

PHL340 Handout 5: The Reach of Perceptual Content

§1 The Question of Perceptual Reach

We've surveyed accounts of perceptual experience that answer the metaphysical problem of perception: indirect realism, direct realism, representationalism, and so on. These are accounts of *what it is* for a subject to undergo a perceptual experience as of an external world.

Today our focus shifts to the question of *which* properties does perceptual experience reveal (or, in Siegel's lingo, which properties does perceptual experience represent)? In particular, philosophers ask: which properties does *visual* experience represent?

Philosophers of perception forge a rough distinction between *low-* and *high-level properties*.

Low-level properties include shape, size, and colour. Almost everyone agrees that visual experience reveals these features.

Candidate high-level properties include *kind properties*—such as being an elm or being a tiger—as well as the property of being mind-independent.

Another (imperfect) division between these two types of property relies upon perceptual priority. If perception represents high-level properties, many assume that it does so partly in virtue of representing a bunch of low-level properties.

Example: if visual experience represents the property of being a tiger, it does so partly in virtue of representing certain colours, shapes, and sizes.

Note: this division remains imperfect because there may be low-level properties whose representation in visual experience depends in part upon the representation of other low-level properties (e.g. a visual experience could not represent an object's shape without also representing its colour).

Some philosophers adopt a 'conservative' view and insist that visual experience only represents low-level properties. Others, like Siegel and Bayne, argue for a 'liberal' stance and insist that visual experience can represent high-level properties (this is Siegel's *Thesis K*).

Historical Precedent: David Hume famously argued that visual experience does not present genuine causal relations. But this is just a claim about which properties perceptual experience manages to reveal. A good question, moving forward, is whether the argument Siegel provides could extend to causal high-level properties.

This debate about the reach of perception is not confined to properties. *Singularists* and *generalists* about perceptual experience disagree over whether visual experience represents particular objects (singularists say it does, generalists say it does not). Generalists are therefore conservative about the reach of perception: they think that visual experience represents a range of properties, but fails to represent particular objects as such. (Siegel argues for singularism elsewhere, in keeping with her general 'liberal' stance).

§1.1 Direct Realism and Representationalism

All of this talk about which properties visual experience represents can make it seem as though the question only applies to representationalists. *Representationalism*, recall, is the view that the representational content of perceptual experience exhausts its phenomenal character.

Direct realists, by contrast, hold that the phenomenal character of a veridical perceptual experience consists in the perceived objects, properties, and a subject's perceptual perspective upon these objects and properties.

A cursory glance at direct realism may suggest that our question has no application to the view, since direct realists as yet make no mention of representational content.

But there are two ways to resist this suggestion.

First, direct realists needn't deny that perceptual experience possesses representational content. They need merely deny the representationalist claim that this content exhausts the phenomenal character of perceptual experience. For instance, a direct realist that posits a difference in kind between veridical, illusory, and hallucinatory experience might appeal representational perceptual content.

The other option would be to re-construe the original question within a direct realist framework. Instead of asking which properties perceptual experience represents, direct realists can ask which properties can help constitute the phenomenal character of a veridical visual experience.

Siegel's argument for *Thesis K* relies upon a weak claim about the relationship between a perceptual experience's representational content and its phenomenal character: the latter must supervene upon the former.

Facts in a domain A supervene upon facts in a domain B iff necessarily, there exists a difference in the A facts only if there exists a difference in the B facts (e.g. most accept that mental facts supervene upon physical facts, since few can are willing to countenance a possibility in which all the physical facts are the same but the mental facts differ from our own).

A direct realist can accept a version of this supervenience claim: if one holds the object perceived and the perspectival factors fixed, then the phenomenal character of a veridical perceptual experience supervenes upon the properties perceived. *A nice question*: does Siegel's later argument go through even with this direct realist supervenience claim? I'm not aware of anyone who discusses this question in print.

Note: if you read the Bayne paper, you'll see him discuss our question about the reach of perception as a question about the reach of the *phenomenal content* of visual experience. The phenomenal content of visual experience is that aspect of an experience's representational content that supervenes upon the experience's phenomenal character. *Good Question*: is this what Siegel has in mind?

§2 Phenomenal Contrast Method

Siegel's argument for Thesis K—the claim that some high-level properties are represented by some visual experiences—is an instance of a more general argument type: *the phenomenal contrast strategy*.

