts to your visual experiences.

If it were possible for our perceptual experiences to represent high-level properties even when we lack the relevant high-level property concepts or failed to apply those concepts to our perceptual experiences, you would expect high-levelists to have discussed some cases like that. The fact that they don't is some reason to believe that such cases aren't possible. At this point, you might be thinking that high-levelists do discuss such cases – namely, cases in which subjects' perceptual experiences purportedly represent causal, linguistic, and facial properties. But for reasons I will discuss in section 4.4, cases in which subjects' perceptual experiences purportedly represent causal, linguistic, and facial properties don't count (i.e. are not counterexamples to this point).

One final reason that I want to mention for thinking that high-level representations have (or would have) conceptual contents concerns how high-levelists have motivated their view. As I mentioned when presenting the view (in section 2), the primary motivation high-levelists have given

us for thinking there are high-level representations is the existence of contrast cases, i.e. pairs of cases in which subjects seem to have perceptual experiences that are different, yet in which subjects' perceptual experiences don't seem to be representing different low-level properties. High-levelists use these pairs of cases to argue that the best (and perhaps only) way to explain subjects (purportedly) having different perceptual experiences in these pairs of cases is that their perceptual experiences represent different high-level properties. Because subjects' perceptual experiences can't represent different high-level properties unless it's the case that high-level properties are (sometimes) represented in our perceptual experiences.

In all the contrast cases high-levelists discuss, there are obvious cognitive differences between subjects in the relevant pairs of cases (e.g. differences in the high-level concepts subjects applied to their perceptual experiences). And in all the contrast cases these cognitive differences between subjects are the only obvious differences there are. In other words, apart from the cognitive differences between subjects, the relevant cases are supposed to be identical. There's a good reason for that. The phenomenal contrast arguments only work if it's plausible that in the contrast cases subjects would have perceptual experiences that are different and the best (or only) explanation for subjects' perceptual experiences being different is that they were representing different high-level properties. So for these arguments to even get off the ground it has to be plausible that there are differences in subjects' perceptual experiences despite the fact that there are no obvious differences in which low-level properties those experiences were representing.

It is most plausible that both conditions are met if the pairs of cases involve situations that only obviously differ from one another in terms of subjects' cognitive states and yet subjects seem to have different perceptual experiences. Because if the situations only differ from one another in terms of subjects' cognitive states, there doesn't seem to be any basis for difference(s) in lower-level properties perceptually represented.

This poses a dilemma. Either the only difference between the contrast cases is with respect to subjects' cognitive states or there are some noncognitive differences between the contrast cases as well. If the only difference(s) were with respect to subjects' cognitive states (and yet their perceptual experiences were different), then presumably the differences in subjects' cognitive states were somehow responsible for subjects' perceptual experiences being different. But if there were noncognitive differences between the contrast cases, that undermines the claim that there couldn't be differences in low-level properties perceptually represented to explain/account for subjects' perceptual experiences being different.

But it should be uncontroversial that possessing high-level property concepts and applying those concepts to one's perceptual experiences are matters of cognition, i.e. involve or are constituted by having and/or being in certain cognitive states. And if high-level properties can only be represented in one's perceptual experiences if and when she has or is in certain cognitive states, high-level properties can't be represented in her perceptual experiences without her cognitive states having a direct influence on or altering her perceptual experiences. But cognitive states having a direct influence on or altering perceptual experiences just is cognitive penetration. So there can't be perceptual experiences representing high-level properties without cognitive penetration.

1.4.2 Objections and Replies

First Objection. Premise (1) of your argument is false. High-level representations don't (or wouldn't) have conceptual contents. A number of philosophers have argued that our perceptual experiences are unlike our thoughts and beliefs in that they are mental states that represent the world in a *non-conceptual* way, and thus have *non-conceptual contents.*[8] A mental state M has nonconceptual content if

[8] See Bernueacute, 2015 for an overview of the debate about nonconceptual mental content.

and only if (i) M has accuracy conditions and (ii) subjects' of M need not possess or be disposed to apply the concepts that would be used in a canonical specification of M's accuracy conditions.

One motivation these philosophers have given for claiming that perceptual experiences have non-conceptual content appeals to the so-called "richness" of perceptual experiences. Here's Michael Tye on the richness of color experiences.

> Human sensory experience is enormously rich. Take color experience. There is a plenitude of detail here that goes far beyond our concepts. Humans can experience an enormous number of subtly different colors, something on the order of 10 million, according to some estimates. But we have names for only a few of these colors, and we also have no stored representations in memory for most colors.… This is why I cannot go into a paint store and reliably identify a color on a chart as exactly matching the precise hue of my dining room walls. I possess the concept red, of course, and I exercise it when I recognize something as red, but I lack the concepts for determinate hues. (Tye, 2000:11)

Tye's point is just that our perceptual experiences of color far outstrip our color concepts—we have highly determinate color experiences (e.g. of red_{17}) even though we lack highly determinate color concepts (e.g. the concept RED_{17}). But if our perceptual experiences of color outstrip our color concepts (i.e. our color experiences are more determinate than our color concepts), we obviously don't need to possess and deploy color concepts for our perceptual experiences to represent those colors. And thus, our perceptual experiences of color must have non-conceptual contents.

Philosophers have also motivated the view that our perceptual experiences have nonconceptual contents by appealing to apparent similarities between our perceptual systems and experiences, and the perceptual systems and experiences of babies and some non-human animals (Peacocke, 2001). We know by direct comparison that babies and some non-human animals have perceptual systems that are substantially similar to our own. And we have a fair amount of evidence that they also have perceptual experiences that are similar (at least in certain respects) to our own.

However, it's plausible that babies and non-human animals don't possess all the concepts that would be needed to specify the contents of those perceptual experiences. So, the possession of concepts and disposition to apply them must not be necessary for having perceptual experiences. Since high-level representations are perceptual experiences representing high-level properties, the motivations for thinking that perceptual experiences have non-conceptual contents are also motivations for thinking that high-level representations have (or would have) conceptual contents.

Reply. The motivations philosophers have given for thinking that our perceptual experiences have non-conceptual contents don't apply to perceptual experiences representing high-level properties. Take the motivation stemming from the apparent richness of our perceptual experiences. We seem to perceptually experience far more colors and shapes, for example, then we have color and shape concepts. This suggests that we don't need to possess and apply color and shape concepts in order to have perceptual experiences of color and shape properties. But richness considerations don't seem to apply when it comes to perceptual experiences representing high-level properties. We don't seem to perceptually experience more high-level properties then we have high-level property concepts.

The motivation arising from empirical evidence concerning the perceptual systems and experiences of babies and nonhuman animals also doesn't seem to apply when it comes to perceptual experiences of high-level properties. We don't have empirical evidence concerning the perceptual representation of high-level properties. Any evidence we have comes from the contrast cases and the associated phenomenal contrast arguments.[9]

Second Objection. Premise (6) of your argument is false. It assumes if the subjects need to have and/or be in certain cognitive states in order to perceptually represent high-level properties, (i.e.

[9] Some people contend there is empirical evidence that babies perceptually represent facial properties. But for reasons I will discuss in section 4.4, I don't consider this a counterexample to my claim.

having and/or being in certain cognitive states is *necessary* for the perceptual representation of high-level properties), then those cognitive states must be directly *influencing* their perceptual experiences of high-level properties. But even if having and/or being in certain cognitive states is necessary for the perceptual representation of high-level properties, you can't conclude that those cognitive states are influencing subjects' perceptual experiences of high-level properties. It's not generally true that if A is necessary for B, that A influences B (directly or otherwise). To see that, just consider that being at least 6 feet tall is necessary for being 7 feet tall. But it's clearly *not* the case that being at least 6 feet tall influences being 7 feet tall.

R*eply*. I grant in general, that you can't conclude from A's being necessary for B, that A influences B. However, I think the conclusion is warranted here. Statements of the form "A is necessary for B" do not express primitive facts – i.e. if "A is necessary for B" is true, it's true because or in virtue of the fact that there is some underlying relationship between A and B. Being at least 6 feet tall is necessary for being 7 feet tall because in being 7 feet tall, one is by definition, also at least 6 feet tall. What, then, is the underlying relationship between having and/or being in certain cognitive states and perceptual experiences representing high-level properties? It's certainly not a matter of definition. My contention is that subjects need to have and/or be in certain cognitive states in order to perceptually represent high-level properties because the perceptual representation of high-level properties requires or involves influence from cognition.

1.4.3 Argument from Explanation

High-levelists have concluded that the best (and perhaps only) explanation for certain perceptual experiences being different (e.g. subjects' perceptual experiences in the contrast cases) is that those experiences represent different high-level properties. But even if we grant that the perceptual experiences in question are different (and, in many cases, that's up for debate), I take it the

conclusion that those perceptual experiences are different because they represent different high-level properties can only be sustained if there's a plausible account of how perceptual experiences come to represent (different) high-level properties.

Up to this point, high-levelists have said remarkably little about how high-level properties come to be represented in our perceptual experiences. They have described a variety of cases in which it supposedly happens, but they have said little about *how* it supposedly happens. Some high-levelists have suggested an account in terms of cognitive penetration. I take it their basic proposal is that high-level properties come to be represented in our perceptual experiences as a result of our cognitive states directly influencing or altering the contents of our perceptual experiences.

Even though some high-levelists have suggested that our perceptual experiences come to represent high-level properties via cognitive penetration, the general consensus seems to be that high-levelists aren't committed to cognitive penetration, and that there can be high-level representations without cognitive penetration. However, I take it, that assumes there's a plausible account of how our perceptual experiences come to represent high-level properties *not* in terms of cognitive penetration. My goal in this section is to argue there is no such theory in the offing.

Broadly speaking, there are only two types of accounts high-levelists could give: internalist or externalist. Internalists about perceptual representation (i.e. the contents of our perceptual experiences) contend that what our perceptual experiences represent is entirely determined by (or supervenes on) our internal states or intrinsic features. And that means, for internalists, whether and what (high-level) properties our perceptual experiences represent is entirely determined by our internal states. And that also means, for internalists, there can't be differences in what properties our perceptual experiences represent, unless there are also differences in our internal states or intrinsic features.

So, an internalist account of how subjects' perceptual experiences come to represent (different) high-level properties in the contrast cases would have to explain (differences in) high-level representations purely in terms of (differences in) subjects' internal states. As far as I can see, there are only three possible types of internal states such an account could appeal to: innate features of subjects' perceptual systems, noncognitive inputs to subjects' perceptual systems, and cognitive inputs to subjects' perceptual systems. However, the account can't appeal to *cognitive* inputs to subjects' perceptual systems without also appealing to cognitive penetration. So, for there to be an internalist account of high-level representations *not* in terms of cognitive penetration, it would have to be just in terms of innate features of subjects' perceptual systems and/or noncognitive inputs to their perceptual systems.

So the question is whether there is a plausible account just in terms of innate features of and non-cognitive input to subjects' perceptual system(s) that could explain their perceptual experiences representing different high-level properties in the contrast cases. I don't think so. The account needs to explain how subjects' perceptual experiences could have represented different high-level properties in the contrast cases. For example, it would have to explain how your visual experience could have failed to represent the property of being a pine tree before you learned to recognize pine trees, but then came to represent it after you had learned. However, as I mentioned, an internalist account can't explain subjects' perceptual experiences representing different properties, high-level or otherwise, without appealing to some sort of difference(s) in subjects' internal states. And subjects in the contrast cases don't seem to differ with respect to either the innate features of or noncognitive inputs to their perceptual systems. But the contrast cases, by design, are pairs of cases that are nearly identical – only differing from one another with respect to some aspect(s) of subjects' cognition. When it comes to Siegel's contrast cases, for example, we are supposed to be imagining that the

cases are identical apart from the fact that you have the ability to recognize pine trees in one case, but not the other.

But if the only differences between the contrast cases is with respect to aspect(s) of subjects' cognition, then there are no differences in the innate features of or noncognitive inputs to their perceptual systems to which that an internalist account could appeal to explain subjects' perceptual experiences representing different high-level properties in those cases (e.g. your perceptual systems have the same innate features before and after you learned to recognize pine trees, and if, in fact, the difference(s) in your recognitional ability is the only difference between the cases there is, then the noncognitive inputs to your perceptual systems are also the same before and after you learned to recognize pine trees). To put it differently, since subjects in the contrast cases are the same with respect to the innate features of and noncognitive inputs to their perceptual systems, it's not possible for there to be an account just in terms of innate features and noncognitive inputs that would be capable of explaining subjects' perceptual experiences representing different high-level properties.

But suppose we were to allow that the contrast cases differ from one another with respect to more than just some aspect(s) of subjects' cognition – suppose we were to allow that the contrast cases differ with respect to both innate features of and noncognitive inputs to subjects' perceptual systems. You might wonder whether there could be an internalist account under those conditions. The question then is whether differences in innate features of and /or noncognitive inputs to subjects' perceptual systems could plausibly explain how their perceptual experiences came to represent different high-level properties. Once again, I think the answer is no.

We have innate abilities to perceptually represent properties like shapes, colors, tones, and textures. That is to say that the innate features of our perceptual systems are such that they can process information coming in through our sense organs and produce perceptual experiences representing those properties. Consider, for example, visual representations of color properties. We

don't have to learn anything, know anything, or do anything for color properties to be represented in our visual experiences. If our eyes are open and the lights are on we can't help but have visual experiences that represent color properties – even if we would prefer not to.

However, we don't have any innate abilities to perceptually represent high-level properties (even high-levelists agree with that). In all the cases purportedly involving high-level representations, subjects had to learn, know, and/or do something for their perceptual experiences to come to represent high-level properties. But if we have to learn it, know, or do something for our perceptual experiences to represent high-level properties, that suggests there can't be an explanation of differences in high-level representations in terms of innate features of subjects' perceptual systems, even if those innate features were different.

Could there be an explanation for subjects' perceptual experiences representing different high-level properties in terms of different noncognitive inputs into their perceptual systems? I take it that noncognitive inputs could only be three things: sensory inputs from the environment to the relevant perceptual system (e.g. for visual experiences, sensory inputs to the visual system), sensory inputs to the relevant perceptual system from other perceptual systems (e.g. auditory inputs to the visual system), or non-sensory (and noncognitive) inputs from any of the perceptual systems. So, the question is whether differences in any of those kinds of inputs could plausibly explain how subjects' perceptual experiences came to represent different high-level properties.

For differences in any of those kinds of inputs to explain subjects' perceptual experiences coming to represent different high-level properties, it would have to be possible for those kinds of inputs to explain high-level representations. However, I don't see how that would be possible. Consider your visual experience representing the property of being a pine tree. According to Siegel (2010) your visual experience came to represent that property as a result of your learning to recognize pine trees by sight, together with cognitive penetration. But could you have come to

represent the property of being a pine tree as a result of noncognitive inputs to your visual system? I don't see how. It's not as though your visual experience coming to represent that property is like your visual experience coming to represent a new color. Changes in sensory inputs to your visual system could certainly explain your visual experience coming to represent a new color (as well as other low-level properties). But how could changing (or adding) sensory inputs to your visual system explain your visual experience coming to represent the property of being a pine tree? Presumably, having the property of being a pine tree is to possess a certain kind of essence and/or a certain kind of history. But presumably, there is no sensory information about essences or histories. Those are not features that are perceptual systems have evolved to detect. In other words, our perceptual systems do not have the ability to detect something's essence or history. So I take it that no change to or additional sensory inputs could possibly explain your visual experience coming to represent the property of being a pine tree after you had learned to recognize pine trees.

Plausibly, our perceptual experiences can represent colors (and other low-level properties) even if we don't possess the relevant color concepts, because our perceptual systems evolved to detect such properties. And thus our perceptual systems can produce experiences that represent color properties with just inputs/information coming in from our sensory organs.

But presumably our perceptual systems did not evolve to detect the property of being a pine tree (or other high-level properties). According to Siegel, you come to perceptually represent that property (via cognitive penetration) because of learning how to recognize pine trees. But how could learning to recognize pine trees, which I take it is a cognitive matter, plausibly give you the ability to *visually represent* a new property, namely the property of being a pine tree?

Plausibly, in learning to recognize pine trees, you gained the concept PINE TREE and/or the disposition to apply your concept PINE TREE when undergoing a visual experience of a pine tree. But suppose you can't recognize pine trees and that you even lack the concept PINE TREE. It

doesn't seem like there's any way your perceptual system(s) just using information from your sensory organs, could produce a perceptual experience representing the property of being a pine tree. It's plausible that possessing high-level property concepts and being disposed to apply those concepts to your perceptual experiences are necessary for your perceptual experiences to come to represent high-level properties because your perceptual systems can't generate representations of high-level properties just using sensory information (poverty of the stimulus). However, if you possess high-level property concepts and abilities to apply those concepts to your perceptual experiences, then your perceptual systems can make use of that information to make up the difference (e.g. perhaps your perceptual systems can use those concepts, along with whatever comes along with your abilities to apply those concepts, to group or chunk sensory inputs in a way that makes it possible for your perceptual systems to generate experiences that represent high-level properties). The idea is that your perceptual system(s) take in sensory inputs/information, which allows them to generate representations of low-level properties, and then your perceptual system(s) use those representations together with your concept PINE TREE, to generate a perceptual experience that represents the property of being a pine tree.

