

Perceptual Learning and the Contents of Perception

By Kevin Connolly

Abstract: Suppose you have recently gained a disposition for recognizing a high-level kind property, like the property of *being a wren*. Wrens might look different to you now. According to the Phenomenal Contrast Argument, such cases of perceptual learning show that the contents of perception can include high-level kind properties such as the property of *being a wren*. I detail an alternative explanation for the different look of the wren: a shift in one's attentional pattern onto other low-level properties. Philosophers have alluded to this alternative before, but I provide a comprehensive account of the view, show how my account significantly differs from past claims, and offer a novel argument for the view. Finally, I show that my account puts us in a position to provide a new objection to the Phenomenal Contrast Argument.

1. Introduction

Some philosophers hold that we perceive high-level kind properties (in addition to low-level properties like colors, shapes, size, orientation, illumination, textures, and bare sounds). High-level kind properties include natural kind properties like *being a wren* and artificial kind properties like *being a table*. Susanna Siegel has recently offered an argument (the *Phenomenal Contrast Argument*) for the conclusion that such properties are represented in perception (2006; 2010). Roughly and briefly, the idea is that since tables or wrens can look phenomenally different to someone once they become disposed to recognize them, those properties are represented in perception.¹ I detail an alternative explanation for the phenomenal difference: a shift in one's attentional pattern onto other low-level properties. Philosophers have alluded to this alternative before (the first being Price, 2009), but I provide a comprehensive account of the view, show how my account significantly differs from past claims, and offer a novel argument for the view. Finally, I show that my account puts us in a position to provide a new objection to Siegel's Phenomenal Contrast Argument.

2. The Phenomenal Contrast Argument

¹ "Disposed to recognize" and "recognitionally disposed" are Siegel's locutions. I adopt this language throughout in presenting her argument and in replying to it.

Plausibly, when looking at a wren, the perception of an expert birdwatcher is phenomenally different from the perception of a layperson, even when viewed under the exact same background conditions. That is to say, it seems plausible that the expert birdwatcher exhibits what Eleanor Gibson calls *perceptual learning*: “[a] relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array” (1963, p. 29). Due to practice or experience with properties in their expert domain, experts in some fields perceive the world differently from non-experts. Cabernet Sauvignon tastes different to a wine connoisseur than to a novice. Beethoven’s Ninth Symphony sounds different to a conductor than to an untrained listener. What it is like to see a wren is different for an expert birdwatcher than a non-expert.

In such cases, Susanna Siegel argues that there is a phenomenal difference in what she calls the “sensory phenomenology,” that is, in the phenomenology pertaining to the properties that sensory experience represents, properties like colors, shapes, and perhaps also high-level kind properties. She contrasts sensory phenomenology with phenomenology associated with imagination, with emotions, with bodily sensation, with background phenomenology (as with drunkenness or depression), and with non-sensory cognitive functions (as with a feeling of familiarity) (2006, p. 492). There may be changes in those kinds of phenomenology as well, but Siegel’s concern is just with sensory phenomenology. Very roughly, Siegel seems to have in mind that for vision, the sensory phenomenology is that visual phenomenology that typically changes when you move your head from side to side (excluding, say, the proprioceptive changes that also occur). The idea is that with some cases of expertise, this phenomenology differs between an expert and a layperson.

One plausible explanation for the difference in sensory phenomenology is as follows. Some experts possess recognitional dispositions for each of several properties in their expert domain. The expert birdwatcher is disposed to recognize a House Wren and a Marsh Wren. The wine connoisseur is disposed to recognize Merlot and Pinot Noir. The symphony conductor is disposed to recognize a bassoon and a clarinet. The layperson, on the other hand, is not disposed to recognize those properties. So, suppose that an expert birdwatcher and a layperson both see a wren. Suppose that their perceptions differ in terms of their sensory phenomenology, even though they both view the wren under the exact same background conditions. One plausible explanation for this is that the expert possesses and exercises a recognitional disposition for wrens, while the layperson lacks that recognitional disposition altogether.