The strategy always begins with an introspective claim that two experiences possess distinct phenomenal characters (hence 'phenomenal contrast'). These are meant to be experiences that do not differ with respect to the low-level properties they represent. Armed with the intuitive introspective claim, the next move is to argue that the best explanation of the phenomenological distinction is that one experience, but not the other, includes a visual experience that represents a high-level property.

Siegel uses the phenomenal contrast strategy to argue that the kind property of being a pine tree can be represented by visual experience. Here is how she motivates the initial introspective claim:

'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 capacity you gain, the salience of the trees emerges gradually. Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experience you had before and after the recognitional disposition was fully developed. [Siegel p. 491, emphasis added]

Siegel provides a rather handy reconstruction of the phenomenal contrast strategy on p. 491 of her paper. But to understand this formal statement, we need some additional terminology:

Visual experiences are assumed to have a proprietary phenomenal character—what Siegel calls a *sensory* phenomenal character. But a subject's sensory phenomenal character does not necessarily exhaust what it is like for her at a given time. So distinguish a subject's *overall experience* at a time from an *aspect of her total experience that possesses a sensory phenomenal character* (i.e. her visual experience).

Following Siegel, let's say E1 is the subject's sensory experience of pine trees before she learns to recognize them, and E2 the sensory experience of pine trees that occurs after the subject learns to recognize these trees as pines.

1. The overall experience of which E1 is an aspect (or part) differs phenomenologically from the overall experience of which E2 is a part [*Introspective Claim*]
2. If these two overall experiences differ phenomenologically, then E1 and E2 differ phenomenologically.
3. If E1 and E2 differ phenomenologically, they differ in content.
4. If E1 and E2 differ in content, there is a difference with respect to the kind properties represented by these experiences.

1-4 entail Thesis K, so the task of a proponent of the phenomenal contrast strategy is to motivate 1, and provide arguments for 2-4.

Siegel and her allies generate different versions of 1-4 by substituting different high-level properties for the pine tree kind property. In her later work she uses this strategy to argue for the claim that visual experience represents particular objects, causal relations, mind-independent, and a variety of kind properties.

Note: the plausibility of 1-4 will depend greatly upon our choice of high-level property. So even if the argument works for one sort of high-level property, it may not work for others.

Siegel argues for 2-4 by inference to the best explanation. In each case, she argues that the difference mentioned in the antecedent of the conditional (e.g. that E1 and E2 differ in content) cannot be adequately explained except by the fact mentioned in the consequence (e.g. that there is a difference with respect to the kind properties represented by E1 and E2).

Inference to the best explanation (or *abduction*) should be familiar from other philosophical and scientific contexts. You can think of it as having three steps:

- A. Identify the *explanandum* (i.e. that which must be explained).
- B. Agree upon criteria for what would count as an adequate explanation of the *explanandum*.
- C. Apply the criteria to candidate explanations of the explanandum to find the explanation that best satisfies the criteria.

If you read Siegel's paper, you'll rarely see her explicitly state the criteria mentioned at B. Instead, she introduces criteria when she wishes to argue against a competing explanation of a given *explanandum*.

Questions to ask whenever faced with an inference to the best explanations:

- (i) do all parties agree upon the *explanandum*?
- (ii) Do all parties agree upon the criteria mentioned at B (and if not, *ought* one side agree with the other)?
- (iii) Are there unexplored alternative explanations that could, given the mutually accepted criteria, outperform an opponent's explanation?

Siegel spends much of her paper arguing against alternative explanations of the antecedents of 2-4. Here are the broad options:

- Deny 2 on the grounds that while the overall experiences differ phenomenologically, this difference ought to be explained by a phenomenological difference between non-sensory aspects of the overall experiences.
 - Siegel considers two non-sensory aspects of our overall experience: implicit or explicit cognitive attitudes (e.g. beliefs, hunches, etc.); background non-cognitive non-sensory aspects (e.g. moods, feelings, etc.).
- Deny 3 on the grounds that while E1 and E2 differ phenomenologically, this difference ought to be explained by a non-representational difference between E1 and E2.
 - Hardcore representationalists will simply deny that there could be non-representational differences between E1 and E2 that make a difference to their respective phenomenal characters. But other philosophers accept a variety of non-representational differences. For instance, it might be that the two experiences differ only with respect to how a subject deploys perceptual

attention. Another option is that the experiences possess different ‘raw feels’ (i.e. non-representational qualitative aspects of perceptual experience).