Since subjects' cognitive states are internal states, there could be an internalist account appealing to cognitive penetration that plausibly explains high-level representations. But I conclude there is no internalist account that plausibly explains high-level representation without appealing to cognitive penetration.

Is there an externalist account of high-level representation that could plausibly explain subjects' perceptual experiences representing different high-level properties in the contrast cases? Externalists about perceptual representation contend that what our perceptual experiences represent is partly determined by (or minimally supervenes partly on) factors *external* to subjects of those experiences. The most popular externalist account of perceptual representation is *tracking*

representationalism or *tracking intentionalism*.[10] There are a few different versions of this view, but the basic thesis is that our perceptual experiences represent those features or properties (of our environments) that our perceptual states normally track (or track under normal conditions). So the question is whether there is a plausible account of high-level representations in terms of what our perceptual states normally track that can plausibly explain how subjects' perceptual experiences came to represent different high-level properties in the contrast cases.

The primary source of disagreement among proponents of the tracking account is over what "normally tracks" or "tracks under normal conditions" means. One of the main competing views, most associated with the work of Fred Dretske, cashes out the meaning of "normally tracks" in teleological terms (Dretske, 1995). On this view, the features or properties that our perceptual states normally track are understood to be the ones that our perceptual states evolved or were designed by nature to track. And since our perceptual experiences only represent those features or properties our perceptual states normally track, according to this view, our perceptual experiences only represent those properties our perceptual states evolved or were designed to track.

Teleological interpretations of the tracking account, however, have been criticized for (among other things) having the consequence that Swampman – an accidentally created micro-physical duplicate of a human being – has no perceptual experiences. Since Swampman was created by accident, his perceptual states did not evolve and were not designed by nature to track anything. And thus, on this interpretation, Swampman's perceptual states don't track anything, and thus he doesn't have any perceptual experiences at all!

Even if we set aside worries about the teleological version of the tracking account, it is implausible high-levelists can explain subjects coming to perceptually represent high-level properties

[10] For an overview and discussion of the key versions of the tracking views see Bourget and Mendelovici, 2014.

(in the contrast cases). On teleological tracking account, our perceptual experiences represent those properties that our perceptual states evolved or were designed by nature to track. So for there to be a teleological tracking account of high-level representations, it would have to be the case that our perceptual states evolved to track high-level properties. But *no one*—not even high-leveliststhinks that our perceptual states evolved to track high-level properties. Plausibly, the properties our perceptual states were evolved to track are identical to the properties that we have innate abilities to perceptually represent—and those are low-level properties, like colors and shapes.[11] And if our perceptual states did not evolve to track high-level properties, there can't be a teleological tracking account of high-level representations.

An alternative to the teleological view, most associated with the work of Michael Tye, cashes out "normally tracks" in counterfactual terms. On Tye's view, our perceptual experiences represent those features or properties they would track under optimal conditions, where optimal conditions are those conditions under which an experience representing property P would be tokened in subjects if and only if P were the case, and because of P (Tye, 2000).

So the question is whether there is an optimal conditions version of the tracking theory that could plausibly explain subjects' perceptual experiences coming to represent high-level properties (in the contract cases). At least on the face of it, this seems more promising than that teleological version. After all, high-levelists claim that subjects come to perceptually represent high-level properties as a result of learning to recognize them by sight. So, for example, your perceptual experiences purportedly come to represent the property of being a pine tree after you've learned to recognize pine trees by sight. But presumably having the ability to recognize pine trees by sight means that you can reliably identify pine trees when they are around, and that you don't (often) mis-

[11] With the possible exception of facial and linguistic properties. But it's unclear those should count as high-level properties. See section 4.4.

identify non-pine trees (e.g. "recognize" birch trees as pine trees). But if that's the case, your perceptual experiences will represent the property of being a pine tree when you are in the presence of pine trees and because of pine trees. And that sounds awfully lot like the optimal conditions requirement.

The problem is that presumably your perceptual experiences won't just represent the property of being a pine tree in the presence of (and because of) pine trees. But also in the presence of (and because of) *twin* pine trees and other pine tree *look-alikes*. In other words, your perceptual representation of the property of being a pine tree fails to meet the "only if" side of the requirement. So I conclude that it's also implausible there is an optimal conditions tracking account of high-level representations.

Since tracking theories are not the only externalist theories of representation, I haven't shown definitively that there can't be a plausible externalists account of high-level representations. However, I think I've done enough to shift the burden of proof onto high-levelists — if there is a plausible account (not in terms of cognitive penetration) for high-level representations, then what is it?

1.1.2 Objections and Replies

First objection. Premise (4) of your argument is false. There is substantial empirical evidence that we have innate abilities to perceptually represent some high-level properties – namely, facial and linguistic properties. So there could be a plausible explanation for subjects' perceptual experiences representing different high-level properties in terms of differences in the innate features of their perceptual systems.

Reply. I accept that there is substantial empirical evidence that we have innate abilities to perceptually represent facial and linguistic properties. However, for this objection to work, we also have to accept that facial in linguistic properties are high-level. And I'm not prepared to do that.

If we assume that "high-level properties" and "low-level properties" name categories that are both mutually exclusive and collectively exhaustive, then facial and linguistic properties have to be one or the other. People engaged in debate over high-level representation have tended to assume they are high-level. But I see no reason to accept that.

Whether facial and linguistic properties are high-level or low-level depends on how those categories are defined. Unfortunately, the typical way this has been done in the literature is via lists of canonical or paradigmatic examples of each. Low-level properties are defined as properties like colors and shapes. And high-level properties are defined as properties like being a duck and being a pine tree.

On the face of it, there does seem to be a difference between properties like shape and color, on the one hand, and properties like being a pine tree and being a mate in six, on the other. But it's not at all obvious which of these two categories facial and linguistic properties fall into.

Those definitions may be fine for some purposes. However, unless there's a clear way to "go on" or continue the list, they are virtually useless when it comes to categorizing new properties or properties not already on one of the lists.

Aware of the inherent flaws of "list definitions", a couple of philosophers have attempted to define the categories by articulating their essential feature(s). Here is one such attempt.

> Some properties are high-level compared to quintessential low-level properties in the sense that the awareness of the former depends on neural processing of the latter, for instance our awareness of the high-level property of being a face or the property of gazing in a particular direction depends on the neural processing of lower-level properties, such as shape, texture, and direction and brightness (Brogaard & Chomanki, 2015: 473-4).

That way of defining the categories does give us a principled way of categorizing properties not already on the canonical lists. And it also has the result that facial and linguistic properties are high-level. The problem is that this way of drawing the distinction or defining the categories can't be right.

It seems to me that one desideratum on any definitions of low-level and high-level properties is that they have to correctly categorize or "get right" most or all of the properties already on the lists – i.e. that definitions must count as low-level those properties on the lists of paradigmatic low-level properties, and must count as high-level those properties on the list of paradigmatic high-level properties. The proposed definition is good to the extent that it achieves this goal, and bad to the extent that it fails.

But Brogaard and Chomanski's definition does poorly in this regard. Presumably, our perceptual awareness of color properties depends on the neural processing associated with our perceptual awareness of shape properties. We can't visually represent a color without visually representing something (an object, an area) that is extended space, and we can't visually represent things as extended in space without representing shapes. So it seems to follow on Brogaard and Chomanski's definitions that color and shape properties are high-level! But that's unacceptable. Color and shape properties are the paradigmatic low-level properties.

And just in case you're the least bit tempted to accept that shape and color properties are high-level, I want to emphasize that the problem here is not just that that definition fails to respect the canonical lists. If shape and color properties are categorized as high-level, since everyone already agrees that our perceptual experiences represent color and shape properties, this way of defining the properties has the consequence that the question of whether our perceptual experiences represent

high-level properties is no longer interesting. However, it leaves unanswered the question as to whether our perceptual experiences represent properties like being a duck and being a pine tree. [12,13]

So where does that leave us? As I said, for the high-levelist's objection to premise (4) to work, facial and linguistic properties would have to be high-level. But I've argued whether facial and linguistic properties are high-level depends on how "high-level properties" is defined. And I've argued none of the proposed definitions for "higher-level properties" are both plausible and support the claim that facial in linguistic properties are high-level.

I am satisfied with the above reply, but just for good measure, I also want to offer a reply to the objection that doesn't rely on quibbling over the definition of "high-level properties". My second reply is that even if we grant that facial and linguistic properties are high-level (and that we have innate abilities to perceptually represent those properties), at most, that shows my argument needs to be restricted to perceptual representations of non-linguistic and non-facial high-level properties.

---

[12] You might be wondering whether Brogaard and Chomanski's definitions can be rescued with a minor tweak – namely changing 'dependence' to 'asymmetric dependence'. High-level properties would then be defined as properties, our perceptual awareness of which asymmetrically depends on the neural processing of low-level properties. On this definition, the properties of being a duck and being a pine tree would come out as high-level since, presumably, our perceptual awareness of those properties depends on the neural processing of color and shape properties. But presumably our perceptual awareness of color and shape properties does not depend on the neural processing of the property of being a duck or being a pine tree. But the advantage of this definition is that it gets shape and color properties right because the dependencies between color and shape properties are symmetric.

Even though the definition in terms of asymmetric dependence gets color and shape properties right – namely, as low-level properties, it seems to get other canonical low-level properties, like motion and depth properties, wrong – according to this definition, motion and depth properties are high-level. Presumably, we can't be perceptually aware of motion or depth properties without being perceptually aware of something in motion or at a distance. And presumably we can't be perceptually aware of something in motion or at a distance without being perceptually aware of color and shape properties. If that's right, then presumably our perceptual awareness of motion and depth properties asymmetrically depends on the neural processing of color and shape properties.

[13] A proposal I find compelling draws the distinction between low-level and high-level properties in terms of innateness. To put it differently, the proposal I prefer uses innateness as the dividing line between low-level and high-level properties. The low-level properties are the properties that our perceptual experiences represents as a result of innate features alone. The high-level properties, in contrast, are the properties that our perceptual experiences represent, but not as a result of innate features alone. But I take it high-levelists won't be happy with this way of drawing the distinction because it puts causal, linguistic, and facial properties on the low-level side.

Much of the debate over high-level representations has centered on perceptual representations of natural and artifactual kind properties (e.g. being a duck, being a pine tree, and being a stethoscope). And everyone (including high-levelists) agrees that we don't have innate abilities to perceptually represent properties like that. If our perceptual experiences come to represent those properties, it's as a result of subjects learning something, coming to know something, and/or doing something. And my claim is that we can't gain abilities to perceptually represent high-level properties (i.e. properties we don't have innate abilities to perceptually represent) without cognitive penetration.

1.5 CONCLUSION

Two of the most popular debates in the philosophy of perception in recent years have been over whether high-level properties are represented in our perceptual experiences (i.e. the existence of high-level representations), and whether our cognitive states can directly influence or alter our perceptual experiences (i.e. the existence of cognitive penetration). Despite some overlap between these debates (e.g. when high-levelists have appealed to cognitive penetration as an explanation for subjects' perceptual experiences coming to represent high-level properties), the general sense has been that these issues are entirely independent – i.e. that there can be high-level representations without that cognitive penetration, and cognitive penetration without high-level representation.

I don't deny that there can be cognitive penetration without high-level representations. In fact, the alleged empirical evidence of cognitive penetration that people have found most compelling comes from experiments purportedly finding that subjects' perceptual representations of *low-level* properties had been directly influenced or altered by their cognitive states. For example, experiments purportedly found that children's needs and desires had directly influenced or altered their visual

experiences of pennies, dimes, nickels, quarters, and half dollars—made them look about 30% bigger (in diameter) than they actually were.

However, my aim in this essay has been to argue that the existence of (some) high-level representations depends on the existence of cognitive penetration. If my arguments are successful, they show that the existence of cognitive penetration is necessary for the existence of (some) high-level representations.

1.6 REFERENCES

[Bayne 2011] Bayne, Tim: Perception and the Reach of Phenomenal Content. In: Hawley, Katherine (Hrsg.) ; Macpherson, Fiona (Hrsg.): The Admissible Contents of Experience. Wiley-Blackwell, 2011

[Bermueacute; und Cahen 2015] Bermueacute; ; Cahen, Arnon: Nonconceptual Mental Content. In: Zalta, Edward N. (Hrsg.): The Stanford Encyclopedia of Philosophy. Fall 2015. Metaphysics Research Lab, Stanford University, 2015

[Bourget und Mendelovici 2014] Bourget, David ; Mendelovici, Angela: Tracking Representationalism. In: Bailey, Andrew (Hrsg.): Philosophy of Mind: The Key Thinkers. Continuum, 2014, S. 209–235

[Brogaard und Chomanski 2015] Brogaard, Berit ; Chomanski, Bartek: Cognitive Penetrability and High Level Properties in Perception: Unrelated Phenomena? In: Pacific Philosophical Quarterly 96 (2015), Nr. 4, S. 469–486

[Bruner 1947] Bruner, Jerome: Value and Need as Organizing Factors in Perception. In: Journal of Abnormal and Social Psychology 42 (1947), S. 33–44

[Chudnoff forthcoming] Chudnoff, Elijah: Phenomenal Contrast Arguments for Cognitive Phenomenology. In: Philosophy and Phenomenological Research XCI (forthcoming), July, Nr. 1, S. 82–104

[Deroy 2013] Deroy, Ophelia: Object-sensitivity versus cognitive penetrability of perception. In: Philosophical Studies 162 (2013), S. 87–107

[Dretske 1995] Dretske, Fred: Naturalizing the Mind. MIT Press, 1995

[Firestone 2013] Firestone, Chaz: How "Paternalistic" Is Spatial Perception? Why wearing a Backpack Doesn't–and Couldn't–Make Hills Look Steeper. In: Perspectives on Psychological Science 8 (2013), Nr. 4, S. 455–473

[Firestone und Scholl 2014] Firestone, Chaz ; Scholl, Brian J.: "Top- Down" Effects Where None Should Be Found: The El Greco Fallacy in Perception Research. In: Psychological Science 25 (2014), Nr. 1, S. 38– 46. – URL http://pss.sagepub.com/cgi/content/abstract/25/1/38

[Firestone und Scholl 2016] Firestone, Chaz ; Scholl, Brian J.: Cognition does not affect perception: Evaluating the evidence for "top- down" effects. In: Behavioral and Brain Sciences 39 (2016)

[Fodor und Pylyshyn 1981] Fodor, Jerry A. ; Pylyshyn, Zenon W.: How Direct is Visual Perception? Some Reflections on Gibson's 'Eco- logical Approach'. In: Cognition 9 (1981), Nr. 2, S. 139–96

[Hawley und Macpherson 2011] Hawley, Katherine (Hrsg.) ; Macpherson, Fiona (Hrsg.): The Admissible Contents of Experience. Wiley-Blackwell, 2011

[Koksvik 2015] Koksvik, Ole: Phenomenal Contrast: A Critique. In: American Philosophical Quarterly 52 (2015), October, Nr. 4, S. 321–334

[Kuhn 1962] Kuhn, Thomas S.: The Structure of Scientific Revolutions. The University of Chicago Press, 1962

[McClelland 2016] McClelland, Tom: Gappiness and the case for liberalism about phenomenal properties. In: The Philosophical Quarterly 264 (2016), S. 536–558. – URL http://pq.oxfordjournals.org/content/early/2016/01/08/pq.pqv128.abstract

[Peacocke 2001] Peacocke, Christopher: Does Perception Have a Non- conceptual Content? In: Journal of Philosophy 98 (2001), Nr. 5, S. 239–264

[Price 2009] Price, Richard: Aspect-Switching and Visual Phenomenal Character. In: Philosophical Quarterly 59 (2009), Nr. 236, S. 508–518

[Pylyshyn 1999] Pylyshyn, Zenon. W.: Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception. In: Behavioral and Brain Science 22 (1999), Nr. 3, S. 341–365

[Raftopoulos 2009] Raftopoulos, Athanassios: Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy? Mit Press, 2009

[Siegel 2006] Siegel, Susanna: Which Properties Are Represented in Perception? In: Gendler, Tamar S. (Hrsg.) ; Hawthorne, John (Hrsg.): Perceptual Experience. Oxford University Press, 2006,

S. 481– 503

[Siegel 2007] Siegel, Susanna: How Can We Discover the Contents of Experience? In: The Southern Journal of Philosophy XLV (2007), S. 127–142

[Siegel 2010] Siegel, Susanna: The Contents of Visual Experience. Oxford University Press, 2010

[Siegel 2016] Siegel, Susanna: The Contents of Perception. In: Zalta, Edward N. (Hrsg.): The Stanford Encyclopedia of Philosophy. Winter 2016. Metaphysics Research Lab, Stanford University, 2016

[Siewert 1998] Siewert, Charles: The Significance of Consciousness. Princeton University Press, 1998

[Stokes 2013] Stokes, Daniel: Cognitive Penetrability of Perception. In: Philosophy Compass 8 (2013), Nr. 7

[Tye 2000] Tye, Michael: Consciousness, Color, and Content. MIT Press, 2000

[Wu 2013] Wu, Wayne: Visual Spatial Constancy and Modularity: Does Intention Penetrate Vision? In: Philosophical Studies 165 (2013), Nr. 2, S. 647–669

[Zeimbekis 2013] Zeimbekis, John: Color and Cognitive Penetrability. In: Philosophical Studies 165 (2013), Nr. 1, S. 167–175

Chapter 2: Evidence of Perceptual Adaptation is not Evidence of Cognitive Penetration

2.1 INTRODUCTION

Over the last several years, a number of philosophers and psychologists have endorsed the thesis that there's *cognitive penetration*, i.e. the thesis that our cognitive states can (and sometimes do) directly influence or alter our perceptual experiences (i.e. the way things look, sound, smell, taste, and feel to us). Even though no one thinks that it's *impossible*[14], the thesis that there *is* cognitive penetration (i.e. that it actually exists) is both radical and controversial.