So far, we have been discussing cases of full-blown expertise. But a similar point can be made in cases where subjects fall well short of that standard. As Gary Hatfield reminds us, “We are not born recognizing books and tables, but we learn to categorize these artifacts and to determine at a glance that a table is an old one of good quality” (2009, p. 125). This applies to experts and non-experts alike. And Charles Siewert takes the point further. Regardless of whether one is an expert or not, there is “a difference between the way things look to us when they merely look somehow shaped, colored, and situated, and how they look to us when they look recognizable to us as belonging to certain general types...” (1998, p. 256). As an example, consider someone who has a recognitional disposition for wrens, but is not an expert birdwatcher. According to Siewert, when looking at a wren, even though the subject is not an expert birdwatcher, her sensory phenomenology might differ from someone who lacks a recognitional disposition for wrens. Supposing this is right, forget about the expert birdwatcher—who likely has a large repertoire of recognitional dispositions for birds—and just

consider an ordinary subject instead. Also, to further simplify things, instead of comparing our subject's perception with the perception of another subject, think about an intra-personal case across time.

Now consider Siegel's core argument that perception represents high-level kind properties: the Phenomenal Contrast Argument. (Note that Siegel uses pine trees as her example, while I am using wrens, and she calls such properties as pine trees and wrens "K-properties," while I am calling them "high-level kind properties"). The argument runs as follows. Suppose our ordinary subject acquires a recognitional disposition for wrens. Contrast her perceptions before and after she gains that disposition. Plausibly, after she gains the disposition, even if she looks at exactly the same scene of a wren, the sensory phenomenology of her perception has changed. Given that the perceptions differ in their sensory phenomenology, the argument continues, they differ in their content, that is, in what the perceptions represent. Specifically, they differ with respect to the high-level visual property that she is now disposed to recognize, namely, the property of *being a wren*. The argument generalizes, *mutatis mutandis*, to other high-level kind properties. High-level kind properties can be represented in perception.

Here is Siegel's more formal expression of the argument. To preface the argument, she writes:

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. Gaining this recognitional disposition is reflected in a phenomenological difference between the visual experiences had before and after the recognitional disposition was fully developed. (2010, p. 100)

To formalize her argument, she starts by labeling each of the visual experiences in the previous example:

Let E1 be the visual experience had by a subject *S* who is seeing the pine trees before learning to recognize them, and let E2 be the visual experience had by *S* when *S* sees the pine trees after learning to recognize them. E1 and E2 are visual parts of *S*'s overall experiences at each of these times. The overall experience of which E1 is a part is the contrasting experience, and the overall experience of which E2 is a part is the target experience. (p. 100)

The Phenomenal Contrast Argument then runs as follows:

- (0) The target experience differs in its phenomenology from the contrasting experience.
- (1) If the target experience differs in its phenomenology from the contrasting experience, then there is a phenomenological difference between E1 and E2.
- (2) If there is a phenomenological difference between E1 and E2, then E1 and E2 differ in content.
- (3) If there is a difference in content between E1 and E2, it is a difference with respect to K-properties represented in E1 and E2. (p. 101)

Siegel's final summation of the argument is helpful: "I've argued that gaining a disposition to recognize K-properties can make a difference to visual phenomenology, and that this difference is accompanied by a representation of K-properties in visual experience" (p. 113).

3. The Attentional Reply to the Phenomenal Contrast Argument

3.1 The Attentional Reply

Richard Price (2009) replies to the Phenomenal Contrast Argument in the following way:

After one learns to recognize pine trees, one starts to attend to those features of pine trees that distinguish them from other trees, for instance, the colour or thickness of the bark. Acquiring a recognitional disposition for pine trees will cause one's patterns of attention to shift when one looks at a grove containing pine trees and other sorts of trees. (p. 516)

