Expert Knowledge by Perception

MADELEINE RANSOM

Abstract
Does the scope of beliefs that people can form on the basis of perception remain fixed, or can it be amplified with learning? The answer to this question is important for our understanding of why and when we ought to trust experts, and also for assessing the plausibility of epistemic foundationalism. The empirical study of perceptual expertise suggests that experts can indeed enrich their perceptual experiences through learning. Yet this does not settle the epistemic status of their beliefs. One might hold that the background knowledge of experts is the cause of their enriched perceptual experience – what is known as cognitive permeation – and so their subsequent beliefs are only mediately justified because they are epistemically dependent on this background knowledge. I argue against this view. Perceptual expertise is not the result of cognitive permeation but is rather the result of perceptual learning, and perceptual learning does not involve cognition in a way that entails cognitive permeation. Perceptual expertise thus provides a means of widening the scope of the immediately justified beliefs that experts can form.

1. Introduction

We rely on experts to guide us to the best restaurants and exhibits, provide us with guidelines to separate the edible from the poisonous plants, keep our populations safe from external military threats, and run quality control on the tools and devices we use, amongst other things. While not all expertise is perceptual, the study of what is known in psychology as 'perceptual expertise' suggests that some of it is. Perceptual expertise is the ability to categorize or recognize objects reliably, automatically, and exceedingly rapidly. Almost all of us are experts at recognizing faces, birds, dogs, and cars, and some of us develop expertise in categorizing these objects at the 'subordinate' level, where these are more fine-grained subcategories such as 'Northern thrush,' 'Boston terrier,' and 'Lada'.

Perceptual experts are able to form beliefs on the basis of perception – henceforth, 'post-perceptual beliefs' – that novices cannot form without testimony or inference. Suppose a birding novice is out for a walk in the forest with his friend the expert ornithologist. The ornithologist recognizes a bird in a tree and points it out – 'there, a yellow warbler', she exclaims. The novice looks and

doi:10.1017/S0031819120000157 © The Royal Institute of Philosophy, 2020
First published online 29 May 2020

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

recognizes a yellow bird, but he would not be able to form the belief that it is a yellow warbler without his friend's assistance and expert testimony. Likewise, a radiologist can readily point out tumours in X-rays, whereas the novice must rely on the radiologist's testimony to form a belief that some shadowy patch is a tumour.

In order to advance our understanding of the nature of perceptual expertise, it is crucial to understand both what the expert's ability consists in and the justificational structure of their beliefs. On one understanding of expert ability, a change in perceptual experience between novice and expert is the source of the expert's ability to form beliefs such as that a bird is a yellow warbler, or that a shadow in an image is a tumour. There are two ways of elaborating this view. On the first, the change in perceptual experience between novice and expert is entirely understood in terms of low-level properties, where these are confined to the most basic building blocks needed to construct perceptual experience, such as colours, illumination, shapes, motion, and spatial properties (Tye 1995; Dretske 1981; Connolly 2019). On the second – sometimes labeled the *rich content thesis* – the change between novice and expert is understood in terms of high-level properties, where high-level properties are simply all those properties that fall outside the category of low-level properties (Siegel 2006; Siewert 1998). In our examples, the perceptual expert represents the high-level property 'yellow warbler' or 'tumour' in her perceptual experience, whereas the novice does not. In what follows I assume the rich content thesis to be true and explore what follows for the epistemology of expert post-perceptual belief.[1] While we readily defer to perceptual experts, in order to make sense of this deference we must get clear on the structure of their justification: does it differ in kind from that which justifies novice beliefs, or does it merely differ in degree of reliability?

Perception is often thought to provide a distinctive source of justification for our beliefs, with some maintaining that post-perceptual beliefs are basic, or immediately justified (Huemer 2001; Brewer 1999; Goldman 2008). A belief is immediately justified when it does not depend on another belief for justification, but is nevertheless

[1] I provide detailed argument elsewhere (Ransom 2020) for the claim that perceptual experts can come to represent high-level properties in perceptual experience via a process that I shall discuss below: perceptual learning. This is why I do not discuss the possibility here that expert knowledge is instead the result of an unconscious or implicit inferential process (see for example (McGrath 2004)).

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

justified (Alston 1983). Mediately justified beliefs are those that do depend on one or more other beliefs for justification. In the case of post-perceptual beliefs, the source of justification is not a further belief, but rather our perceptual experience. Post-perceptual beliefs are therefore plausibly immediately justified (though see Kvanvig and Riggs 1992).