There are several reasons for that. Among them is the fact that the existence of cognitive penetration is inconsistent with our best theories about how human perception works (Marr, 1982). It also conflicts with assumptions built into current computer models of perception (in particular vision) according to which cognition plays no role (i.e. don't include cognitive variables) (Firestone and Scholl, 2016). The existence of cognitive penetration also doesn't square with the fact that we continue to experience optical illusions (e.g. the Phi effect, the Müller-Lyer lines, and the Ames room) long after we've learned both that the illusions are not as they appear, and how they really are (e.g. that the Müller-Lyer lines are really the same length) (Fodor, 1984). It also seems opposed to the fact that we can't change our perceptual experiences simply by changing our minds. We can change our minds, so if there's cognitive penetration and our minds can change our perceptual experiences, then why can't we change our perceptual experiences simply by changing our minds?

[14] No one denies there *could* have been cognitive penetration. It seems both metaphysically and nomically possible. What people have argued is that the existence of cognitive penetration is inconsistent with our actual mental architecture. But, presumably, our mental architecture could have been different than it actually is.

Proponents of cognitive penetration don't really deny that there are reasons—perhaps even very good reasons—to doubt there is cognitive penetration. They just think those reasons are outweighed by what they take to be compelling empirical evidence of its existence. Given that empirical evidence of cognitive penetration is the only reason we have to believe it exists, and we have lots of reasons to believe it doesn't, we need to be asking two questions: just how compelling is empirical evidence of cognitive penetration and is it compelling enough to override all the reasons we have to doubt? Ultimately, here's where I come down on those two questions: not very and definitely not.

Most of the empirical evidence that proponents of cognitive penetration have presented in defense of their view is from studies reportedly finding *top-down effects* on subjects' perceptual experiences, i.e. effects on their perceptual experiences that were a result of direct influence from their cognitive state(s). Among these are a study that reportedly found "valuable" coins looked bigger to children than worthless cardboard discs of the same size as a result of influence from their desires (Bruner and Goodman, 1947), a study that reportedly found that heart-shaped cutout looked redder to subjects than oval-shaped cutouts of the same color as a result of direct influence from their color expectations (discussed in McPherson, 2012), and a study that reportedly found that African-American faces looked darker to subjects than Caucasian faces with the same average luminance as a result of direct influence from subjects' racial stereotypes (Levin and Banaji, 2006).[15]

[15] It's worth noting that these reported "findings" were not the observations or data collected from the studies. The observations for data from these studies concern subjects behaviors and self reported judgments – e.g. subjects' judgments of the size of coins were larger than their judgments of the worthless cardboard discs. Researchers then inferred from subjects behaviors and/or judgments, that they must have had different perceptual experiences. But because of the similarities in the experimental settings, researchers inferred that the (alleged) differences in subjects' perceptual experiences must have been a result of differences in their cognitive states.

There's no question that the top-down effects reportedly found in these studies would be cases of cognitive penetration (i.e. they meet all four of the conditions I discussed in the last section). So if it could be established that those top-down effects were *actually found*, that wouldn't just be compelling evidence of cognitive penetration, that would be proof! So I take it the reason these studies are presented as *evidence* of cognitive penetration, and not proof of it, reflects uncertainty with respect to whether their reported findings hold up—i.e. whether top-down effects were actually found.

And opponent of cognitive penetration have effectively argued this empirical evidence is unconvincing precisely because it's implausible that top-down effects were actually found (e.g. Bitter, 2014; for a comprehensive treatment see, Firestone and Scholl, 2016; Gross et al, 2014; O'Callaghan, 2010; Zeimbekis, 2013). Broadly speaking, they have challenged these reported findings in two ways: either by denying there really was the reported difference in subjects' perceptual experience (e.g. that coins *looked* bigger than worthless cardboard disks of the same size), or by denying the reported differences in subjects' perceptual experience were a result of direct influence from their cognitive state(s) (e.g. direct influence from children's' desires for objects of social value).

But some of the empirical evidence they discuss is from studies that have reportedly found, what I'll call, *adaptational effects* on subjects' perceptual experiences, i.e. effects on their perceptual experiences that arose after (or perhaps over) an extended period of time and were a result of durable changes in the way perceptual systems process information (e.g. sensory inputs) or, what we might call, *processing adaptations*. Proponents of cognitive penetration have also claimed compelling empirical evidence of cognitive penetration has come from studies that have reportedly found adaptational effects on subjects perceptual experiences. Among these are a study that reportedly found that after subjects had worn a pair of inverting glasses (continuously) for almost two weeks, the upside-down-look of everything (which the glasses had initially induced) had faded away—i.e.

everything looked right-side-up again (Churchland, 1988), a study that reportedly found that subjects made faster and better visual discriminations among a set of target stimuli after numerous intensive training sessions on those stimuli (Cecchi, 2014), and a study that reportedly found that the subjects in the study from non-Western cultures were less susceptible to certain visual illusions (e.g. the Müller-Lyer lines) than the subjects in the study from Western cultures (McCauley and Henrich, 2006).

Unlike the empirical evidence coming from studies reportedly finding top-down effects, the empirical evidence coming from studies reportedly finding adaptational effects has largely flown under the radar and escaped careful scrutiny. In what follows, I will argue the empirical evidence coming from studies adaptational effects is even *less* compelling than the empirical support coming from studies that have reportedly found top down effects.

The remainder of this essay will proceed as follows. In section 2, I will say more about cognitive penetration. In section 3, I will discuss the empirical evidence from studies reportedly finding adaptational effects, in particular the inversion lenses experiment. Even if we assume that those reported findings hold up (i.e. adaptational effects were found in those studies), that's only empirical evidence of cognitive penetration if adaptational effects are cases of cognitive penetration. But in section 4, I argue that's completely implausible. In section 5, I consider the possibility that adaptational effects are cases of weak cognitive penetration. In section 6, I show how my reasoning applies to two further empirical studies reportedly finding adaptational effects. I finish with a few concluding remarks (Section 7).

2.2 COGNITIVE PENETRATION

Is there cognitive penetration? To answer that question we first need to know what cognitive penetration is. We know that it's a kind of cognitive influence on perception, or from the other side,

a kind of effect in perception resulting from cognitive influence. But there is a lot of disagreement over how to define the phenomenon more precisely (Stokes, 2015; Wu, 2013; Macpherson, 2012)—and, unfortunately, the result of that has been that there are about as many different definitions of cognitive penetration as there are people engaged in the debate over its existence.

But the significance of the disagreement over how to define cognitive penetration should not be overblown. Everyone knows that definitions are hard. But as Supreme Court Justice Potter Stewart once said regarding hard-core pornography, and I'm paraphrasing, "I may not be able to define it, but I know it when I see it." And despite the disagreement over the definition, there has, in fact, been a fair amount of agreement between people's judgments about whether particular effects on perceptual experience would count as cases of cognitive penetration. So instead of trying to define the phenomenon, I will try to articulate the central or core features of cases of cognitive penetration.

I take it that, at a minimum, cases of cognitive penetration are (or would be) effects on perceptual experiences that are, in some sense, a result of influence from cognition, where an *effect* on perceptual experience is an aspect (or content) of a perceptual experience that has been changed, altered, or is different than it otherwise would be (i.e. if there had been no influence from cognition).

But there are all sorts of effects on our perceptual experiences that are, in some sense, a result of influence from cognition, but are definitely *not* cases of cognitive penetration. For example, suppose I cup my hand behind my ear *because* I want to hear you better and believe that cupping my hand behind my ear will enable me to hear you better. And suppose it does, in fact, alter (or enhance) my auditory experience. In that situation, the change in my auditory experience was, in some sense, the result of influence from my cognitive states — namely my desire to hear you better and my belief about how to accomplish that. But *no one* thinks that effect on my auditory experience would count as a case of cognitive penetration.

There is some disagreement among people as to *why* it wouldn't count. But the reason has something to do with the *indirect* or *mediated* nature of the influence from my cognitive states. My cognitive states only contributed to the change in my auditory experience *by* bringing about a change in a non-cognitive dimension of the situation – namely, the position of my hand behind my ear. But it was that change in the non-cognitive dimension of the situation that was *directly* or *immediately* responsible for the change in my auditory experience. To put it differently, it wasn't my cognitive states *per se* that did the relevant work in bringing about the affect on my auditory experience. The relevant work was done by the change in the position of by hand. My cognitive states were neither necessary nor sufficient for the effect on my auditory experience. If a muscle spasm, and not my belief and desire, had caused my hand to cup behind my ear there still would've been a change in my auditory experience. But if my hands had been restrained, my belief and desire wouldn't have changed anything.

The upshot is that to be a case of cognitive penetration, effects on perceptual experiences have to be a result of *direct* or *unmediated* influence(s) from cognition. I'll call this the *directness requirement*.
It also generally thought that to be cases of cognitive penetration, the resulting effects on perceptual experiences have to "make sense" given the cognitive state(s) that did the influencing—i.e. be semantically or logically related to the cognitive states doing the influencing. Here is how Zeon Pylyshyn states, what I'll call, the *semantic constraint*.

> …if a system is cognitively penetrable then the function it computes is sensitive, in a semantically coherent way, to the organism's goals and beliefs, that is, it can be altered in a way that bears some logical relation to what the person knows. (Pylyshyn, 1999: 343)

The semantic constraint ensures there's an intelligible connection between the resulting effects on perceptual experiences and the cognitive states doing the influencing. Fiona MacPherson helpfully illustrates the point with the following example.

> I believe that today is the day of an important exam. This belief causes stress and brings on a migraine. The migraine causes disturbances to my vision and I now experience flashing lights at the periphery of my visual field in addition to experiencing the scene in front of me. In this case, my visual experience has been altered on account of my belief but there is no intelligible connection between the content of the belief – my exam is today – and the content of my visual experience – the apparent flashing lights…(MacPherson, 2012: 26)

No one thinks that the flashing lights in this scenario would count as a case of cognitive penetration. Now, arguably, that's because the flashing lights also failed to meet the direct this requirement — the influence from cognition is mediated by the migraine. However, presumably even if we were to relax the directness requirement – and allow (some) effects on perceptual experience resulting from indirect cognitive influence to count as cases of cognitive penetration — the flashing lights *still* wouldn't count as a case of cognitive penetration. And we need the semantic constraint to get that result.

At this point we've covered what are widely considered to the core features of cases of cognitive penetration. And taken together they tell us that a case of cognitive penetration would have to meet the following four conditions:

(i) be an effect on a perceptual experience;
(ii) that was the result of *cognitive* influence(s);
(iii) where the cognitive influence(s) is *direct* or *unmediated* (e.g. by changes in the stimuli, sensory organs, or attention);
(iv) and where the effect is *semantically* or *logically* related to (the content of) the cognitive state(s) doing the influencing[16].

The claim that there's cognitive penetration, then, can be understood as the claim that *there are actual cases of cognitive penetration*.

[16] Adapted from (Bitter, 2014:4)

2.3 THE EMPIRICAL EVIDENCE OF ADAPTATIONAL EFFECTS

Most of the empirical evidence of cognitive penetration comes from studies reportedly finding top-down effects. There's no question that top-down effects would be cases of cognitive penetration.[17] So *finding* top-down effects wouldn't just be empirical evidence that there's cognitive penetration, it would be proof. But the problem with empirical studies reportedly finding top-down effects is that their reported findings don't hold up.

But proponents of cognitive penetration have also claimed empirical evidence comes from studies reportedly finding *adaptational effects* on subjects' perceptual experiences, i.e. effects on perceptual experiences that are a result of changes in the way perceptual systems process information. One of these studies that proponents of cognitive penetration have found particularly compelling is what I'll call *the inverting lenses experiment*. Here's Paul Churchland's discussion of the study.[18]

> Recall the effects of chronically worn 'inverting lens' on the visual perception of normal humans. Such lenses have the effect of inverting the orientation of all visual information relative to the body's tactile and motor systems. In short, they turn the visual world upside down…
>
> The initial effect [of wearing the lenses] is profoundly disorienting, but with little more than a week's practice, subjects adjust to the new perceptual regime. The subjects are not confined to a chair or bed for the duration of the experiment, but are forced by practical necessity to continue to interact with familiar objects and to engage in the normal forms of motor behavior. The result is that the subjects slowly manage to recoordinate their vision with the rest of their sensory and motor systems, and the illusion of the world's being upside down is said to fade away, all on a time scale of roughly a week…
>
> In similar experiments on animals, training produces a reversal in the character of what one might have presumed to be endogenously specified reflexes, such as the vestibular-ocular reflex, which directs one's eyes, when fixated on a target, to move

[17] The fact that top-down effects are clear cases of cognitive penetration is no surprise. The studies reportedly finding top-down effects were set to look for cognitive penetration.

[18] Dustin Stokes (XXX) and Robert McCauley & Joseph Henrich (XXX) also mention this study in arguing for cognitive penetration.

> an appropriate amount or right in order to compensate for head movements in the opposite direction. Here the brain seems literally to rewire the relevant neural mechanism under the pressures imposed by left-right inverting lenses (Gonshor and Jones 1976).
>
> Cases like these are important, for they reflect the plasticity of some very deep 'assumptions' implicit in visual processing, such as the specific orientation of the visual world relative to one's other sense modalities and to one's motor systems. If assumptions as deep as these can be re-shaped in a week or two, then our perception begins to look very plastic and very penetrable indeed. (Churchland, 1988: 174-5)

When subjects first put on the inverting lenses (that had "the effect of inverting the orientation of all visual information relative to the body's tactile and motor systems"), everything apparently looked upside down (as you would expect). However, according to Churchland, it was found that after subjects had been wearing the lenses for roughly a week, things no longer looked upside down—"the illusion of the world being upside down is said to fade away".

As I mentioned in the introduction, one reason people have for doubting the existence of cognitive penetration is the fact that illusions persist even after we come to know that the illusions are not as they appear and how they really are. Take, for example, the Müeller-Lyer lines, I know full well that the lines are the same length, and I've known it for some time. And at this point, I've spent dozens or hundreds of hours reading about, thinking about, arguing about, and talking about the illusion. And yet, despite that, the illusion has not dissipated in the least—*the lines still look* different me!

The passage I quoted above is part of Churchland's response to the persistence of illusions objection to the existence of cognitive penetration. He presents the experiment as a way of undermining that objection. The inverting lenses experiment, he claims, shows empirically that illusions do fade away — just not quickly or easily.

As a response to the persistence of illusions objection, I think Churchland's response is wrongheaded and confused in a number of ways. But, for the purposes of this paper, I'm only

interested in this response insofar as it is an expression of his apparent belief — a belief that has subsequently been shared by several others — that the illusion fading away is empirical evidence of cognitive penetration.[19]

Opponents of cognitive penetration could respond to Churchland's claim that the illusion fading away is evidence of cognitive penetration in the same way they have responded to empirical evidence from studies reportedly finding top-down effects — namely, by denying the reported findings hold up. In this case, that would mean denying that the illusion really faded away after subjects had been wearing the inverting lenses for a week.

Churchland's presentation invites us to think that the illusion fading away was a straightforward and uncontroversial (if surprising) experimental observation. So you might not have thought denying that the illusion faded away was an option. But that's a mistake. The thesis that the

---

[19] Technically, Churchland and others have presented it as empirical evidence of *diachronic* cognitive penetration. Churchland says the debate over cognitive penetration has (implicitly) been focused on whether cognition directly influences perception *quickly*, over very short periods of time. But that's just about *synchronic* cognitive penetration. And he thinks that makes the discussion too narrow. He says the debate over cognitive penetration is a debate over whether cognition directly influences perception, *period.* And that means the debate is also about whether cognition influences perception *over a long period of time* (e.g. days, weeks, months, or even years), i.e. diachronic cognitive penetration.