- Deny 4 on the grounds that the difference in content between E1 and E2 ought to be explained purely in terms of their representation of low-level properties.
 - Siegel’s example: the representation of ‘gestalt’ properties (i.e. complexes of lower-level properties that are not themselves high-level properties).

Question: can you think of potential explanations that Siegel does not consider? Also, take a look at her paper to see which criteria she uses to exclude alternative explanations.

§3 Alternative Arguments

The phenomenal contrast method, as deployed by Siegel, does not exhaust the potential arguments for and against Thesis K.

§3.1 *Argument from Agnosia*

Tim Bayne appeals to cases of associative agnosia to generate his own contrast argument (cf. 390-392). ‘Agnosia’ is a general term for perceptual impairments that cannot be blamed on simple sensory malfunction. Associative agnosia occurs when subjects are otherwise unimpaired, but nevertheless systematically fail to recognize objects as belonging to categories with which they are familiar.

Associative agnosia provides the materials for a contrast argument because subjects with it suffer from a genuinely *perceptual* impairment, but their perception of low-level properties remains unaffected. We know this because the subjects retain practical abilities that rely upon perception of these low-level properties.

Bayne thus provides something like an empirical argument for premiss 1 from the phenomenal contrast strategy. What we know about associative agnosia also promises to help Bayne exclude the sort of alternative explanations that threaten 2-4 (see pp. 392-396 of his paper).

§3.2 *Argument From The Limits of Illusion*

Alex Byrne (2009) provides something like the following argument against Thesis K:

1. If visual experience represents that an object possesses a property F, then it will be possible for the experience to count as illusory (i.e. if the object is not F).
2. But we do not have visual experiences that are illusory with respect to high-level properties
3. So visual experience does not represent high-level properties

Why accept 2?

Byrne claims that vision science, whose job it is to study visual illusions, does not recognize illusions of high-level properties. In some ways, this move looks like a variant of what Tye says on p. 63 of his chapter about the reach of perception.

Byrne also argues that putative illusory experiences involving high-level properties lack phenomenological markers of genuine illusory experience. For example, physical illusions persist even after a subject learns about them (e.g. the Muller-Lyer illusion), but the same does not apply to an elm tree that looks to be a pine. See also p. 482 of Siegel's paper for a discussion of Thesis K and illusory experience.

Of course, a proponent of Thesis K could try to turn Byrne's argument against him by arguing that 2 is false (a result that when combined with 1 would entail that 3 is false). So his argument brings up nice questions about the extent to which vision science can fruitfully inform an inquiry into the reach of perception.

§4 The Significance of The Reach of Perception

Why bother with this question about which properties visual experience represents? Siegel discusses a handful of implications of Thesis K, but we're going to look at only two of them.

§4.1 Informational Encapsulation and Cognitive Penetration

Cognitive scientists, psychologists, and empirically-minded philosophers have long thought that a large portion of the visual system—the system responsible for converting sensory input into visual information that could be used for action and cognition—cannot be directly influenced by cognition (i.e. belief, judgment, desire, hope, expectation, etc.). Cognition can influence where we look and (to some extent) how we direct our perceptual attention, and it participates in the 'downstream' manipulation of visual information, but (so the story goes) it cannot otherwise influence the workings of the visual system (cf. Pylyshyn and Fodor).

In cognitive science lingo, this autonomous portion of the visual system—often called the 'early visual system'—is informationally encapsulated.

If Siegel is right, her pine tree case involves a cognitive capacity—in this case a recognitional capacity—influencing the phenomenology of visual experience. This result could be interpreted in at least two ways. Either the case constitutes an exception to the informational encapsulation of the early visual system (cf. p. 501 of Siegel), or the outputs of the early visual system are not visual experiences (i.e. the cognitive influence must occur between the exit of the early visual system and visual experience). Exceptions to informational encapsulation are cases of cognitive penetration.

If Thesis K opens the door to widespread cognitive penetration, a number of hard questions arise about how the visual system works, and about implications of cognitive penetration for the epistemological significance of visual experience.

§4.2 Immediate Perceptual Justification

Next week we're going to look at Jim Pryor's influential account of how perceptual experience *immediately justifies* the beliefs we form on the basis of experience. A belief is immediately justified iff its justification does not depend upon justification for any other beliefs. Given that the scope of immediate perceptual justification depends in part upon which properties perception represents, where one stands on Thesis K will constrain which beliefs could count as immediately justified by perceptual experience. See pp. 487-488 of Siegel's paper for discussion.