If the number and kinds of properties we represent in perceptual experience are fixed, then so too is the scope of immediate perceptual justification. However, if our perceptual experience can change and grow with learning, then it seems that the scope of immediate perceptual justification can also change. Call this the *wider scope* thesis: the scope of immediately justified perceptual beliefs can be enlarged or broadened through learning. This change in scope would have particular significance for foundationalism, the view that there is a class of immediately justified beliefs that justify all mediately justified beliefs (for example Bonjour 1978). The truth of the wider scope thesis would render foundationalism more plausible because an expandable base of immediately justified beliefs could in turn ground more mediately justified beliefs. Beyond foundationalism, any view that allows that perception provides us with immediately justified beliefs should care whether the wider scope thesis is true. If it were not true, then we would need to create and be mindful of a further distinction between perceptual experiences that provide immediate justification and those that do not.

At first glance, perceptual expertise supports the wider scope thesis: perceptual experts come to have richer perceptual experiences, and since perceptual experience is a source of immediate justification, then their resulting beliefs are also immediately justified. In fact, some foundationalists do appeal to learning as a way of increasing the number of basic beliefs one has (Goldman 2008; Brewer 1999). Likewise, rich content theorists tend to hold that many high-level properties come to be represented in perception via some form of learning (Macpherson 2012; Siegel 2010; 2006). While our perceptual systems may allow us to represent certain low-level properties like colours and shapes without any learning – or even limited high-level properties such as causation – the idea that the visual system represents high-level properties like 'yellow warbler' from the get go is less plausible. Some of us recognize these objects, some of us do not, and learning explains the difference.

In both debates, however, the nature of the learning process must be further specified so that the epistemic consequences can be better understood. Understanding what sort of learning is at issue is particularly important because the wider scope thesis and the rich

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

content thesis can come apart – we might represent high-level properties in perceptual experience without its being the case that this provides a new source of immediate justification.

One way in which perceptual expertise may fail to vindicate the wider scope thesis is if expert perceptual experience is the result of cognitive permeation. This would render the justification of expert post-perceptual beliefs mediate rather than immediate. In such cases, 'perceptual' expertise would be a matter of acquiring the right background theory, and our evaluation of expertise would consist largely in evaluating the merits of the learned theory. If this were true, then it would have profound implications for our understanding of perceptual expertise, under what conditions we ought to trust perceptual experts, and how we design training programs in disciplines from medicine, to the military, to art criticism. While perceptual experts would still have reliable post-perceptual beliefs, the structure of their justification for such beliefs would be radically different from the post-perceptual beliefs one can form on the basis of unlearned perceptual experience.

In what follows, I argue that perceptual expertise is not the result of diachronic cognitive permeation but is rather the result of perceptual learning, and perceptual learning does not involve cognition in a way that entails cognitive permeation. Perceptual expertise thus provides a means of widening the scope of the immediately justified beliefs that experts can form. This provides us with an added understanding of why we ought to trust expert testimony in many domains: not only is it based on perceptual experience, it is also epistemically independent of background theory. In section two I present the cognitive permeation thesis in more detail, and elaborate what I take to be the most plausible hypothesis for how it might result in perceptual expertise. This advances the debate by filling in the details of a process that has remained largely underspecified in the literature. In section three, I discuss the epistemic significance of cognitive permeation, arguing that post-perceptual beliefs formed via permeated perceptual experiences are only mediately – rather than immediately – justified. In section four, I introduce an alternate conception of how perceptual expertise is acquired: through perceptual learning. I then argue that perceptual learning can accommodate cases of perceptual expertise previously thought to involve cognitive permeation. This paper thus also provides a challenge to defenders of cognitive permeation: they must find cases of perceptual expertise that cannot be explained away in terms of perceptual learning along the lines I have proposed.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

2. Perceptual expertise as cognitive permeation

Cognitive permeation – also commonly referred to as cognitive penetration – is hypothesized to occur when the contents of perceptual experience are altered in some way by one's cognitive states, such as beliefs and desires (Pylyshyn 1999; Siegel 2012).[2] Permeating beliefs are sometimes called 'background beliefs'. Susanna Siegel offers the example of Jill, who has a background belief that Jack is angry with her, and on this basis has the perceptual experience of Jack's face as angry.(2012) This experience in turn causes her to form the post-perceptual belief that Jack is angry with her, or reinforces her existing belief.