As a conceptual matter, this is surely right. There's nothing problematic about the *concept* of diachronic cognitive penetration. Churchland and Stokes make a lot of this point. But as far as I know, so does everyone else. I don't know of anyone who argues diachronic cognitive penetration is conceptually incoherent. But remember the conceptual (metaphysical or nomic) possibility of cognitive penetration (diachronic or otherwise) is neither here nor there. No one denies that cognitive penetration is *possible.* The debate is about whether it's actual.

So technically, Churchland and others think the inverting lenses experiment is empirical evidence of diachronic cognitive penetration. Since the reported finding is that illusion faded away over the course of a week, if it was evidence of cognitive penetration, I suppose it would be evidence of diachronic cognitive penetration. However, I will argue the claim that the illusion fading away—and adaptational effects more generally—evidence of cognitive penetration, diachronic or otherwise, is wrong (and possibly confused). So the question of whether it evidence of diachronic, as opposed to synchronic, cognitive penetration is, largely a distraction. It makes no difference whatsoever until I consider whether adaptational effects could empirical evidence of *weak* cognitive penetration (Section 6). So until I get to that point, I'm going to ignore this distinction, and talk simply in terms of cognitive penetration.

illusion faded away after subjects had been wearing the lenses for a week is an *interpretation* of the data and observations collected from subjects verbal reports about their experiences, as well as improvements in their abilities (e.g. abilities to navigate their environment). To be fair, Churchland is not going his own way here, that is an interpretation shared by a number of people, including some psychologists involved with the study. However, that interpretation is far from mandatory and has been rightly criticized.[20]

Even though I'm not convinced that the illusion faded away, I don't think denying the illusion faded away is the best approach to responding to this evidence. The illusion fading away is a (purported) example of an adaptational effect, i.e. an effect on perceptual experience resulting from changes in the way perceptual system(s) process information (i.e. processing adaptations). Even though there is reason to doubt that an adaptational effect was found in this experiment, there's no question that adaptational effects exist. So I'm just going to assume that the illusion actually faded away. My goal is to argue that the illusion fading away isn't evidence of cognitive penetration because to be evidence of cognitive penetration the illusion fading away would have to be a *case* of cognitive penetration, and it's completely implausible that *any* adaptational effects are cases of cognitive penetration.

[20] Eric Schwitzgebel (unpublished blog post XXX) and Jesse Prinz (see Stokes ftntXXX) have both called into question that interpretation. For example, Switzgebel argues that subjects' verbal reports seem to support the conclusion that their visual experience was unchanged. Rather, subjects became accustomed to *things looking upside down*—e.g. learning to associate things looking to be up with things being near the ground. This sort of learned association can also explain why subjects were able to more skillfully navigate their environment after becoming accustomed to the lenses. In other words, none of the empirical findings from this experiment clearly support the claim that the illusion faded away.

2.4 ARE ADAPTATIONAL EFFECTS CASES OF COGNITIVE PENETRATION?

I take the illusion fading away – or any adaptational effect for that matter — can only be evidence of cognitive penetration if it is (or would be) a *case* of cognitive penetration. So in presenting the illusion fading away as empirical evidence of cognitive penetration, Churchland and others are claiming that it is a case of cognitive penetration.

At least first pass, claiming that adaptational effects, like the illusion fading away, are cases of cognitive penetration doesn't seem just *wrong*, it seems *confused*. Despite the ongoing debate over exactly what cognitive penetration is, it should be uncontroversial that cases of cognitive penetration, at minimum, have to be effects on perceptual experiences that were, in some sense, the result of cognitive influence. So adaptational effects can only be cases of cognitive penetration if they were, in some sense, the result of cognitive influence. But proponents of cognitive penetration say explicitly that adaptational effects are the result of *processing adaptations* or changes in the way subjects perceptual systems processed information.[21] And processing adaptations are not cognitive influences. And, at least on the face of it, have nothing to do with cognition. So, at least on the face of it, it's not even *possible* for adaptational effects to be cases of cognitive penetration. And if it's not even possible for adaptational effects to be cases of cognitive penetration, claiming they *are* cases seems confused.

Adaptational effects can only be cases of cognitive penetration if they were, in some sense, the result of cognitive influence(s). And on the face of it, it's not clear that's even possible, the fact that proponents of cognitive penetration present studies reportedly finding adaptational effects

[21] As I said, it's not clear that the illusion really did fade away. But if it did, I take it the only plausible way to explain that would be in terms of changes in the way subjects visual systems processed information. Assuming conditions were more or less the same at the beginning of the experiment, as they were at the end (e.g. subjects didn't begin standing on their heads), what else could explain everything changing from looking upside down to looking right side up?

without any argumentation is baffling. But that aside, you might wonder whether anything can be said on their behalf. In other words, is there any way the illusion fading away *could* be a case of cognitive penetration?

2.5 AN UNWARRANTED ARGUMENT

Is it possible for adaptational effects to be cases of cognitive penetration? Adaptational effects can only be cases of cognitive penetration if they, in some sense, result from (direct) cognitive influence. But adaptation effects, by definition, result from processing adaptation(s).[22] So I take it the only way adaptational effects could, in some sense, result from cognitive influence is if cognitive influence was somehow involved in bringing about the relevant processing adaptations (i.e. the processing adaptations that were responsible for bringing about the adaptational effects). Let's call processing adaptations brought about, at least in part, by cognitive influence *cognitively influenced (processing) adaptations*. The only way adaptational effects could be cases of cognitive penetration if they resulted from cognitively influenced adaptations.

So I assume the basis for proponents of cognitive penetration claiming that empirical studies reportedly finding adaptational effects, like the experiment reportedly finding that the illusion faded away, show there's cognitive penetration, is their un-argued-for conviction that the adaptational effects in the studies they point to are the result of cognitively-influenced processing adaptations. If we assume that's the case, then we can get the conclusion that the adaptational effects in question (e.g. the illusion fading away) are cases of cognitive penetration via the following simple argument::

(1) Adaptational effects are the result of cognitively influenced adaptations.
(2) If adaptational effects are the result of cognitively influenced adaptations, they are genuine cases of cognitive penetration.

[22] I'm assuming the claim that adaptational effects like the illusion fading away are result of processing adaptations is not up for grabs (if the illusion really did feel way, I take it that the only plausible explanation for that would be in terms of processing adaptations).

(3) So adaptational effects are genuine cases of cognitive penetration.

The problem, however, is that neither premise of this argument is warranted.

2.5.1 Re (1): Cognitively-Influenced vs. Bottom-up Adaptations

Let's first consider premise (1), according to which adaptational effects are the result of cognitively influenced processing adaptations. For that to be true, it has to be the case that the relevant changes in perceptual processing (i.e. those changes in processing responsible for bringing about the adaptational effects) were caused (or in some other way brought about) at least partly by subjects' cognitive state(s).

Broadly speaking, there are only two ways that could have happened: either the relevant processing adaptations were caused entirely by influence from subjects' cognitive states or the processing adaptations were caused partly by influence from subjects' cognitive states and partly by something else. Let's call cognitively influenced adaptations caused entirely by influence from cognition top-down adaptation, and cognitively influenced adaptations caused partly by influence from cognition and partly by something else mixed adaptations. For premise (1) to be true, it has to be the case that adaptational effects they point to were either the result of top-down or mixed adaptations.

Whether subjects' cognitive states were involved in bringing about the processing adaptations responsible for bringing about the adaptational effects they point to, is ultimately an empirical question. But I want to bracket the empirical question, for the moment, and talk about the space of possibilities.

I grant that it's *possible* that subjects' cognitive states were involved. But it's also *possible* that they were not. It's possible that the relevant processing adaptations were caused or otherwise brought about entirely by non-cognitive factors (e.g. sensory deprivation and training on perceptual stimuli).

Let's call processing adaptations brought about entirely by non-cognitive factors *bottom-up adaptations*. I take it it's at least *possible* that adaptational effects in question, like the illusion fading away, were entirely the result of bottom-up adaptations. And if it's possible that the adaptational effects in question were brought about by bottom-up adaptations, proponents of cognitive penetration are not entitled to premise (1) unless we have reason to think all processing adaptations are (in fact) cognitively-influenced (i.e. cognition is *always* involved) or we have some special reason to think that subjects cognitive states were involved in the particular experimental cases in question. Let's consider each of these options in turn.

It has long been known that the way our perceptual systems process information can (and does) change over time – i.e. that there are processing adaptations. For example, deprivation experiments done on animals have shown that long periods of sensory deprivation (especially during critical periods of development) have a profound impact on the animals' perceptual systems. In one experiment, kittens' eyelids were sewn shut for six months, depriving their visual systems of sensory stimulation during a critical period of development. Brain imaging and dissection revealed gross morphological differences between the brain areas associated with visual processing between the experimental and control kittens. And given the behavioral differences between these groups, the inference that these gross morphological differences resulted in (visual) experiential differences is well supported.

The evidence from animal experiments is corroborated by observational studies with human subjects. Some babies are born with opaque cataracts that effectively make them blind. Those cataracts can be surgically removed in infancy with little to no long-term side effects. However, there have been cases where the cataracts are not removed (due primarily to lack of funds for the surgery). These individuals then grow up blind—and thus their visual systems are deprived of sensory stimulation during critical periods of development. But here's where the relevant data comes in. A

select few of these individuals have had their cataracts removed in late adolescence or early adulthood (but after a childhood of blindness) . These individuals (currently, their numbers are in the double digits) constitute "natural" deprivation experiments on humans. And the results of those experiments are the same as they are in animals. Namely, neural imaging of the visual areas of their brains shows gross morphological differences. And their verbal reports and behaviors indicate that their visual experiences (after the cataracts are removed) are unlike the visual experiences of the rest of us (Dengenaar, 2014; Scholl, 1999).

And all that evidence is further corroborated by the results from perceptual training experiments done with humans. Neural imaging on "trained" areas of subjects perceptual systems before and after intensive periods of training, reveal significant structural changes within those areas. And the fact that subjects in these experiments improve on tasks like visual discrimination, suggests that the structural changes within their perceptual systems resulted in experiential differences (which allowed them to perform better on discrimination tasks).

I don't know of any reason to think that cognitive influences are *always* involved in bringing about processing adaptations. In fact, if anything, the evidence from sensory deprivation and perceptual training experiment suggests that processing adaptations are brought about entirely by bottom-up adaptations. And if there are some bottom-up adaptations, it's false that cognitive influence is always involved.

Do we have any reason to believe that cognitive influences had to be involved in the processing adaptations responsible for bringing about adaptational effects like the illusion fading away? Not that I can see. In fact, in the passage I quoted earlier, Churchland suggests the processing adaptations responsible for the illusion fading away were caused by "continue[d]…interact[ion] with familiar objects and engage[ment] in normal forms of motor behavior." He only mentions subjects' cognitive states to say this:

> When the lenses were first put on, and the world is made to appear upside down, the subjects are of course quite aware of what the lenses are doing. They may even know how they do it. But the illusion is not banished by the mere possession of this information. It would clearly be wrong, however, to draw from this any conclusion about the impenetrability of our visual processors. (Churchland, 1988:174)

For it to even be possible – never mind actual – that adaptation effects like the illusion fading away are genuine cases of cognitive penetration, it has to be the case that they were the result of cognitively influenced adaptations (i.e. either top-down or mixed adaptations). I can't categorically deny that they were. However, my point is that, proponents of cognitive penetration have given us no reason whatsoever to believe that they were. And since we also have no reason to believe that cognitive influences are always involved in bringing about processing adaptation, premise (1) of the argument is unwarranted. To put it another way, unless my opponents rule out the possibility that the illusion faded away as a result of bottom-up adaptations, they can't rely on premise (1).

2.5.2 Re (2): Cognitively-Based Processing Adaptations vs. Cognitive Penetration

But even if we assume that premise (1) is true (i.e. that the adaptational effects like the illusion fading away are the result of cognitively influenced adaptations), it doesn't follow that the adaptational effects are cases of cognitive penetration. If premise (1) is true, the adaptational effects meet the minimum required for being genuine cases of cognitive penetration – namely, they are effects on perceptual experience that are in some sense the result of cognitive influence. But adaptational effects can, of course, meet the minimum required without actually being cases of cognitive penetration. As I discussed in section 2, there are plenty of effects on perceptual experiences that are in some sense the result of cognitive influence that are not cases of cognitive penetration (e.g. effects on perceptual phenomenology that are the result of indirect cognitive influence). So the question now is whether adaptational effects resulting from cognitively influenced

adaptations are (or would be) genuine cases of cognitive penetration. To get from the claim that the adaptational effects in question were the result of cognitively influence adaptations the conclusion that adaptational effects are genuine cases of cognitive penetration, we need premise (2).

Even if we assume the adaptational effects in question resulted from cognitively influenced adaptations, those adaptational effects are only genuine cases of cognitive penetration and if they also meet the direct this requirement and semantic constraint. Do it the adaptational effects in question meet these further conditions? I'll consider each in turn.

2.5.2.1. The Directness Requirement

Suppose you're walking home late at night and start to fear you're being followed; to assuage your fear, you turn your head to look. Unless circumstances are highly unusual, after you turn your head, your visual experiences will be different (than it was before). Since you turned your head because you were afraid, there's a clear sense in which the effect on your visual experience was, at least in part, the result of influence from your cognitive states.

However, as I pointed out in section 2, it's widely agreed that effects on perceptual experience like this do not count as cases of cognitive penetration. The diagnosis for this is that such effects are only the result of indirect influence from cognition. In other words, in this situation, the influence from your cognitive states was indirect or mediated. Your fear only influenced your visual experience in the sense that it prompted you to turn around. But it was turning around that changed what you were looking at (i.e. your visual stimuli), and thus your visual experience. The influence from your cognitive states was neither necessary nor sufficient for the affect on your visual experience. In other words, that effect on your visual experience fails to meet the directness requirement.

If we maintain that directness requirement, adaptational effects can't be cases of cognitive penetration unless they, in some sense, resulted from direct influence from subjects' cognitive states. But that doesn't seem possible. Even if the adaptational effects were the result of top-down or mixed adaptations, it seems to me that the influence from subjects cognitive states would be indirect. At least on the face of it, adaptation effects resulting from top down or mixed adaptations result from cognition in the same sense that the effect on your visual experience resulted from cognition in the above scenario. And we've already seen that the effect on your visual experience in the above scenario does not count as a case of cognitive penetration.

Figure 1: Variations of (In)direct Cognitive Influence

Since the effect on your visual experience doesn't count as a case of cognitive penetration because of the indirect nature of the influence from cognition, unless we have some reason to believe the indirectness of the cognitive influences when it comes to adaptational effects is different, then even if they result from top-down or mixed adaptations, arguably, they could not be cases of cognitive penetration.

2.5.2.2. The Semantic Constraint

Even if we relax the directness requirement – and allow that some effects on perceptual experience resulting from indirect cognitive influence are cases of cognitive penetration – adaptational effects can't be cases of cognitive penetration unless they also meet the semantic constraint. For an effect on perceptual experience resulting from direct cognitive influence is only a

case of cognitive penetration if it is semantically or logically related to the cognitive states doing the influencing. For example, all the empirical studies reportedly finding top-down effects reported uncovering differences between subjects' perceptual experiences as a result of some sort of cognitive differences. For example, studies reportedly finding that faces with stereotypical African-American features looked darker (in skin tone) than faces with stereotypical Caucasian features (that were just as dark), that expecting bananas to be yellow made images of bananas look yellower, and expecting lips to be red made lip shaped cutouts look redder, that being depressed made things look darker, and even that wearing heavy backpack made hills look steeper. In each of these examples, there's a non-arbitrary connection between what the cognitive state(s) are about and the effect on perceptual experiences their influence brings about.

Processing adaptations are changes in the way perceptual systems process information. It's not clear that such changes have semantic content. And if they don't, there can't be a semantic connection between adaptational effects and either top-down or mixed adaptations. But perhaps we can look to the semantic contents of the influencing cognitive states

To summarize, if cases of cognitive penetration have to meet both the directness requirement and semantic constraint, arguably, adaptational effects wouldn't be cases of cognitive penetration even if they were the result of top-down or mixed adaptations. And since resulting from top-down or mixed adaptations is the only way adaptational effects could be the result of cognitively influenced adaptations, premise (2) is false.

2.6 WEAK COGNITIVE PENETRATION

I argued in the previous section that at adaptational effects like the illusion fading away are not cases of cognitive penetration even if they resulted from top-down or mixed adaptations because they fail to meet both that directness requirement and semantic constraint. However, both of those

conditions have called into question (e.g. Stokes, 2015; Wu, 2013). Since I don't want to adjudicate those dispute here; and it may still be interesting if there are adaptational effects that are the result of top-down or mixed adaptations. So let's just assume that cases of cognitive penetration need not meet those further conditions (or any other similar conditions).