Price's reply to Siegel's pine tree case makes a compelling suggestion. The claim is perhaps best put as follows: The phenomenal difference in the pine tree case is explicable in terms of a shift in one's attentional pattern onto other low-level properties, and so, contrary to what Siegel concludes, we do not need to appeal to high-level kind properties. While Price's attentional reply

is brief, I will provide a robust account of the attentional difference that occurs in the pine tree case. My account draws on the psychology of attention and learning, and I will show that it provides a compelling explanation of changes in expert perception without an appeal to high-level kind properties. My view is novel in the following way. Price writes, “*After* one learns to recognize pine trees, one starts to attend to those features of pine trees that distinguish them from other trees, for instance, the colour or thickness of the bark” (2009, p. 516, italics added for emphasis). The central claim of my argument is that it is not *after*, but *before* one learns to recognize pine trees that one starts to attend to the distinguishing features of pine trees (like the color or thickness of the bark). On my account attending to those features for the first time is part of the very process that enables you to develop a recognitional disposition.

My account has application beyond the Phenomenal Contrast Argument. Take Zenon Pylyshyn’s defense of the *cognitive impenetrability* of visual perception, the view that “an important part of visual perception, corresponding to what some people have called early vision, is prohibited from accessing relevant expectations, knowledge, and utilities in determining the function it computes...” (1999, p. 341). Pylyshyn defends cognitive impenetrability in part by arguing that some putative counter-examples are explicable just in terms of “the allocation of attention to certain locations or certain properties *prior to* the operation of early vision” (p. 344). Chicken sexers, for instance, train for years to learn the lucrative skill of being able to identify the sex of day-old chicks. Pylyshyn argues that the case of chicken sexers is *not* a case where their knowledge directly affects the content of what they see. Instead, their case is explicable fully in terms of the way they allocate their attention (p. 359). The model of perceptual learning that I will offer can help to illuminate how the allocation of attention occurs in such cases,

thereby building on the attentional strategy that Pylyshyn has offered to defend the cognitive impenetrability of perception.

3.2 The “Blind Flailing” Model of Perceptual Learning

When someone learns to recognize a pine tree or a wren, this is a case of perceptual learning, “[a] relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array” (Gibson, 1963, p. 29). Siegel’s pine tree case, for instance, involves a relatively permanent change in your perception of pine trees following experience with them. But one question that Siegel overlooks is why exactly these changes occur in the first place. What purpose does a change in one’s perception serve? In the psychology literature, the answer is fairly straightforward. Perceptual changes occur so that we can better perform the cognitive tasks that we need to do. To ideally perform cognitive tasks, it is better for perceptual systems to be flexible, rather than hardwired. As Robert Goldstone explains:

One might feel that the early perceptual system ought to be hardwired—it is better not to mess with it if it is going to be depended upon by all processes later in the information processing stream. There is something right with this intuition, but it implicitly buys into a “stable foundations make strong foundations” assumption that it is appropriate for houses of cards, but probably not for flexible cognitive systems. For better models of cognition, we might turn to Birkenstock shoes and suspension bridges, which provide good foundations for their respective feet and cars by flexibly deforming to their charges. Just as a suspension bridge provides better support for cars by conforming to the weight loads, perception supports problem solving and reasoning by conforming to these tasks. (2010, p. v)

Perceptual systems are flexible rather than hardwired so that they can better support cognitive tasks. This flexibility allows for perceptual learning to occur, that is, for there to be relatively permanent perceptual changes such as the one that happens in Siegel’s pine tree case (see Connolly, 2013, and Connolly, forthcoming, for more on other such changes). These changes

happen for a reason—so that they can enable cognitive processes, processes like the ability to recognize a pine tree.

The perceptual learning process is often a low-level process, in which “our perceptual abilities are altered naturally through an automatic, non-conscious process” (Goldstone, Landy, and Brunel, 2011, p. 5). One key feature of the perceptual learning process is that the learning often occurs randomly. As Goldstone, Landy, and Brunel put it, “If a random change causes important discriminations to be made with increasing efficiency, then the changes can be preserved and extended. If not, the changes will not be made permanent” (2011, p. 5). Goldstone, Landy, and Brunel advocate a simple model of perceptual learning, by which perceptual learning occurs through a process of random variation, followed by reinforcement. They call this the “blind flailing” model, named after the fact that infants flail their arms randomly during the process of learning motor control. As Goldstone, Landy, and Brunel summarize the infant motor learning process (drawing from Smith and Thelen, 1993), “The flails that are relatively effective in moving the arms where desired are reinforced, allowing an infant to gradually fine-tune their motor control” (2011, p. 5). Analogously, in perceptual learning, those random changes that cause important perceptual discriminations to be made are reinforced and selected, allowing us to fine-tune our perceptual systems.