The basic idea of cognitive permeation is easy enough to grasp. Providing an adequate characterization of the phenomenon, however, has proven more challenging (Stokes 2013; Zeimbekis and Raftopoulos 2015, pp. 27–32). Nevertheless, there are common strands amongst several definitions. First, there is a *causal* condition: the background belief must be a cause of the resulting perceptual experience. Second, there is an *internal* condition: in order for cognitive permeation to occur, the background belief must provide direct input to perceptual processing, which then modulates perceptual experience. Third, there is a *semantic coherence* condition: the content of the input provided by the background belief must bear a semantically non-arbitrary relationship to the resulting content of the perceptual experience. As Zenon Pylyshyn puts it, the influence of the cognitive input on the contents of perception must be 'coherent or quasi-rational when the meaning of the representation is taken into account' (Pylyshyn 1999, p. 365 n. 3). Together, these conditions suffice for cognitive permeation, and rule out common cases of benign cognitive influence on perception.[3] For example, suppose you have a belief that a friend is coming to the department to visit, and so you attend to the open door of your office, expecting to see her any minute. When your friend finally appears, your belief has causally altered the contents of your perceptual experience by

[2] Here I follow Becko Copenhaver (personal conversation) in using 'cognitive permeation' to avoid gendering reason as masculine.

[3] While some theorists have rejected the semantic coherence condition in order to make room for the permeation of experience by desires and moods, the sort of cognitive permeation hypothesized to occur with expertise must be explicable in terms of beliefs or concepts closely tied to the resulting perceptual content, and so would meet the semantic coherence condition.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

directing your attention to the door. This alteration is semantically coherent because the content of the attention-guiding belief – that your friend will appear in the doorway – is semantically connected to the relevant content of your perceptual experience: your friend standing in the doorway. However, the internal condition has not been satisfied because cognition has provided only indirect input to perceptual processing, through the reorientation of your attention. For this reason, the case fails to count as cognitive permeation.

At first glance, the route to perceptual expertise seems quite different from cognitive permeation – psychological studies have found that the change in perceptual abilities from novice to expert requires long periods of training involving many perceptual exemplars, not the simple acquisition of background beliefs or knowledge (Gauthier et al. 1998). No matter how firmly a novice might believe based on the testimony of an expert that the bird pictured is a yellow warbler, this does not alter her ability to recognize yellow warblers in the short term – acquiring such knowledge is not sufficient for perceptual expertise.

Nevertheless, some have argued that perceptual expertise should count as a form of cognitive permeation – *diachronic* cognitive permeation – where this signifies that the permeation of perception by cognition happens over time, with cognition gradually reshaping the structure of the perceptual system and only eventually leading to a change in perceptual experience (Churchland 1979 ch. 2; Kuhn 1962 ch. V, X; Churchland 1988, p.179). This is unlike the case of *synchronic* cognitive permeation described above, where Jill's belief that Jack is angry is able to permeate experience without delay. Recent defenders of cognitive permeation have discussed cases of perceptual expertise such as the radiologist's ability to detect tumours, the art connoisseur's ability to recognize categories of artworks, and the herpetologist's ability to recognize copperhead snakes, assuming these to be cases of diachronic cognitive permeation (Stokes 2014; Siegel 2012; Stokes 2020; Lyons 2011).

If achieving perceptual expertise involves a form of cognitive permeation, then it must be the case that it is the learning of the relevant theory, or some subset of these beliefs, that is responsible for altering one's perception (see also Fodor 1988, p. 195). Given that we are interested here in high-level properties that correspond to categories, the background beliefs will concern the nature of these categories. Novices may begin by acquiring background beliefs roughly of the form 'objects of category C possess properties x, y, z' where x, y, z are low-level perceptual properties such as size, colour, and shape, and C corresponds to a high-level property such as being a yellow

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

warbler. These beliefs must then cause the change in perceptual experience over time. Again, as with synchronic cognitive permeation, the content of the background beliefs must also bear a semantically non-arbitrary relationship to the resulting content of the perceptual experience.

The condition that there must be an internal connection between the belief and the perceptual experience is more difficult to adapt for the diachronic case. It is hard to see how the causal link between perception and cognition could be direct and internal and yet cognition only exert its influence on perceptual processing – and perceptual experience – over time. That is, if perceptual experience is not synchronically cognitively permeable by the learned background beliefs then it remains to be explained how it could ever come to be diachronically cognitively permeable. A gradual effect, if it is to be internal, still requires a direct pathway for providing input.