Presumably, however, adaptation will affect that don't meet those further conditions would be weaker cases of cognitive penetration than ones that do. Let's call them weak cases of cognitive penetration. The question I want to explore in this section is whether adaptational effects resulting from top-down or mixed adaptations would be cases of weak cognitive penetration. I'll first address the conceptual question of whether such adaptational effects could be cases of weak cognitive penetration. I conclude that it is, indeed, possible. But ultimately, all that matters for our purposes is the empirical question of whether actual adaptational effects (e.g. the illusion fading away) are cases of the weak cognitive penetration. And I argue that we have no reason to believe that any of the actual adaptational effects are the result of top-down or mixed adaptations. As I suggested, the empirical evidence for processing adaptations suggests they are the result of entirely noncognitive factors (i.e. they are bottom-up adaptations).

2.6.1 The Conceptual Question

I take it whether adaptational effects resulting from cognitively influenced adaptations would count as cases of weak cognitive penetration depends on how subjects' cognitive states influenced the relevant processing adaptations.

In particular, I would argue that to count as cases of weak cognitive penetration, subjects' cognitive states had to play a direct role in bringing about the relevant processing adaptations. To allow for the possibility that adaptational effects are cases of cognitive penetration, we have to relax the directness requirement, and allow that some effects on perceptual experience resulting from

indirect cognitive influence our cases of cognitive penetration. However, we know for certain that some effects on perceptual experience resulting from indirect cognitive influence definitely would not even count as weak cases of cognitive penetration, e.g. effects resulting from a head turn. So the trick will be to distinguish between those cases where the indirect influence is problematic and those cases in which it is not. Even though I'm relaxing the directness constraint (i.e. allowing that indirect influences from thought can count as cognitive penetration), I'm not getting rid of it entirely. Some kinds of indirect influence from cognition definitely would not be cognitive penetration, that's not up for debate. So if thought only plays an indirect role in bringing about the processing adaptations in question, presumably any resulting adaptational effects would not be even cases of weak cognitive penetration. In other words, I'm only willing to grant that adaptational effects are (or would be) cases of weak cognitive penetration, if cognition played a direct role in bringing about the processing adaptations in question.

I'm willing to grant that adaptational effects resulting from top-down or mixed adaptations are cases of weak cognitive penetration as long as cognition played a *direct* role in bringing about the relevant processing adaptations. I don't have a precise way of cashing out the difference between cases where cognition plays a direct role and cases where cognition only plays an indirect role. But here's an intuitive gloss. Suppose I learn that daily eye exercises can, over time, improve my visual acuity by changing the way my visual system processes visual stimuli; and that prompts me to do the daily exercises. And suppose that doing those exercises does, in fact, change the way my visual system processes visual stimuli, which, in turn, improves my visual acuity. Changing my visual acuity, will change my visual experiences. In this scenario, my cognitive states definitely played a role in bringing about the processing adaptations responsible for improving my visual acuity. But my cognitive states only played an indirect role in bringing about the processing adaptations in question.

And arguably that's not enough for the adaptational effects in question to count as cases of weak cognitive penetration.

I'm willing to just grant that adaptational effects resulting from top-down adaptations would be cases of weak cognitive penetration. But when it comes to adaptational effects resulting from mixed adaptations, things are a bit more complicated. Let's draw a distinction between mixed adaptations where cognition played a direct role in bringing about the processing adaptations and mixed adaptations where cognition only played an indirect role in bringing them about. Let's call the former *direct mixed adaptations* and the latter *indirect mixed adaptations*. If what I've said is correct, adaptational effects resulting from direct mixed adaptations are (or would be) cases of cognitive penetration, but not adaptational effects resulting from indirect mixed adaptations. Presumably, adaptational effect resulting from indirect mixed adaptations would fail to be cases of weak cognitive penetration for the same reason that effects on perceptual experiences resulting from head turns fail to be cases of full blown cognitive penetration.

Since it's in principle possible that the processing adaptations responsible for the adaptational effects in question were either top-down or direct mixed adaptations if we can see it that adaptation all effects resulting from top-down and direct makes adaptations are (or would be) cases of weak cognitive penetration, then it's at least possible that the adaptational effects are cases of weak cognitive penetration.

If we concede that adaptation all effects resulting from top-down or direct mixed adaptations are (or would be) weak cases of cognitive penetration, and we grant that it's possible the adaptational effects like the illusion fading away were the result of top-down or direct mixed adaptations, then it follows that proponents of cognitive penetration weren't confused in claiming that adaptational effects are cases of cognitive penetration.

However, proponents of cognitive penetration need to establish that adaptational effects like the illusion fading away *are* cases of cognitive penetration, not just that it's possible. But that's not something my opponents get for free. It's also possible that the illusion faded away as a result of bottom-up or indirect mixed adaptations, and that's clearly not a case of cognitive penetration. And to do that they need to establish that the adaptational effects were the result of top-down or direct mixed adaptations.

2.6.2 The Empirical Question

So far I've argued it's in principle possible that adaptation all effects like the illusion fading away our cases of weak cognitive penetration. However, the mere possibility is not enough for proponents of cognitive penetration. Remember, they don't just claim that it's possible that adaptational effects like the illusion fading away our cases of (weak) cognitive penetration, they claim such effects actually are cases of cognitive penetration.

So the question I want to address now is whether we have any reason to only believe that adaptational effects like the illusion fading away are a result of either top-down or direct mixed adaptations. In a passage I quoted earlier, Churchland explicitly says that the illusion didn't fade away as a result of subjects' cognitive state alone.

> …the subjects are of course quite aware of what the lenses are doing. They may even know how they do it. But the illusion is not banished by the mere possession of this information. (Churchland, 1998:174)

So presumably it's safe to conclude that the illusion didn't fade away as a result of top-down adaptations. And that means the only way it could be a case of cognitive penetration is if it was the result of direct mixed adaptation(s).

I admit that it's possible. But as I said before, it's also possible that it was the result of bottom-up or indirect mixed adaptations. Whether the illusion fading away was the result of direct mixed

adaptations – and thus a case of the cognitive penetration – or the result of bottom-up or indirect mixed adaptations – and thus not a case of cognitive penetration – is an empirical question. Proponents of cognitive penetration give some reason to think that what actually happened. In other words, they need to establish that subjects' cognitive states played a direct role in bringing about the processing adaptations responsible for the illusion fading away.

I can't show that the illusion didn't fade away as the result of direct mixed adaptations, but I also don't think we have any reason to think it did. All the empirical evidence suggests that processing adaptations are the result of non-cognitive or bottom-up factors. And the illusion fading away seems to be no exception. Churchland himself suggests that the illusion faded away as a result of subjects continued engagement with their environment and familiar objects.

So I conclude that even if it's conceptually possible for the illusion fading away to be a case of cognitive penetration, we have no reason to believe it actually is. And the claim that the inverting lenses experiment is empirical evidence of cognitive penetration is completely unconvincing.

2.7 GENERALIZING

Up to this point I've been focused on the inverting lenses experiment and whether the illusion fading away is a case of (weak) cognitive penetration. But the inverting lenses experiment is only one of the empirical studies reportedly finding adaptational effects that proponents of cognitive penetration contend show there's (diachronic) cognitive penetration.

There is nothing unique about my argument against the claim that the inverting lenses experiment shows there's cognitive penetration. It applies *mutatis mutandis* to all other empirical studies reportedly finding adaptational effects that have been presented as empirical cases of cognitive penetration.

Proponents of cognitive penetration have also claimed an empirical study reportedly finding cross-cultural differences in subject susceptibility to visual illusions (e.g. the Muller-Lyer lines) shows there is (diachronic) cognitive penetration (McCauley and Henrich, 2006).

> The results for the Muller--Lyer stimuli show substantial differences among these social groups in their susceptibility to the illusion. American adults in Evanston, Illinois are the most susceptible. On average, these adults require that segment to be about a fifth longer than bbefore they perceive them as equal in length (PSE= 19%). At the other end of the "susceptibility spectrum," hunter -gathers from the Kalahari Desert are virtually immune to the "illusion." (They probably would not even recognize it as an illusion.) This population, on average, requires that segment to be only one percent longer than segment bbefore seeing them as equal (PSE = 1%). (McCauley and Henrich, 2006: 93)

This study apparently exposed subjects from around the world to various visual illusions (e.g. the Muller-Lyer lines), and reportedly found that subjects from Western cultures were more susceptible to the illusions than subjects from the non-Western cultures. So, for example, when it comes to the subjects from Western cultures, the Muller-Lyer lines must be very different links before subjects will visually experience them as the same. But apparently, subjects from non-Western cultures visually experienced the lines as the same length when the lines were actually very close to being the same length – that's the sense in which these individuals are less susceptible to the illusion.

McCauley and Hendrich claim these findings show that subjects' visual experiences (of the illusions) were cognitively penetrated. I haven't looked into this study myself, so I don't know whether the authors claim that there were cross-cultural differences in the way subjects visually experienced the illusions is an interpretation warranted by the experimental findings. However, let's assume it is. For cross-cultural differences in subjects' visual experiences of the illusions to show that subjects' visual experiences were cognitively penetrated (over a long period of time), it would

have to be the case that the differences in subjects' visual experiences were the result of direct influence from their cognitive states.

But McCauley and Henrich explicitly say that the differences in subjects' susceptibility to illusions were a consequence of variations in the way their visual systems processed visual (and perhaps other sensory) stimuli/information. And the claim is that there were differences in the way subjects' visual systems processed information across cultures because they had grown up in different environments. And those different environments had shaped their perceptual systems differently (NB the carpentered environments hypothesis).

> Explanations for the observed cultural variation in people's susceptibility to visual illusions center around the...notion that the human visual processing system will somehow adapt to the local visual environment by building up biases that tend to produce useful inferences in that environment. (ibid., p. 91)

Let's assume that that adaptation to their local visual environments is the reason there were cross-cultural differences in subjects' visual experiences. That would mean the cross-cultural differences are adaptational effects. So this experiment only shows there's cognitive penetration if those adaptational effects are cases of cognitive penetration.

A third example comes from a recent paper by Ariel Cecchi (2016) in which she argues that reported findings from a visual training experiment show there's, what she calls, *architectural cognitive penetration*.[23]

[23] Here is how Cecchi defines architectural cognitive penetration: "Architectural cognitive penetration is the process whereby the behaviour or the structure of the perceptual system is affected by the cognitive system. Cognitive influences can affect the architecture of the system either by guiding the function of the system – e.g., saccadic eye movements may be influenced by intentions (Wu 2013) – or by modulating its structure – e.g., cognitive influences may elicit neural reorganization in the visual cortex (Churchland 1988). Architectural cognitive penetration has an indirect impact on the content of perceptual experience. The subject's cognitive background first influences the

> The task consisted in an intensive monocular training requiring subjects to report, simultaneously, both the identity of an element appearing in the centre of the visual field and the orientation of an object located in the periphery...Participants in the experiment were divided into two groups. One group of subjects trained the left eye whereas the other trained the right...subjects underwent a total of 1,760 trials...performance in the detection of the peripheral element's orientation for the trained and untrained eyes was tested. The results indicate a significant improvement in peripheral target detection for the trained but not the untrained eye...the perceptual learning task in the target study produced a rapid improvement in visual detection. *This progress occurred as a result of cognitive influences on the visual system facilitating visual detection: electrical activity from higher brain areas affected early stages of the visual processing* [emphasis added]. The constant repetition of the detection task produced gradual neural modulations in the trained region of the primary visual cortex improving performance. (Cecchi, 2014: 74)

According to Cecchi, the experiment found that repeated training on a set of visual stimuli significantly improved subjects abilities to detect target stimuli. She also claims the improvement in subjects' abilities to detect stimuli "occurred as a result of cognitive influences on the visual system facilitating visual detection".

Let's assume the study actually found that subjects' abilities to detect the target stimuli improved after a period of training. Does that finding show, as Cecchi claims, that there's architectural cognitive penetration?

That claim seems problematic on numerous levels. First, the study reportedly found there were improvements subjects' abilities to detect stimuli. But to show there's cognitive penetration (architectural or otherwise), the study would need to show there was some kind of an effect on subjects visual experiences, i.e. that the training changed how the target stimuli looked to subjects. Now perhaps Cecchi has in mind that there couldn't have been an improvement in subjects' abilities

---

architecture of the visual system, which consequently may have an impact on the perceptual experience. (Cecchi, 2014:63)"

to detect stimuli without a corresponding change in their visual experience. That's not crazy. But it is a substantive claim in need of defense (which Cecchi does not provide).

But even setting that aside, for the changes brought about by the training to be cases of cognitive penetration, it would have to be the case that those changes were a result of (direct) cognitive influences. In the passage I quoted, Cecchi claims that "the progress [improvement] occurred as a result of cognitive influences on the visual system facilitating visual detection". It's not clear the results from the trainings support her claim. But even if they do, that shows, at best, that cognition influenced the visual system leading to that improvement in subjects' abilities to detect stimuli. But this is an adaptational effect.

> The neural structure of the brain is not rigid and static but plastic and dynamic. Visual cortical areas and their functions can be adapted to different perceptual conditions; such adaptation is possible thanks to neural reorganization in the visual system. (Cecchi, 2014: 65)

This is another example of an empirical study purportedly finding an adaptational effect. As before, I'm not questioning whether the effect on subjects' visual experiences was actually found, or that the effect was the result of changes in the structure of subjects visual systems. But effects on visual experiences resulting from changes in the structure of perceptual systems are *adaptational effects*. And, as I've argued, Cecchi would need to establish that adaptational effects are cases of cognitive penetration.

Cecchi contends that the adaptational effects found in the training experiment are cases of cognitive penetration because subjects' attention was involved. More specifically, the study reportedly found that there were only changes in subjects visual systems and abilities to detect the target stimuli *if they were paying attention during the training.* And that, she contends, shows subjects *attention* was required to get the relevant visual adaptations (*ibid.*, p.75).

Let's grant that the purported differences in the visual adaptation are real and that those differences were the result of influence from subjects' attention. If we grant those two things it follows that influence from subjects attention made a genuine contribution to bringing about the visual adaptations in question. But that, by itself, isn't enough to establish that the visual adaptations were cases of cognitive penetration. Just because influence from subjects' attention was needed to bring about the visual adaptations in question (i.e. they would not have come about if it weren't for the influence from attention), it doesn't follow that the influence from attention played the right kind of role (i.e. direct). To put it another way, the fact that influence from subjects attention contributed to bringing about the visual adaptations in question, is consistent with the visual adaptations being instances of direct mixed adaptations – and therefore cases of cognitive penetration – and instances of indirect mixed adaptations – and therefore not cases of cognitive penetration.

2.8 CONCLUSION

Certainly the primary, if not *only*, motivation for believing there is cognitive penetration is the supposedly compelling empirical evidence of it. But all the empirical evidence of cognitive penetration is either from studies reportedly finding top-down effects or from studies reportedly finding adaptational effects.

Proponents of cognitive penetration don't need to convince us that empirical studies finding top-down effects would show there is cognitive penetration. That's not even up for debate. Top down effects are (or would be) clear cases of cognitive penetration. So there couldn't be top-down effects unless there was also cognitive penetration. What they need to convince us of is that top-down effects were actually found in the empirical studies they discuss.

However, when it comes to the empirical studies reportedly finding adaptational effects things are the exact opposite. When it comes to these studies, proponents of cognitive penetration don't need to convince us that adaptation effects were actually found. Since no one doubts there are adaptational effects, we can just assume such effects were actually found in the studies in question. The challenge is to convince us that the adaptational effects found in the relevant studies are cases of cognitive penetration.

Surprisingly, proponents of cognitive penetration offer no such argument. Given that their claim that adaptational effects are evidence of cognitive penetration turns on this very issue, the fact that no argument *whatsoever* is offered to substantiate this claim is baffling. However, this would be explained if proponents of cognitive penetration think that all adaptational effects, even ones resulting from processing adaptations caused by "wider perceptual experience" or, what I've called, bottom-up adaptational effects count as cases of cognitive penetration. Even if we are willing to count effects from top-down and mixed adaptations as cases of cognitive penetration, I take it claiming that effects of bottom-up adaptations are also cases of cognitive penetration is a non-starter.

2.9 REFERENCES

Bitter, David (2014). "Is Low-Level Visual Experience Cognitively Penetrable?" In: The Baltic International Yearbook of Cognition, Logic and Communication 9.1, pp. 1–26.

Bruner, Jerome (1947). "Value and Need as Organizing Factors in Perception." In: Journal of Abnormal and Social Psychology 42, pp. 33–44.

Cecchi, Ariel S. (2014). "Cognitive Penetration, Perceptual Learning and Neural Plasticity." In: Dialectica 68.1, pp. 63–95.

Churchland, Paul M. (1988). "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry

Fodor." In: Philosophy of Science 55.June, pp. 167–87.

Degenaar, Marjolein and Gert-Jan Lokhorst (2014). "Molyneux's Problem." In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Spring 2014.

Firestone, Chaz and Brian J. Scholl (2016). "Cognition does not affect perception: Evaluating the evidence for "top-down" effects." In: Behavioral and Brain Sciences 39.