While the blind flailing model uses infant flailing as an analogy, the model itself is intended to be a model for visual perception. And while Goldstone, Landy, and Brunel never explicitly clarify this, there is no reason to think that the blind flailing model applies only to infants. The model is fully consistent with the development of all kinds of perceptual expertise exhibited by adults, including the changes that occur when you learn to recognize wrens or pine trees.

Applying the “blind flailing” model to the wren case, suppose you are not disposed to recognize wrens, but a wren is in your visual field. You attend to the wren in various ways. If one of those ways of attending causes an important and efficient discrimination, then that way of attending is preserved and extended. For instance, if you attend to the color of the wren’s plumage, and that way of attending enables you to discriminate wrens from other birds, then that way of attending is reinforced. On the other hand, many other ways of attending cause no important discrimination, in which case, those ways of attending are discarded. The same story can be told, *mutatis mutandis*, for Siegel’s pine tree case. I will go on to argue that once we accept the blind flailing hypothesis, we end up with a problem case that Siegel’s view is unable to handle.

The blind flailing process involves the random variation of attentional patterns plus the selection of a useful pattern. One attends to a wren in all sorts of ways, and the useful attentional pattern gets selected and reinforced. The upshot of this process is a shift in one’s attentional pattern. Psychologists refer to this as a change in “attentional weighting.” The idea is that the weight of attention can change by “increasing the attention paid to perceptual dimensions and features that are important, and/or by decreasing attention to irrelevant dimensions and features” (Goldstone, 1998, p. 588). For instance, you might increase attention to the color of the wren’s plumage, while decreasing attention to the shape of its tail. If so, you have changed your original attentional weighting.

Since there are many different kinds of attention, shifts in attentional weighting can be quite sophisticated. Visual attention can be focal, acting like a spotlight (or several spotlights) on a stage, but it can also be diffuse at times, not simply centered on single points. When we look at the starry sky, for instance, we need not attend to just single stars. We can attend to large

portions of the sky. There may be a diffuse kind of attention in addition to focal attention (see Prinz, 2010, p. 318). Furthermore, attention can follow eye movements, but it does not always follow eye movements (for example, see Carrasco, Ling, and Read, 2004, on the perceptual effects of covert attention without eye movement). Visual attention can be overt and follow eye-movements, or it can be covert and not follow eye-movements. Since there are many kinds of attention, this makes it more plausible that changes in perceptual phenomenology can be explained in terms of changes in the way one attends.

On the blind flailing model, perceptual learning occurs in a way similar to how natural selection occurs. Where the process of natural selection selects a trait, the process of blind flailing selects a way of attending. In the case of perceptual learning, at first your attentional pattern is varied. Those ways of attending that are helpful get selected and preserved, while those ways of attending that are unhelpful are discarded. Like natural selection, the process begins with random variation, and ends with the selection of something useful.

One significant feature of the blind flailing model is that subjects are typically unaware of the changes that are occurring at the level of attention (see Goldstone, Landy, and Brunel, p. 5). This is because the way that attentional changes get selected is through a result that happens in a different domain—at the phenomenal level. If attentional changes yield an important discrimination at the phenomenal level, they are preserved. If not, they are discarded. Since attentional changes are selected at the level of phenomenology, one can notice a phenomenal change without noticing that the source of that change is an attentional change.

With this in mind, recall Pylyshyn's strategy of explaining some putative cases of cognitive penetration by arguing that the cases are explicable just in terms of the allocation of attention prior to perception (1999, p. 344). One such putative case is the case of chicken sexers.