One mechanism by which diachronic cognitive permeation may occur involves appealing to attention. This proposal should be confusing to readers at first. The discussion above rules out attention, as it serves only as an indirect modulator of perceptual experience and so violates the internal condition. Another reason for ruling it out is on the basis of the semantic coherence condition because attention ostensibly determines only the input to perceptual experience and not the resulting content (Pylyshyn 1999). In the example discussed above of awaiting your friend in your office, semantic coherence happens to be satisfied. However this satisfaction is 'lucky' in the sense that it would not still obtain in the case that another colleague had appeared at the door instead. In such a case, one's perceptual experience would be of this other colleague, not the visiting friend. Nevertheless, several theorists have recently argued that some varieties of attention should be understood as a source of cognitive permeation (Mole 2015; Wu 2017; Stokes 2018).

Explaining how attention might allow for diachronic cognitive permeation requires making some more fine-grained distinctions between different varieties of attentional phenomena (Carrasco 2011; Armstrong 2011). Our attention can be oriented on the basis of spatial location, features, or objects. Spatial attention is attention to a location, irrespective of what is there. Object and feature-based attention are attention to objects or features, irrespective of their location. Exogenous attention is a 'capturing' of our attention by certain features, irrespective of our aims: we involuntarily turn to look at sudden loud bangs and bright flashes. Endogenous attention is driven by cognition, where this is usually understood in terms of our goals and tasks: we attend only to people wearing red shirts in a

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

crowd because we are looking for a friend wearing a red shirt. Endogenous attention can also occur more automatically, though it is still driven by what we have learned: we reorient our attention effortlessly on the basis of arrows or the finger of a friend pointing in a given direction, where this reorientation can also conflict with our current task. Overt attention occurs when there is an attentional shift that involves a reorientation of one's body, where this includes eye movements. Covert attention occurs when there is a shift in attention without any such bodily change, such as when we tune out the person we are conversing with in order to eavesdrop on the conversation at the table next to us. All of these distinctions are orthogonal to one another. For example, endogenous attention may be covert or overt, and to features or to spatial location.

The sort of attention ruled out from counting as cognitive permeation in the example above is overt, endogenous, spatial attention. In awaiting your friend in your office, you purposefully reorient your head or your eyes to a spatial location in response to a belief and desire to see your friend. However, covert, endogenous, feature, or object-based attention is not so easy to rule out. This is because in certain cases such attentional phenomena appear to satisfy the three conditions that suffice for cognitive permeation. Endogenous attention allocated on the basis of features or objects will always satisfy the semantic coherence condition: the feature or object that is guiding one's attention will also be the feature or object that one comes to represent in perceptual experience. When such attention is also covert, then the internal condition is also satisfied, as attentional allocation cannot be explained by assigning attention the mere orienting role of providing input to perceptual processing. There is some empirical evidence that supports the claim that covert feature or object-based attention can indeed bias perceptual processing itself (Mole 2015).

These distinctions provide the detail needed to elaborate an account of how diachronic cognitive permeation occurs – call this the *belief-driven hypothesis*. First, the novice acquires the relevant background beliefs: she learns via testimony or inference that a given object possesses characteristic features x, y, z ('robins have red breasts, brown feathers, and a yellow ring on their beaks').[4]

[4] Some might object here that what the novice acquires is a concept rather than a belief. The challenge for those who nevertheless endorse cognitive permeation is then to articulate how theory or belief can lead to concept acquisition in a way that does not appeal to the pattern-driven hypothesis that I defend below.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Second, on the basis of these learned background beliefs, the person overtly and endogenously directs her attention to these characteristic features one by one to make an identification. This is in part a post-perceptual process, likely effortful at first, whereby the person accumulates perceptual evidence for categorization. It is a process somewhat akin to checking boxes on a list ('Does the bird have a yellow ring on its beak? Check'). Third, over time repeated attention to these features causes the restructuring of the perceptual system: the pattern of attentional tuning becomes automated and covert, no longer requiring the deployment of endogenous or overt attention – it directly privileges the grouping of these features and renders them simultaneously salient, perhaps by biasing the competition for perceptual processing resources (Desimone 1998). This automatic perceptual grouping may even be said to correspond to the representation of high-level properties in perceptual experience – the relevant low-level properties have been grouped in a way that reliably tracks the relevant natural or artifactual kinds.[5] With this theory of how perceptual expertise occurs, we may now turn to the epistemic status of expert post-perceptual beliefs.