Fodor, Jerry A. (1984). "Observation Reconsidered." In: Philosophy of Science 51.March, pp. 23–43.

Gross, Steven et al. (2014). "Problems for the Purported Cognitive Penetration of Perceptual Color Experience and Macpherson's Proposed Mechanism." In: Though and Perception. Ed. by Edouard Machery and Jesse Prinz, pp. 01–30.

Levin, Daniel T. and Mahzarin R. Banaji (2006). "Distortions in the perceived lightness of faces: The role of race categories." In: Journal of Experimental Psychology: General 135.4, pp. 501–512.

Macpherson, Fiona (2012). "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism." In: Philosophy and Phenomenological Research 84.1, pp. 24–62.

McCauley, Robert N. and J. Henrich (2006). "Susceptibility to the Muller-Lyer Illusion, Theory-Neutral Observation, and the Diachronic Penetrability of the Visual Input System." In: Philosophical Psychology 19.1, pp. 79–101.

O'Callaghan, Casey (2010). "Experiencing Speech." In: Philosophical Issues 20, pp. 305–332.

Pylyshyn, Zenon. W. (1999). "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception." In: Behavioral and Brain Science 22.3, pp. 341–365.

Scholl, Brian J. and Alan M. Leslie (1999). "Modularity, Development and "Theory of Mind"." In: Mind and Language 14.1, pp. 131–153.

Stokes, Dustin (2015). "Towards a Consequentialist Understanding of Cognitive Penetration." In: Cognitive Penetrability (Oxford University Press). Ed. by A. Raftopoulos and J. Zeimbekis.

Wu, Wayne (2013). "Visual Spatial Constancy and Modularity: Does Intention Penetrate Vision?" In: Philosophical Studies 165.2, pp. 647–669. Zeimbekis, John (2013). "Color and Cognitive Penetrability." In: Philosophical Studies 165.1, pp. 167–175.

Chapter 3: Against the Perceptual Representation of High-Level Properties

3.1 INTRODUCTION

What are the admissible contents of experience? This question has garnered a fair bit of attention in recent years.[24] If we suppose experiences are representational, then the *contents of experience* are (or determine) what the experience represents. This question is specifically about what *properties* feature in these contents. That is, what properties are representable by experience. If any experience can represent a property, the property is *admissible* in the contents of experience; otherwise not.

It will be helpful to have a concrete example. Suppose you're looking at a bowl of fruit. Under normal conditions, you will have an experience that is quite complex—perhaps involving the smell of oranges, pangs of hunger, memories of lunch, etc. Ignore everything except for the *visual* aspect of your experience—just what would be missing if you were completely blind. I'll call this your *visual experience*. From here on, visual experience will be my sole concern.

What properties does your visual experience, on this occasion, represent? Traditionally, the answer would have been exclusively basic or *low-level* properties, such as colors and shapes. But some philosophers, most notably Susanna Siegel, have been arguing the answer also includes some sophisticated or *high-level* properties, such as being a natural kind (e.g. being an apple).[25] Just to be clear, that we visually represent low-level properties is not in dispute. The only issue up for debate is

[24] See (Hawley and Macpherson 2011) for a recent collection of papers on the topic. See also (Brogaard 2012), (Campbell 2013), (Crutchfield 2012), (Deroy, 2013), (Logue, 2013), (Masrour, 2011), (Matey, 2012), (Siegel, 2007), (Silins, 2013), (Sorensen, 2013). For a classic argument that we see causation, see (Michotte, 1963).

[25] I don't have a nice way carving up properties into low and high. Fortunately, for my purposes that won't much matter. One possibility I'm attracted to is using cognitive penetration to draw the distinction—those properties representable in experience without using cognitive penetration are low-level properties, otherwise they're high-level properties. On this way of carving things up, it may be that some properties heretofore thought of as paradigmatically high-level (e.g. being a face) turn out to be low-level.

whether we visually represent high-level properties too. Proponents of the *high-level view* say yes; and proponents of the *low-level view* say no.

My objective in this paper is to make a case for rejecting the high-level view. There are two components to my strategy. First, I'll argue that the most prominent argument for the high-level view, the so-called *phenomenal contrast strategy*, doesn't offer us any reason to endorse the view. Second, I'll argue that we should reject the view because there isn't, and can't be, a plausible explanation for how we come to visually represent high-level properties.

Before moving on, I want to address a question/comment I sometimes hear: isn't it obvious that the high-level view is correct? We definitely *see* more than, e.g., colors and shapes—we see kinds like apples and bananas. Of course, in one sense, that's undeniable—the objects of our experience are (typically) things like apples and bananas, not colors and shapes.[26] *But no one* (in this debate) *denies that*. Easy considerations like this are beside the point; and definitely do not settle the matter in favor of the high-level view. You can see an apple without visually representing the property of being a apple. Just as you can have beliefs about an apple without cognitively representing the property of being an apple.

3.2 THE PHENOMENAL CONTRAST ARGUMENT

The most prominent argument for the high-level view, the *phenomenal contrast strategy* (PCS), is, at bottom, an inference to the best explanation, where the thing to be explained is a phenomenological difference brought about by gaining a recognitional ability.[27] It's argued that the visual representation of high-level properties (or *high-level representation*) is the best explanation of this difference. If that's

[26] However, some philosophers have suggested that, in hallucination, uninstantiated properties are the objects of our experience.

[27] The argument was explicitly put forward first by Susanna Siegel; and subsequently, Tim Bayne advocated a variation on it.

right, we have good reason to think high-level properties are admissible in the contents of visual experience; and thus for endorsing the high-level view.

I'll argue there's no apparent reason to think the explanation is in terms of high-level representation rather than *low-level representation* (i.e. the visual representation of low-level properties); and thus no apparent reason to endorse the high-level view. Let's start by looking at the motivation for the claim about phenomenal difference.

> Suppose you have never seen a pine tree before and are hired to cut down all the pine trees in a grove containing trees of many different sorts. Someone points out to you which trees are pine trees. Some weeks pass, and your disposition to distinguish the pine trees from the others improves. Eventually, you can spot the pine trees immediately: they become visually salient to you...Gaining this recognitional disposition is reflected in a phenomenological difference between the experiences had before and those had after the recognitional disposition was fully developed (Siegel, 2010:100).

If Siegel is to be believed, learning to recognize pine trees (i.e. gaining the recognitional disposition) changes your experience—what it's like for you to look at pines. In fact, she argues that it changes your visual experience, that pine trees *look* different when you're disposed to recognize them. *This* is the difference it's argued can best be explained in terms of a high-level representation. Let's call the experience had before you gain the recognitional ability e_1 and the experience after e_2; and specify the experiences of interest as follows:

> A pair of experiences e_1 and e_2 such that in the situation leading to e_1, S lacks the ability to recognize the property F, whereas in the situation leading to e_2, S has the ability. Otherwise, the situations are as similar as possible.

I refer to these as *recognitional cases*. Of particular interest are recognitional cases where F is a high-level property. These are the cases where it's plausible the phenomenal difference is best explained in terms of high-level representation. Here is my rendering of PCS.

(1) There's a phenomenal difference between e_1 and e_2
(2) If there's a phenomenal difference between e_1 and e_2, there's a difference in the

properties visually represented.

(3) If there's a difference in properties visually represented, it's a difference either in low-level or high-level properties.[28]

(4) A difference in high-level properties offers the best explanation of the phenomenal difference.

(5) So it's a difference in high-level properties (and thus, the high-level view is true)

I'm prepared to grant premises (1), (2), and (3) for the sake of argument. In other words, I'm willing to concede there's a phenomenal difference brought about by recognition that needs to be explained in terms of a difference in properties visually represented. But I'll argue there's no reason to think a difference in high-level representation (or a *high-level difference*) is better suited to explain it than a difference in low-level representation (or a *low-level difference*). Let me quickly run through why the first three premises are needed, then I will discuss premise (4) at length. To my mind the success of the entire argument comes down to what can be said on behalf of premise (4). And what little has been said cannot withstand the pressure.

The first premise is just the intuition, hopefully elicited in the pine tree case, that there's a phenomenal difference brought about by gaining a recognitional ability. Accepting (1) is the minimum necessary to get PCS off the ground. Premise (2) is actually the amalgamation of three claims, which can be unpacked as follows:

(2a) If there's a phenomenal difference between e_1 and e_2, it's a difference in visual phenomenology.

(2b) If there's a difference in visual phenomenology, there's also a difference in visual representation.

(2c) If there's a difference in visual representation, it's a difference in the properties visually represented.

According to Siegel, merely reflecting on the pine tree case or ones similar to it should be enough to

[28]Of course, it might be that *both* low-level and high-level properties make a difference. I'm including this with the latter because the debate is over whether we represent low-level properties exclusively or along with high-level properties.

convince you of (1). But she acknowledges it may not be enough to convince you that it's a difference in *visual* phenomenology. To get you there, she argues other sources of phenomenology (e.g. feelings familiarity and imagery associated with memory) don't offer adequate explanations the difference. But this only helps her cause if there's a connection between visual phenomenology and visual representation (i.e. the contents of visual experience). This is indeed the case, if *intentionalism* is true—intentionalism is the thesis that the phenomenology of experience supervenes on its content. If so, any difference in visual phenomenology must correspond to, or as I've been saying—*explained in terms of* —a difference in visual representation.[29] This takes us through (2a) and (2b). The purpose of (2c) is to establish that the difference in representation is with respect to properties, as opposed to some other type of constituent (e.g. objects). If we assume properties visually represented are either low or high-level (premise (3)), there are only two options to explain (1): a low-level difference or a high-level difference. To establish that high-level differences offer the best explanation, it either has to be shown that low-level differences can't do it or that they can't do it as well. I'll consider each of these options in turn.

Why might proponents of PCS think low-level differences can't explain (1)? Siegel is not as forthcoming about this as one might hope. Here's what she says.

> In this sort of [recognitional] case, it seems implausible to suppose that there must be a change in which color and shape properties are represented before and after one learns...[to recognize a property like being a pine tree]. (112)

As I interpret her, Siegel is saying it's implausible *there are* low-level differences in every recognitional case. And if there are no low-level differences in some recognitional cases, then at least sometimes

[29] Siegel doesn't appeal to intentionalism herself. Presumably because it's a controversial thesis and the fewer controversial claims you have to rely on the better. Instead, she argues for the specific claim that *this* phenomenal difference is a representational difference. See Siegel (2010:110 -111). But this doesn't matter for my purposes, because I grant this premise.

high-level differences will offer the *only*, and thus the best, possible explanation. But the question is why think there aren't any low-level differences? Siegel isn't explicit about this either, but an answer is suggested by the very nature of recognitional cases.

The last line in the specification of recognitional cases reads: *Otherwise, the situations are as similar as possible*. This is included to ensure that recognitional cases are, what Ole Koksvik (2015) calls, *minimal pairs*—cases where the experiences (e_1 and e_2) arise in situations that only vary with respect to recognitional ability and everything else (from the lighting conditions to the subject's viewing perspective and sensory organs) stays the same. This ensures that (1)—which is immediately explained by low and/or high-level difference(s)—is ultimately explained by the recognitional difference.[30] This is important to the success of PCS because, for reasons that are not entirely clear, the fact that representational differences have to be explained in terms of the recognitional difference makes it plausible that it must be high-level differences.

I'm willing to grant that the only difference in the situations giving rise to e_1 and e_2 is with respect to recognitional ability. How is this relevant to the issue at hand—why it's implausible there are low-level differences? Here's the thought: it's hard to see how recognitional differences could explain low-level differences. If that's right, it's implausible there are such differences because there's nothing around to explain them.[31][32]

---

[30] It may be helpful to think of this on the model of randomized controlled studies in medicine. To figure out the impact of a given treatment, the treatment will be given to one group and withheld from another. But it's important that these groups otherwise be as similar as possible. That way there's nothing besides the treatment to explain any differences between the groups. Similarly, if the recognitional ability is the only difference in the situations giving rise to e1 and e2 , there's nothing besides the recognitional ability to explain the representational difference.

[31] Nothing around to explain low-level differences would be good evidence there aren't any, because, if there were, there would be something to explain them.

[32] Siegel does consider one possible low-level difference in, what she calls, a *pine-tree-shape gestalt*.

> Suppose that when you learn to recognize pine trees by sight, your experience comes to represent a complex of shapes—leaf shape, trunk shape, branch shape, and overall pine-tree shape. This complex is an overall pine-tree gestalt. The pine-tree-shape gestalt is general enough that it can be

Ultimately, however, I don't think this reasoning can be right. Admittedly, it's not immediately apparent how just learning to recognize pine trees could explain a low-level difference. But *that cuts both ways*—it's not immediately apparent how it could explain a high-level difference either. In other words, the fact that recognitional ability is the only variable is no more (or less) reason to suspect there aren't low-level differences, than to suspect there aren't high-level differences. Since, given our other assumptions, it must explain one or the other (or both), that the explanation is not immediately apparent isn't grounds for suspicion.[33]

It's worth noting that appealing to cognitive penetration won't help. You might think cognitive penetration makes it plausible that recognitional differences can explain high but not low-level differences. But this is a mistake. Cognitive penetration as a mechanism to explain differences in visual representation doesn't favor one kind of difference over the other. If there is cognitive penetration, it could bring about low-level differences as easily as high. In fact, the experimental data purportedly supporting cognitive penetration (that I'll discuss in section 3), involves low-level differences. Moreover, it's actually *far more* plausible that cognitive penetration can be used to explain a low-level difference. As I'll argue in section 4, an explanation involving cognitive penetration and high-level differences is extremely implausible.

However, the fact that (1) must ultimately be accounted for in terms of recognitional differences might help shore up PCS by making it plausible that high-level differences make for a better explanation (than low-level differences). If that's right, it doesn't matter if there are (or could

---

> shared by different-looking pine trees. But it is specific enough to capture the look shared by exemplary pine trees. (2010:111)

She thinks this won't do because a shape gestalt abstract enough to be invariant over the shapes of actual pine trees will be very abstract and "the more abstract a shape-gestalt is, the less reason there is to think that experience fails to represent it prior to one's gaining a recognitional disposition (Siegel, 2010:112)." I don't find a difference in shape gestalt particularly plausible either.

[33] Tim Bayne explicitly allows that there are (or might be) low-level difference(s) in every recognitional case.

be) explanations in terms of low-level differences—explanations in terms of high-level differences would still be best. But is it right? Is there any reason to think high-level differences make for a better explanation?

Again, Siegel doesn't address this directly, so we will have to do some reconstructive work. I mentioned at the beginning of this section that Siegel argues that the phenomenal difference is a difference in the way things look. Charles Siewert makes the same point in the following passage from *The Significance of Consciousness*.

> ...your visual experience of sunflowers differs phenomenally depending on whether or not you can recognize sunflowers as sunflowers...generally I think we can say that the way things look to us alters when we come to recognize what types of things they are....There is a way it seems to us to see sunflowers not just as some more shaped and colored things, but as what has a distinctively sunflowery look (Siewert, 1998: 256)

I'm not entirely sure I know what a "distinctively sunflowery" look is, or for that matter what a "look" is. But setting those issues to one side, you might think the sunflowery look sunflowers take on can better be explained by a high-level difference than a low-level one. The thought being that the visual representation of being a sunflower can better explain the distinctive look than the visual representation of any other property or properties. Moreover, the reason sunflowers take on this look once you've learned to recognize them is that you come to visually represent the property of being a sunflower for the first time.

Once again, I'm willing to just grant that there is a distinctive sunflowery look that sunflowers take on when you've learned to recognize them. However, I deny that this is better explained by the visual representation of the property of being a sunflower, than any other property or properties. Or at least I'm not willing to grant it without argument. Here's at least one reason to think the sunflowery look is not best explained by the visual representation of being a sunflower. The visual representation of being a sunflower should, at least first pass, be cashed out as looking to

be a sunflower. But, as Alex Byrne (2009) has pointed out, looking *to be* F is importantly different than *looking* F. He gives the example of a naked mole rat looking old. It's looking old is decidedly *not* it's looking *to be* old—in fact, it may even look to be young! The rat looks old, if anything, because of it's wrinkly and hairless appearance—an appearance mole rats have regardless of age. Far from being the best, an explanation of the mole rat looking old in terms of the visual representation of the property being old is not even a reasonable explanation. Why should we think the sunflowery look is any different?

You might be tempted to think the sunflowery look of sunflowers is different than the old look of mole rats because of how the former came about, namely by learning to recognize sunflowers. Plausibly, this consists in learning to recognize things *as being* sunflowers.; and thus the look things take on is looking to be sunflowers.

It may be the case that the ability to recognize sunflowers is an ability to recognize things as being sunflowers.[34] But it's not plausible that the look they take on as a result is looking to be sunflowers. After you've learned to recognize sunflowers, both real ones and obvious fake ones will presumably share the elusive sunflowery look. But obviously fake sunflowers don't look *to be* sunflowers. So looking sunflowery is not looking to be a sunflower.