Expert chicken sexers train for years to learn to accurately identify the sex of day-old chicks. But as Pylyshyn explains, research shows that the case of chicken sexers is explicable fully in terms of the way they allocate their attention (p. 359). Interestingly, however, chicken sexers are wholly unaware of this fact. They are entirely unaware of what has happened at the level of attention. The case of chicken sexers would seem to be a paradigm case of blind flailing. They attend in all sorts of ways, and get feedback. Their attentional patterns get tested on whether they yield important discriminations between the sexes of chicks. The helpful patterns are selected and reinforced. And this entire process happens without any knowledge of the attentional changes that are occurring.

Siegel's pine tree case has many of the features of a standard blind flailing example as well. Recall her description of the case:

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. (2010, p. 100)

Here are three features that suggest the pine tree case is a standard case of blind flailing. First, in Siegel's description, someone points out the pine trees to you, but then you are left to your own devices. There is no overt direction. This suggests that the perceptual improvements that then occur are achieved through a blind process—blind flailing. Second, the phenomenal change occurs over a long timeframe (Siegel says “some weeks”). As evidenced by the chicken-sexing case, attentional shifts without overt direction take time. For a difficult task like chicken-sexing, chicken sexers estimate that it takes about 2.4 months for them to attend in a way that allows them to identify the sexes at a 95% success rate (Biederman and Shiffrar, 1987, p. 643).

Reasonably, an attentional shift for pine trees might take some weeks, as Siegel's description

stipulates. Third, while the subject in the pine tree case is aware of a change at the phenomenal level, there is no awareness of the source of that change. This is standard for blind flailing, where subjects are typically unaware of the changes that are occurring at the level of attention.

The fact that one can be entirely unaware of attentional changes may go some way toward explaining why the attentional explanation is not the first explanation to come to mind in cases like the pine tree case. We may remember what something used to look like, but we remember how we used to attend to it much less frequently. We often do not notice how we are attending at all, let alone notice it, remember it, and compare it to how we are attending at a later time.

3.3 The Revised Attentional Reply

The Phenomenal Contrast Argument trades on a contrast in your perceptions before and after you acquire a recognitional disposition (say, for wrens). But when we understand the wren case in terms of the “blind flailing” model, it is not your disposition to recognize wrens that improves your perception. Rather, your perception of wrens improves through a random change, and that improvement enables you to become disposed to recognize them.

Notice now that we have a novel attentional reply to the Phenomenal Contrast Argument—one that is unique when compared with Price’s original reply. His claim was that “*After* one learns to recognize pine trees, one starts to attend to those features of pine trees that distinguish them from other trees...” (2009, p. 516, italics added for emphasis). But according to the “blind flailing” model, one starts to attend to those distinguishing features of pine trees *before* one learns to recognize pine trees. In fact, this is part of the very process that enables one to recognize pine trees perceptually in the first place. Your attentional pattern to pine trees changes. This gives the pine tree a new look to you. And the new look of the pine tree is part of

what enables you to become disposed to recognize pine trees. Before you had never much noticed them. But now that they have a new look, this helps you to become disposed to recognize them.

The same attentional pattern can be cued in more than one way. A recognitional disposition is one way to cue a particular attentional pattern. If you have a recognitional disposition for pine trees, for instance, this might cue you to attend to pine trees in a particular way. But that is not to say that the same attentional pattern cannot be cued through some other means. The pattern arises first through blind flailing. It gets selected because it is useful, enabling you to recognize a pine tree. Then you use that recognitional disposition as a shorthand cue to redeploy that attentional pattern. Just as the same thought can be cued in two different ways, through a long-winded way such as “The teacher of Alexander the Great,” or through a short-winded way like “Aristotle,” the same attentional pattern can be cued through the long process of blind flailing, or the shorter process of deploying a recognitional disposition.