I've tried to find reason for thinking that an explanation of (1) in terms of high-level differences is the best and come up empty handed. It's consistent with what I've said so far that there is a reason that I've overlooked. But unless and until we get one, PCS fails as an argument for endorsing the high-level view. In what remains, I will make the case that we have good reason to reject the high-level view—the existence of high-level differences in recognitional cases can't be plausibly explained.

[34] I'm not sure about this. Wouldn't we count someone as being able to recognize sunflowers even if they live in fake sunflower country? One reason to think yes is that, upon seeing actual sunflowers, she would identify them as such. But the issue is complicated, is a precondition on recognizing something as a sunflower that you possess the sunflower concept?

3.3 EXPLAINING HIGH-LEVEL REPRESENTATIONS

Like the differences in phenomenology they explain, differences in visual representation must also be explained; and I've already granted that this must be done in terms of gaining a recognitional ability. So far I've been relying on an implicit understanding of recognitional ability, but at this point it will be helpful to have something slightly more concrete. Here's Siegel.

> A perceiver who can recognize trees by sight seems to have some sort of memory representation, and some sort of perceptual input, such that the input "matches" the memory representation, and the cognitive system of the perceiver registers that this is so. (2010:110)

Assuming this is basically right, gaining a recognitional ability consists in difference(s) in one or more of the following: stored representations, the matching process, and/or sensory input. The first two are, roughly, *cognitive components*—involving higher level mental processing—and the third is a *sensory component*—involving basic sensory stimulation. So gaining a recognitional ability might involve changes to either of these components. However, as I understand it, in recognitional cases, the assumption is that the sensory component doesn't change (I'll consider relaxing this assumption later); so in the relevant cases, there must be change(s) to the cognitive component (or *cognitive differences*) and they must ultimately explain the difference in visual representation (or a *visual difference*).

Broadly speaking, there are two ways a change in the cognitive component might do this: either directly or indirectly. If a change to the cognitive component explains a change to the sensory component[35] and *that* explains the difference in visual representation, the cognitive difference *indirectly* explains the visual difference. The following is a simple example. Suppose you believe your car is parked off to the left. If you want to see it, you will probably look left *because that's where you*

[35] As I'm understanding it, a change in *attention* is a change in the sensory component.

believe your car is parked. If you do, your visual representation will change and be different than it was before. There's a clear sense in which your belief explains this difference—if you hadn't believed as you did, it's substantially less likely you would have looked left and thus that your visual representation would have changed in just the way it did. However, if your belief hadn't explained the change to the sensory component, it wouldn't have explained the difference in visual representation.

A more interesting possibility is a change in the cognitive component explaining a difference in visual representation *directly*—without explaining a change in the sensory component. If the cognitive component can do this, there's *cognitive penetration*. Apparent examples are discussed in empirical psychology.[36] In one study, subjects learned a color-letter association and then, upon viewing the letters, reported they looked tinted with color. In another, subjects reflected on negative words and actions and then reported that the room looked darker than it did before.[37] In both studies, it appears that changes in cognition *alone* (without changing the sensory input), explain the visual difference. The letters looked one color, then some associations were learned, and they looked another color.

High-level differences have to be explained in one of these two ways. But for there to be an indirect explanation, there would need to be a difference in the sensory component and that has already been ruled out. So the explanation must be direct, in terms of cognitive penetration. Siegel acknowledges as much in the following passage.

---

[36] A wave of discussion from the 1950's is associated with Jerome Bruener and the New Look school. But there is currently another wave of discussion happening. See (Firestone and Scholl, in press) for references to studies I'll mention and several others.

[37] Firestone is skeptical that the findings in these studies are genuinely visual. I share his skepticism. MacPherson (2012) discuss an experimental finding that a subject's beliefs about the characteristic color of a shape influences the color it looks to be. For example, your belief that heart shapes are characteristically red makes heart shaped cutouts look redder to you.

> One kind of argument for the [High-level] View relies on the idea that non- perceptual mental states can influence the contents of visual experience. Accord- ing to this kind of argument, a [high-level] property such as the property of being a pine tree can come to be represented in visual experience partly in virtue of your ability to recognize pine trees when you see them. If so, then your visual experience would be *cognitively penetrated* [emphasis added] by the mental states associated with your recognitional ability. (p.10-11)

The fact that high-level differences are explained this way is significant because cognitive penetration is extremely contentious. Traditionally, it has been thought that visual experience is cognitively *im*penetrable or informationally encapsulated—the result of a (mostly) bottom-up process, taking sensory input and processing it through a hard-wired visual system.[38] Cognitive penetration would require a fundamentally different mental architecture—where visual experience results from a top-down process that takes cognitive states as inputs. If this were the case, cognitive states would be direct inputs to visual processing, and thus changes in cognition could directly explain differences in visual representation.

I won't try to convince you that there's no cognitive penetration (though I myself am dubious), but I will mention two very different reasons cognitive penetration has been viewed with suspicion. The first is the evident imperviousness of experience in cases of illusion. Despite our knowledge to the contrary, the lines in the Müller-Lyer illusion continue to look different lengths. If visual experience is cognitively penetrable, this kind of recalcitrance is puzzling. The second reason has to do with the possible implications for perceptual justification and knowledge. If there's cognitive penetration, it may be that *believing really* is *seeing* or that, to use a phrase from philosophy of science, "perception is theory laden". This raises serious questions about whether there is, as Quine thought, "a tribunal of the senses" to which we must subject our beliefs for *independent* confirmation (or disconfirmation).

[38] See (Marr, 1982) for a classic discussion.

Given the contentiousness of cognitive penetration, you might wonder whether high-level differences can be explained without it. If we maintain that visual differences have to be explained in terms of recognitional differences (and I see no reason to give this up), for there to be an alternative explanation we have to relax the assumption that there's no difference in the sensory component.[39] Then it's at least possible there's an explanation just in terms of the sensory component or indirectly in terms of the cognitive component. However, I'm not optimistic an explanation of either sort can be pulled off.

If Siegel is right, you don't fail to visually represent being a pine tree in e_1 because you lack appropriate stimulation (like Mary in her black and white room). You fail to visually represent it because you are, in some sense, *blind to it*. When you're looking at the grove of trees before you've learned to recognize pines, instances of the property are *right there in front of you*, staring you in the face; and yet, you fail to see them. But somehow, by gaining the recognitional ability your are able to see. To put it poetically, your blindness was a consequence of your ignorance and, thus, your knowledge enables you to see. This suggests that you must have a cognitive component to be able to visually represent high-level properties.[40] If that's right, an explanation just in terms of the sensory component won't do. No amount of difference to your sensory input could explain you coming to visually represent being a pine tree in e_2.

Since no change in just sensory input can explain it, plausibly, no change in cognition can do it indirectly. For that to be possible, the visual difference would have to be *independent* of the cognitive component, and the visual representation of being a pine tree is not. Consider again the car example I discussed earlier, but add the further detail that when you look left a blue car comes

[39] Doing so might undermine the motivation to think there are high-level differences that need to be explained. But that is beside the point at the moment.

[40] This suggests that at least certain versions of the high-level view (e.g. those which include properties like being a pine tree) commit you to being a conceptualist about experiential content.

into view; and suppose there's literally nothing else blue anywhere around. Then your belief explains your visual representation of blue. But there's an important sense in which, nevertheless, your visual representation of blue is *independent* of your belief. Your belief brought about conditions that allowed you to exercise an ability you already had. In other words, you *could* have visually represented blue even if hadn't believed (e.g. an involuntary neck spasm caused the car to come into view). However, your visual representation of being a pine tree is not independent of the cognitive state(s) you gain when you learn to recognize pines. You *couldn't* have visually represented being a pine tree without those (or perhaps some other) cognitive states.[41] In other words, your visual representation *depends* on the cognitive states, without them you have pine tree blindness. This kind of dependence on cognition rules out an indirect explanation.

If high-level differences can't be explained in terms of just changes in the sensory component or indirectly by changes in the cognitive component, we are left with explanations in terms of cognitive penetration. This means endorsing the high-level view comes along with a commitment to cognitive penetration. Moreover, it's unclear that the kind of cognitive penetration, which is apparently supported by empirical findings (e.g. in the letter-color association study), is the same kind of cognitive penetration needed to explain high-level differences.

In the studies I mentioned earlier, cognition directly explains the visual representation of properties *you already have the ability to visually represent*. So, for example, your negative thoughts directly explain your visual representation of a darker color; but you already had the ability to do that prior to having those thoughts. However, in the pine tree case, cognition directly explains the visual representation of a property (being a pine tree) that you didn't have the ability to represent prior to having those states. So in the latter but not the former, cognition explains both your very ability to

[41] I assume Siegel doesn't think the cognitive component of recognition is needed *in particular* (e.g. perhaps the cognitive states associated with the ability to verbally describe pine trees would do just as well).

visually represent the property and your visual representation of it on this occasion. So endorsing the high-level view, not only requires commitment to *ordinary cognitive penetration* but also to something, we might call, *enabling cognitive penetration*; and presumably, this is even more contentious than the ordinary variety. This is a pretty steep price to pay for endorsing a view there's, at best, very little reason to endorse.[42] A commitment to cognitive penetration (enabling or ordinary) will be enough to put some people off. But I take it Siegel, and others like her, won't be moved. After all, she was committed to cognitive penetration before I argued she had to be. In the next section, I'll argue her explanation is extremely implausible and this is not just because it relies on cognitive penetration.

3.4 RECOGNITION TO VISUAL REPRESENTATION

As I understand it, the high-level difference in the pine tree case is explained as follows. In e_1 you don't visually represent the property of being a pine tree. You have various sensory inputs $S_1...S_n$ and cognitive states $C_1...C_n$. But you lack the cognitive states constitutive of the ability to recognize pine trees. Then, you gain the recognitional ability—e.g. by looking at field guides, visiting arboretums, and speaking with experts—acquiring the aforementioned cognitive states, call them *Cp*. When you undergo e_2, everything is the same except now you have *Cp*. But surprisingly, this time around, you do visually represent the property of being a pine tree. It's as if the addition of *Cp* enabled you to visually represent the property. I'll call this the RR explanation (for *recognition to representation*).

According to RR, we can come to see new properties (e.g. being a pine tree)—properties we were previously blind to—just by learning to recognize them. Imagine a colorblind person started seeing red or you starting seeing ultraviolet blue. The claim is that something similar happens to us

[42] Even if you think there is a phenomenal difference in recognitional cases that, at present, we don't have a satisfactory explanation for without high-level representation, you might want to hold out for something a little less costly.

when we learn to recognize things. If this is true, it's surprising (to put it mildly). If you can't visually represent a property, can learning to recognize it really help you overcome your blindness? I'll argue the answer is no for low-level properties, which can be used as the basis for an inductive generalization to *all* properties, low and high.

Suppose Ruby red color blind and (predictably) also lacks the ability to recognize red when she sees it. Imagine she's given a special pair of glasses that make red things glow (i.e. look to her as if they are highly illuminated) and informed of their special power. When she first puts on the glasses, she must deliberately think (a relative slow and cognitively demanding process) about the power of the glasses and the fact that the glowing indicates the presence of red. However, after some weeks or months of wearing the glasses continuously, this becomes less and less deliberate and more and more automatic. Eventually she gets so used to the glass that she bypasses the transition from glow to red altogether and is immediately aware of red. By the end of the process (if not substantially before), Ruby has gained the ability to recognize red when she it. However, she has not come to visually represent any new properties, including red, or so it seems to me. She doesn't visually represent red now anymore than she did before she got the glasses. Moreover, it seems pretty clear that *any* way Ruby could learn to recognize red won't enable her to visually represent it.

If this is right (and I think it is), RR is false for low-level properties. In other words, RR can't explain low-level differences. This is a strong basis the claim that the falsity of RR generalizes to *all* properties, low and high. If RR can't explain high-level differences and they can't be explained in any other way (as I argued in the previous section), then we have good reason to think there are no high-level differences; and thus that the high-level view is false.[43]

[43] Consider an analogy. Suppose I'm trying to convince you that a bear has been in my house. The reason I give is that it's the best explanation of the claw marks on my sofa. You might wonder, how did a bear get in the house? Suppose there are no signs of forced entry, so I say it must have teleported. That's not a plausible

I can think of three ways my opponent might respond: deny that Ruby has learned to recognize red, insist that she has come to visually represent it, or accept that RR is false for low-level properties but block the generalization by arguing that the recognition and/or visual representation of high and low-level properties is different in a way that's relevant. I'll consider the first two responses in this section and the last response in the next section.

The option of denying that Ruby has learned to recognize red is not particularly promising. It might be pointed out that Ruby's way of recognizing red is substantially different than ours. We recognize red *directly* using our visual representations of the property. Whereas Ruby recognizes red, if at all, *indirectly* using her visual representation of the glow plus some background cognitive states that link the glow to the property red.

There is indeed this difference, but it's unclear why it's at all relevant to whether Ruby's recognitional ability is genuine. One reason to think the ability *is* genuine is that Ruby (with her glasses on) can do the sorts of things we associate with having a recognitional ability (e.g. identify, discriminate, group, and categorize). Pointing out that Ruby's recognition is indirect and ours is direct isn't reason to think her ability is less genuine. And my opponents would be hoist on their own petard if they argued that genuine recognition of a property requires visual representation of it. If that were the case, we wouldn't be able to recognize high-level properties before we came to visually represent them and so RR would be false by their own lights. They might try to maintain that visual representation of a property is a requirement for genuinely recognizing low but not high-level properties. The is species of the third option, which I'll consider in section 5. At this point, suffice it to say I'm not optimistic.

A different option is try arguing that Ruby does come to visually represent red after wearing

---

explanation for how the bear got in. If there's no other explanation in the offing, you should conclude that there was no bear in the house, and thus it's not the best explanation of the claw marks.

the glasses for a time. Now, perhaps there's a sense in which Ruby visually represent red—she sees *that* the glow is red. But, if anything, this is a kind of *indirect* visual representation and that is an entirely different beast then what we've been talking about until now. No one denies that we can see *that* something is a pine tree. The question is whether we can see something as a pine tree in the way we can see it as red. For the high-level view to be interesting, the answer has to be yes. In other words, we must visually represent high-level properties in, roughly, the *same way* we visually represent low-level properties. And the way we visually represent something as red is (typically) *directly*, not *indirectly*. What makes the high-level view a genuine *alternative* to the low-level view, is the claim that high-level properties feature in the content of visual experience the same way low-level properties do. The suggestion now is that high-level properties figure in the contents of a different, but related state, we might call *seeing-that*. This threatens to trivialize the high-level view by reducing it to something that everyone already accepts.

A way to maintain the line of indirect visual representation without trivializing the high-level view is to argue that *all* seeing is seeing that. That is, low-level representation is indirect too. What is seeing that or indirect visual representation? Dretske calls it a kind of epistemic seeing, it is a kind of seeing that requires the possession of certain concepts or cognitive states.

The first thing to note is that this way of going seems to commit you to conceptualism about experiential content. But commitment to conceptualism may already be required because of the commitment to cognitive penetration.

Another reason to deny that all seeing is seeing that is that, if the possession of certain concepts is a precondition on having any sort of visual experience *at all* and given that our experiences are the source of many (most?) of our concepts, then as Dretske (1969) says total ignorance would guarantee total blindness" and, furthermore, there would be no way to remedy the

situation. This is one reason Dretske insists there is a more basic kind of seeing than seeing that, which he calls *seeing*$_N$.

3.5 POSSIBLE OBJECTION AND REPLY

The final option is to accept that RR is false for low-level properties, but insist that it's true for high-level properties. This would require blocking the generalization from the failure for low-level properties to a failure for high-level properties. This could be done either by establishing a relevant difference between the recognition of low and high and/or visual representation of them. *Relevant* differences are differences between low and high that would block the generalization. So to successfully use this strategy proponents of the high-level view have to do two things: argue a difference is *relevant* and that *there really is* that difference. I'll argue that it's with regard to the second task that the high-level view runs into problems. Any proposed difference must strike a delicate balance—too much and it threatens trivializing the view, too little and it won't be enough to block the generalization.

If there is high-level representation, one thing that seems undeniable is that high-level representation somehow *depends* on low-level representation. More specifically, the *phenomenology* of high-level representation depends on the *phenomenology* of low-level representation. Why think this? Because you can have low-level representation without high-level representation[44], but not the other way around—if you are low-level blind you will definitely be high-level blind. One thought is that the fact that high-level representations are *phenomenally dependent* whereas low-level representations are not is a relevant difference.

Before considering whether there is this difference between low and high-level representation, let's first discuss why this difference is relevant. That is, how does the fact that high-

[44] According to Bayne, this is what goes on in cases of associative agnosia.

level representation is phenomenally dependent make it plausible that RR is true for high-level properties but not low? The thought might be something like this: your ability to recognize a property F, can only enable you to visually represent F, if the visual representation of F (which you don't yet have) *phenomenally depends* on the visual representations you use in learning to recognize F. The idea is behind this suggestion is that the way you learn to recognize a property is relevant to whether your ability to recognize a property can enable you to visually represent it. And the phenomenal dependence of the ultimate visual representation on the visual representation involved in recognition is required. The motivation behind this suggestion might be that if there's this kind of phenomenal dependence, your visual representations of some properties can help you "get on" to the visual representation of others.

The problem with Ruby, then, is that the visual representation of red is in no way phenomenally dependent on her visual representation of the glow, which she uses to recognize red. And without the required phenomenal dependence, her recognition of red won't enable her to visually represent red. Ruby's case is the rule, not the exception when it comes to low-level properties because the visual representations of low-level properties are *phenomenally basic*—they don't depend on any further visual representations. If that's right, then if you can't already visually represent a low-level property, no way you could learn to recognize it will enable you to visually represent it. If low-level properties are phenomenally basic, trivially recognizing a low-level property won't involve visual representations it depends on (except if the recognition involve the property itself).

When it comes to high-level properties, however, the situation is different precisely because high-level representation phenomenally depends on low-level representation. Suppose, for example, that the visual representation of being a pine tree depends on the visual representation of a

distribution of shapes and colors. Then even if you don't visually represent being a pine tree, using your visual representation of the shapes and colors enables you to "get on" to the visual representation of being a pine tree. This would explain why RR is true for high-level properties despite being false for low-level properties.

Let's grant that phenomenal dependence is a relevant difference between low and high- level representation (i.e. it would explain the why RR differs in truth value). However, as I mentioned earlier, arguing that a difference would be relevant is only half the battle. The other half is to argue there *is* the difference in question. To evaluate the question, we need to figure out the nature of the dependence relation—how exactly is it that the visual representation of high-level properties depends on the visual representation of low-level properties? If there's any kind of dependence at all, the former will supervene the latter. So if it's not even plausible that high-level representation supervenes on low, it's not plausible it depends on low at all.

The *simple supervenience thesis* is just that high-level representations straight up supervene on low-level representations, period. Tim Bayne suggests this possibility in the following passage.

> [Proponents of the high-level view] might argue...that high-level phenomenal content supervenes on low-level phenomenal content, so that any change to high-level phenomenal content requires a change of low-level phenomenal content. I myself do not do find any such view particularly attractive, but I can see nothing in the liberal commitment to high-level phenomenal content which rules it out. (2011:25)

If simple supervenience is true, for there to be a high-level difference there *has to be* a low-level difference. This means there's a natural sense in which the high-level difference is explained by the low-level difference.

As far as I can tell, Bayne is right that the high-level view is *consistent* with simple supervenience. However, endorsing it is not an attractive option for proponents of PCS like him and

Siegel—it undermines the motivation we had for endorsing the view in the first place. Remember, their argument is that we should endorse the high-level view because high-level difference is the best explanation of (1). If simple supervenience true, that's much less plausible. It would ensure, for example, that Siegel's wrong when she says, "it seems implausible to suppose that there must be a change in which [low-level] properties are represented [.]" Far from being implausible, if simple supervenience is true a change in low-level properties is guaranteed.

But an alternative is suggested by the role of cognitive penetration in the explanation of high-level difference—namely, the *cognitive supervenience* thesis that high-level representation supervenes on low-level representation *plus* the cognitive component of recognition.

Cognitive supervenience wouldn't undermine the motivation for the high-level view (because low-level differences wouldn't be sufficient for high-level differences). Nevertheless, this also isn't an attractive option. To see why, suppose you're a lepidopterist (i.e. a butterfly expert) and that you have the abilities to recognize both monarchs and emperors (i.e. normally, when you see them, you can identify them immediately). Moreover, despite their strikingly similar looks (which makes them virtually indistinguishable to novices), you have trained enough so that (normally) you can immediately distinguish them. Presumably, proponents of the high-level view will say that you visually represent both the properties being a monarch and being an emperor. But suppose on this occasion you are looking at, what you know to be, a monarch and an emperor, side by side. Yet, for some reason (e.g. a morphological anomaly in these particular specimen, your viewing angle, etc.) *you are unable to tell them apart*—indeed, they look identical to you. If you visually represented one as a monarch and the other as an emperor, you *would* be able to tell them apart—they would *not* look identical to you. So it follows from the fact that you can't tell them apart that the same properties (or virtually the same—the same for all you can tell) figure in your visual representations of each. At this

point, proponents of the high-level view have four options, they can say you visually represent: (a) both as being monarchs, (b) both as being emperors, (c) neither as being a monarch or being an emperor, (d) both as being a monarch or an emperor. I'll argue none of these options is attractive.

All of these options undermine the cognitive supervenience thesis. Despite your inability to discriminate the butterflies on this occasion, you still *have* the ability to do so in general. So you have the cognitive components associated with both the ability to recognize monarchs and the ability to recognize emperors and you visually represent the low-level properties upon which the high-level ones are supposed to supervene. So it looks like the cognitive components of your recognitional abilities plus the low-level representations are not sufficient for the high-level representations.

An additional problem with (a) and (b) is that both are equally good candidates, and that's one too many. There's no reason to choose one rather than the other, so any choice would be arbitrary. Why say you represent them both as being monarchs rather than as being emperors? You may be tempted to doubt this if you're thinking that monarch are way more common in your environment than emperors. In this situation perhaps it would be arbitrary to pick (a). So let's just suppose that monarchs and emperors are, roughly, just as common in your environment. Moreover, choosing (a) or (b) would require saying that part of your experience is *illusory* and that just seems wrong. It's not as though one of the butterflies looks to you to have some property it doesn't actually have.

But perhaps you weren't tempted by those options because you think the natural thing to say is that in these circumstances, you don't visually represent either one as being a monarch or as being an emperor. But this is bad news for proponents of the high-level view. Despite the fact that you can't discriminate the monarch from the emperor *on this particular occasion*, they look different to you now then they did when you were a novice despite the fact that you're *not* visually representing the

high-level properties. So clearly, the look they take on when you've learned to recognize them *can't* be explained in terms of high-level representation.

The remaining option (d) is to say that in this situation you visually represent them both as having the disjunctive property of being a monarch or emperor. This strikes me as desperate. In addition to standard problems associated with disjunctive properties, saying you visually represent them in this situation is completely unmotivated. High-levelists can't say you normally visually represent the disjunctive property (i.e. *instead* of the natural kind properties) because it's not even remotely plausible *that* the visual representation of that property is the best explanation of the phenomenal difference that came about when you learned to recognize monarchs or emperors. So if you do visually represent the disjunctive property, it's only on occasions like this when you fail to discriminate. But then what's the motivation for saying the property is visually represented? Surely it can't be that the visual representation of the disjunctive property is the best explanation of your failure to recognize the butterflies. That stretches credulity a bit too far. Additionally, there's a pressing question about how that property came to be visually represented. It seems unlikely there's any plausible story in terms of cognitive penetration. And given my argument that the explanation must be in terms of cognitive penetration, this is an extremely pressing problem.

3.6 CONCLUSION

The goal of PCS is to provide us with reason(s) to endorse the high-level view. I have argued it fails to do this in two ways: it doesn't give us any reason to believe and, in conjunction with cognitive penetration, actually gives us reason to disbelieve.

3.7 REFERENCES

[Brogaard 2013] Brogaard, Berit: Do we perceive natural kind properties? In: Philosophical Studies 162 (2013), July, Nr. 1, S. 35–42

[Campbell 2013] Campbell, John: Susanna Siegel's the Contents of Visual Experience. In: Philosophical Studies 163 (2013), Nr. 3, S. 819– 826

[Crutchfield 2012] Crutchfield, Parker: Representing high-level properties in perceptual experience. In: Philosophical Psychology 25 (2012), April, Nr. 2, S. 279–294

[Deroy 2013] Deroy, Ophelia: Object-sensitivity versus cognitive penetrability of perception. In: Philosophical Studies 162 (2013), S. 87–107

[Dretske 1969] Dretske, Fred: Seeing And Knowing. Chicago: University Of Chicago Press, 1969. – 82

[Hawley und Macpherson 2011] Hawley, Katherine (Hrsg.) ; Macpherson, Fiona (Hrsg.): The Admissible Contents of Experience. Wiley-Blackwell, 2011

[Koksvik 2015] Koksvik, Ole: Phenomenal Contrast: A Critique. In: American Philosophical Quarterly 52 (2015), October, Nr. 4, S. 321–334

[Logue 2013] Logue, Heather: Visual experience of natural kind properties: is there any fact of the matter? In: Philosophical Studies 162 (2013), Nr. 1, S. 1–12

[Masrour 2011] Masrour, Farid: Is Perceptual Phenomenology Thin? In: Philosophy and Phenomenological Research LXXXIII (2011), Nr. 2, S. 366–397

[Matey 2013] Matey, Jennifer: You can see what 'I' means. In: Philosophical Studies 162 (2013), Nr. 1, S. 57–70

[Michotte 1963] Michotte, Albert: The Perception of Causality. Basic Books, 1963

[Pylyshyn 1999] Pylyshyn, Zenon. W.: Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception. In: Behavioral and Brain Science 22 (1999), Nr. 3, S. 341–365

[Siegel 2007] Siegel, Susanna: How Can We Discover the Contents of Experience? In: The Southern Journal of Philosophy XLV (2007), S. 127–142

[Siegel 2010] Siegel, Susanna: The Contents of Visual Experience. Oxford University Press, 2010

[Silins 2013] Silins, Nicholas: The significance of high-level content. In: Philosophical Studies (2013),

Nr. 162, S. 13–33

[Sorensen 2013] Sorensen, Roy: The Twin Towers Riddle. In: Philosophical Studies 162 (2013), S. 109–117

Complete List of References

Bayne, Tim (2011). "Perception and the Reach of Phenomenal Content." In: The Admissible Contents of Experience. Ed. by Katherine Hawley and Fiona Macpherson. Wiley-Blackwell.

Bermueacute; and Arnon Cahen (2015). "Nonconceptual Mental Content." In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Fall 2015. Metaphysics Research Lab, Stanford University.

Bitter, D´avid (2014). "Is Low-Level Visual Experience Cognitively Penetrable?" In: The Baltic International Yearbook of Cognition, Logic and Communication 9.1, pp. 1–26.

Bourget, David and Angela Mendelovici (2014). "Tracking Representationalism." In: Philosophy of Mind: The Key Thinkers. Ed. by Andrew Bailey. Continuum, pp. 209–235.

Brogaard, Berit (2013). "Do we perceive natural kind properties?" In: Philosophical Studies 162.1, pp. 35–42.

Brogaard, Berit and Bartek Chomanski (2015). "Cognitive Penetrability and High Level Properties in Perception: Unrelated Phenomena?" In: Pacific Philosophical Quarterly 96.4, pp. 469–486.

Bruner, Jerome (1947). "Value and Need as Organizing Factors in Perception." In: Journal of Abnormal and Social Psychology 42, pp. 33–44.

Campbell, John (2013). "Susanna Siegel's the Contents of Visual Experience." In: Philosophical Studies 163.3, pp. 819–826.

Cecchi, Ariel S. (2014). "Cognitive Penetration, Perceptual Learning and Neural Plasticity." In: Dialectica 68.1, pp. 63–95.

Chudnoff, Elijah (forthcoming). "Phenomenal Contrast Arguments for Cognitive Phenomenology." In: Philosophy and Phenomenological Research XCI.1, pp. 82–104.

Churchland, Paul M. (1988). "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." In: Philosophy of Science 55.June, pp. 167–87.

Crutchfield, Parker (2012). "Representing high-level properties in perceptual experience." In: Philosophical Psychology 25.2, pp. 279–294.

Degenaar, Marjolein and Gert-Jan Lokhorst (2014). "Molyneux's Problem." In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Spring 2014.

Deroy, Ophelia (2013). "Object-sensitivity versus cognitive penetrability of perception." In: Philosophical Studies 162, pp. 87–107.

Dretske, Fred (1969). Seeing And Knowing. Vol. 21. 82. Chicago: University Of Chicago Press, p. 82.

—— (1995). Naturalizing the Mind. MIT Press.

Firestone, Chaz (2013). "How "Paternalistic" Is Spatial Perception? Why wearing a Backpack Doesn't—and Couldn't—Make Hills Look Steeper." In: Perspectives on Psychological Science 8.4, pp. 455–473.

Firestone, Chaz and Brian J. Scholl (2014). ""Top-Down" Effects Where None Should Be Found: The El Greco Fallacy in Perception Research." In: Psychological Science 25.1, pp. 38–46.

Fodor, Jerry A. (1984). "Observation Reconsidered." In: Philosophy of Science 51.March, pp. 23–43.

Fodor, Jerry A. and Zenon W. Pylyshyn (1981). "How Direct is Visual Perception? Some Reflections on Gibson's 'Ecological Approach'." In: Cognition 9.2, pp. 139–96.

Gross, Steven et al. (2014). "Problems for the Purported Cognitive Penetration of Perceptual Color Experience and Macpherson's Proposed Mechanism." In: Though and Perception. Ed. by Edouard Machery and Jesse Prinz, pp. 01–30.

Hawley, Katherine and Fiona Macpherson, eds. (2011). The Admissible Contents of Experience. Wiley-Blackwell.

Koksvik, Ole (2015). "Phenomenal Contrast: A Critique." In: American Philosophical Quarterly 52.4, pp. 321–334.

Kuhn, Thomas S. (1962). The Structure of Scientific Revolutions. The University of Chicago Press.

Levin, Daniel T. and Mahzarin R. Banaji (2006). "Distortions in the perceived lightness of faces: The role of race categories." In: Journal of Experimental Psychology: General 135.4, pp. 501–512.

Logue, Heather (2013). "Visual experience of natural kind properties: is there any fact of the matter?" In: Philosophical Studies 162.1, pp. 1–12.

Macpherson, Fiona (2012). "Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism." In: Philosophy and Phenomenological Research 84.1, pp. 24–62.

Masrour, Farid (2011). "Is Perceptual Phenomenology Thin?" In: Philosophy and Phenomenological Research LXXXIII.2, pp. 366–397. Matey, Jennifer (2013). "You can see what 'I' means." In: Philosophical Studies 162.1, pp. 57–70.

McCauley, Robert N. and J. Henrich (2006). "Susceptibility to the Muller-Lyer Illusion, Theory-Neutral Observation, and the Diachronic Penetrability of the Visual Input System." In: Philosophical Psychology 19.1, pp. 79–101.

McClelland, Tom (2016). "Gappiness and the case for liberalism about phenomenal properties." In: The Philosophical Quarterly 264, pp. 536–558.

Michotte, Albert (1963). The Perception of Causality. Basic Books. O'Callaghan, Casey (2010). "Experiencing Speech." In: Philosophical Issues 20, pp. 305–332.

Peacocke, Christopher (2001). "Does Perception Have a Nonconceptual Content?" In: Journal of Philosophy

98.5, pp. 239–264. Price, Richard (2009). "Aspect-Switching and Visual Phenomenal Character." In: Philosophical Quarterly 59.236, pp. 508–518.

Pylyshyn, Zenon. W. (1999). "Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual Perception." In: Behavioral and Brain Science 22.3, pp. 341–365.

Raftopoulos, Athanassios (2009). Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy? MIT Press.

Scholl, Brian J. and Alan M. Leslie (1999). "Modularity, Development and Theory of Mind"." In: Mind and Language 14.1, pp. 131–153.

Siegel, Susanna (2006). "Which Properties Are Represented in Perception?" In: Perceptual Experience. Ed. by Tamar S. Gendler and John Hawthorne. Oxford University Press, pp. 481–503. – (2007). "How Can We Discover the Contents of Experience?" In: The

Southern Journal of Philosophy XLV, pp. 127–142.

—— (2010). The Contents of Visual Experience. Oxford University Press.

—— (2016). "The Contents of Perception." In: The Stanford Encyclopedia of Philosophy. Ed. by Edward N. Zalta. Winter 2016. Metaphysics Research Lab, Stanford University.

Siewert, Charles (1998). The Significance of Consciousness. Princeton University Press.

Silins, Nicholas (2013). "The significance of high-level content." In: Philosophical Studies 162, pp. 13–33.

Sorensen, Roy (2013). "The Twin Towers Riddle." In: Philosophical Studies 162, pp. 109–117. Stokes, Daniel (2013). "Cognitive Penetrability of Perception." In: Philosophy Compass 8.7.

Stokes, Dustin (2015). "Towards a Consequentialist Understanding of Cognitive Penetration." In: Cognitive Penetrability (Oxford University Press). Ed. by A. Raftopoulos and J. Zeimbekis.

Tye, Michael (2000). Consciousness, Color, and Content. MIT Press.

Wu, Wayne (2013). "Visual Spatial Constancy and Modularity: Does Intention Penetrate Vision?" In: Philosophical Studies 165.2, pp. 647–669.

Zeimbekis, John (2013). "Color and Cognitive Penetrability." In: Philosophical Studies 165.1, pp. 167–175.