Here is the important upshot. By modifying the attentional reply from Price’s original formulation, we have a new argument against Siegel. If wrens (or pine trees) look a new way to us *first*, and that look enables us to then become disposed to recognize them, then this spells trouble for the Phenomenal Contrast Argument. This is because the Phenomenal Contrast Argument tries to explain the new way that a wren looks *in terms of* the recognitional disposition. After all, it is the perceiver’s possession of that recognitional disposition that provides the compelling reason to conclude that the perception represents the property of *being a wren*. As Siegel herself summarizes, “I’ve argued that gaining a disposition to recognize K-properties can make a difference to visual phenomenology, and that this difference is accompanied by a representation of K-properties in visual experience” (p. 113). But if a

perceiver could have that same type-perception without having a recognitional disposition for wrens, then we would need another explanation for the new look of the wren.

When you have a recognitional disposition for wrens, it might seem that the exercise of that disposition is constitutive of your sensory-phenomenal character when you look at a wren. But that is only because there is a strong correlation between the two. The idea is this: typically you get that phenomenal character only when you exercise the right recognitional disposition. But in fact, the exercise of that disposition is only contingently related to your phenomenal character. After all, you can get that phenomenal character through the blind flailing process, even without having the relevant recognitional disposition. The exercise of the recognitional disposition is not necessary for that phenomenal character since you can get that phenomenal character by attending in the right way, even without having the relevant recognitional disposition.

Consider the following analogy. Suppose the existence of some hypothetical factor x , which is involved in the development of lung cancer. Suppose that all causal chains involving smoking go through factor x on their way to lung cancer. Furthermore, suppose that there is a causal chain to lung cancer that does not involve smoking, but does involve factor x . In such a case, it only appears that smoking is directly causing lung cancer, since there is a strong correlation between the two. But actually, smoking causes lung cancer only indirectly, through factor x . Smoking is not necessary for developing lung cancer since you can get lung cancer through factor x , even without smoking.

Analogously, in the wren case, it seems at first glance that your recognitional disposition is directly causing your phenomenal character. But this is only because there is a strong correlation between the two. In fact, that recognitional disposition causes your phenomenal

character only indirectly, through a way of attending. The recognitional disposition is not necessary for that phenomenal character. Again, you can get that phenomenal character by blindly flailing into it, even without having the relevant recognitional disposition.

What then happens when you then acquire a recognitional disposition for wrens? Consider an analogy with learning a backhand in tennis. According to one plausible account, learning a backhand creates a new capacity in you. It gives you a new ability to hit backhands. That might seem reasonable enough. But it stands in contrast to a second account. According to this second account, you don't actually gain a new ability. Rather, you already have the ability to hit backhands. Learning a backhand just selects for and reinforces it.

One reason for holding this second account is that before you learn how to hit a backhand, you might accidentally get it right. If we rapidly hit tennis balls at you and give you a tennis racquet to defend yourself, you might possibly hit a backhand without ever properly learning the skill. Plausibly, this indicates that you already have the ability to hit backhands. Learning just selects for that ability, and enables you to repeat it.

In the way that I am using the terms, there is an important distinction between having an ability and having a skill. On my view, a one-off performance of a task, such as hitting a backhand, qualifies for an ability. You are able to hit a backhand; put another way, you have the capability to hit a backhand. Still, there is a difference between a one-off performance, and repeated performance. A skill is an ability that has been reinforced. It is the competence to consistently repeat an ability. The competence to consistently hit backhands, for instance, is a skill.

Now consider two different accounts of what happens when you acquire a recognitional disposition for wrens. According to one account, acquiring that disposition gives one a new perceptual ability. It enables you to attend to wrens in a new way.

What I want to suggest is that while a recognitional disposition might guide your attention when you see a wren, you could have attended in that way without it, just as you might accidentally hit a backhand without ever properly learning the skill. My positive proposal then is as follows. Acquiring a recognitional disposition does not give you a new ability to attend to wrens in a particular way. It just selects an ability that you already have and enables you to use it repeatedly. It creates a skill.

The claim that I am making is intuitive in its own right. Attending in the same way as a wren-recognizer does not entail that you have a recognitional disposition. In my terms, you can have an ability to attend in that way, but not yet have the skill and be able to repeatedly attend in that way. Someone who arrives at that attentional pattern for the first time through blind flailing has shown an ability to attend in that way, but does not yet have the skill enabling her to repeatedly attend in that way.

In *Reference and Consciousness*, John Campbell tells a similar story about the role that sortal concepts play in attention. On his view, while a sortal concept might be the cause of your attending to an object, there could have been a different cause (if someone had pointed at it, for instance, or if you had just become interested in the object spontaneously) (2002, p. 76). This is not to deny the important role that sortal concepts play in attention. For just as I argued that a recognitional disposition might guide your attention when you see a wren, Campbell argues that a sortal concept might guide your attention to single out one thing rather than another (p. 77). But also just as I argued that the recognitional disposition is dispensable, since you could have

attended in that way without it, so too Campbell argues that a sortal concept is dispensable for singling out one thing rather than another. As he puts it: “You could in principle have your attention oriented towards that object by some other cause” (p. 77).

Through the process of blind flailing, you might exhibit the ability to see wrens as a wren-recognizer would, but not yet have the skill to see wrens that way. In such a case, you would lack a recognitional disposition for wrens, but still see wrens as a wren-recognizer would. In such a case, it is not the high-level kind property that explains your new phenomenology. After all, you lack a recognitional disposition for wrens. Instead, what has happened is that you have blindly flailed into an attentional pattern that is useful for recognition. That attentional pattern allows you to see the wren in a new way, and to form a recognitional disposition for the first time. Once you have that recognitional disposition, you no longer need to blindly flail into that attentional pattern, but can cue it repeatedly and at will. You have a skill.²

² Thanks especially to Mohan Matthen, Charles Siewert, Diana Raffman, Casey O’Callaghan, Rob Goldstone, Sebastian Watzl, Adrienne Prettyman, audiences at the University of Toronto and the Pacific APA, and some extremely helpful anonymous referees.

References

- Biederman, I., & Shiffrar, M. M. (1987). "Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(4), 640.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Carrasco, M., Ling, S., and Read, S. (2004). "Attention Alters Appearance." *Nature Neuroscience* 7: 308–313.
- Connolly, K. (2013). How to test Molyneux's question empirically. *i-Perception* 4(8) 508–510.
- Connolly, K. (forthcoming). "Sensory substitution and perceptual learning." In *Sensory substitution and augmentation*, Fiona Macpherson (ed.). Proceedings of the British Academy: Oxford University Press.
- Gibson, E. J. (1963) "Perceptual learning." *Annu. Rev. Psychol.* 14, 29–56.
- Goldstone, R. L. (1998). "Perceptual learning." *Annu. Rev. Psychol.* 49, 585–612.
- Goldstone, R. L. (2010). Forward to Gauthier, I., Tarr, M. & Bub, D., *Perceptual expertise: bridging brain and behavior*. Oxford New York: Oxford University Press.
- Goldstone R. L., Landy, D. H., and Brunel, L. C. (2011). "Improving perception to make distant connections closer." *Front. Psychology* 2:385. doi: 10.3389/fpsyg.2011.00385.
- Hatfield, G. (2009). "Perception as Unconscious Inference. In *Perception and cognition: essays in the philosophy of psychology*. Oxford Oxford New York: Clarendon Press Oxford University Press.
- Price, R. (2009). Aspect-Switching and Visual Phenomenal Character. *The Philosophical Quarterly*, 59(236), 508-518.
- Prinz, J. (2010). "When is Perception Conscious?" In B. Nanay (Ed.), *Perceiving the World*. Oxford: Oxford University Press.

- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and brain sciences*, 22(3), 341-365.
- Siegel, S. (2006). "Which Properties are Represented in Perception?" In T. Szabó Gendler and J. Hawthorne (Eds.), *Perceptual Experience*. Oxford: Oxford University Press.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.
- Smith, L. B., and Thelen, E. (1993). *A Dynamic Systems Approach to Development: Applications*. Cambridge, MA: MIT Press.