3. The epistemic significance of cognitive permeation

In some cases of cognitive permeation, the resulting perceptual experience seems to be unable to provide any justification at all to the post-perceptual belief. Siegel's case of Jill and Jack is arguably of this variety: given that Jill's perceptual experience of Jack's anger is caused entirely by her background belief that Jack is angry, it seems illicit for the perceptual experience to provide any additional justification for the post-perceptual belief that Jack is angry. Siegel's original analysis of why this is the case rests on the notion of perceptual sensitivity to the incoming stimulus (2012, p. 4, p. 12).[6] In such cases, the relevant content of the incoming stimulus – Jack's actual facial expression – makes no difference at all to Jill's perceptual experience. Instead, her experience is entirely dependent on her belief, in the sense that it – rather than the incoming perceptual

[5] See also McGrath (2017) for discussion of how recognitional abilities may be developed over time.

[6] In her (2012) paper Siegel also provides an analysis in terms of epistemic circularity, and she has since (Siegel 2017) provided an alternate analysis of hijacked experience in terms of poor inference. I adopt the sensitivity analysis because it is particularly fitting for the case of perceptual learning.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

information – determines her perceptual experience. Her perceptual experience would be of Jack as angry, regardless of the actual perceptual input (see also Lyons 2011, p.301). Siegel holds that due to this insensitivity the resulting perceptual experience fails to provide any independent justification at all for the ensuing belief. In this case, while perceptual experience can come to have rich contents as a result of learning, the scope of immediately justified post-perceptual beliefs is not wider as a result.

However, this analysis is compatible with the possibility that there are cases of cognitive permeation that do allow for some sensitivity to the incoming stimulus – it is meant to explain why certain instances of cognitive permeation are epistemically problematic, not to provide a comprehensive diagnosis of all cases of cognitive permeation. Indeed, Siegel and others assume that in some instances cognitive permeation will be virtuous – such as the radiologist who can perceptually recognize tumours, and the herpetologist who can identify copperheads (Lyons 2011; Stokes 2020).

This assumption has received a limited defence. Jack Lyons (2011) proposes that cognitively permeated perceptual experience is virtuous when it increases reliability. Dustin Stokes emphasizes the superior performance of perceptual experts in their domain of expertise, arguing that it is an epistemic good because 'the expert is moving closer to an optimal cognitive stance on the world (or a part of it), where she can better acquire behaviourally relevant category and diagnostic information' (2020, p.18).

We can recast a defence in terms of sensitivity. It is plausible that the background beliefs of perceptual experts – understood as per the belief-driven hypothesis – do help increase sensitivity to the incoming stimulus rather than decrease it. For example, suppose my belief that dogs have tails, fur, and so on, diachronically permeates my perceptual experience in order to allow me to identify something as a dog. This belief may allow me to become more sensitive to certain aspects of my experience, as I may preferentially attend to the animal's tail and fur, say. And of course it allows me to become sensitive to the property of something's being a dog. So long as the background belief does not cause me to perceptually categorize things as dogs irrespective of whether they are dogs, then it would seem that the belief is working in concert with perceptual experience rather than 'hijacking' the experience or reducing my sensitivity to perceptual input. If perceptual expertise results from diachronic cognitive permeation, it is plausible that perceptual experience still provides justification for the resulting post-perceptual beliefs.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Nevertheless, there remains a challenge to the wider scope thesis on this view: perceptual experiences formed as the result of cognitive permeation provide only mediate justification for the ensuing beliefs.[7] One reason for thinking this is because it may be reasonable to endorse what I will call the *dependence* principle. Nicholas Silins (2013, p.25) articulates it thusly:

> When an experience has the content that p in part in virtue of some belief one has, it gives one justification to believe that p only in virtue of one's having justification for the belief.

Silins motivates the dependence principle by providing the example of Alexa, who forms an unjustified background belief that 'things with such and such a look are sheep dogs', which then allows her to represent the property 'sheep dog' in perceptual experience.[8] He appeals to the intuition that it would be illicit for Alexa to gain justification through her perceptual experience for the belief 'that is a sheep dog' given that she was not justified in the background belief that led to her experience in the first place. Those who share this intuition should endorse the dependence principle to explain why expert perceptual experience fails to result in an immediately justified belief. If we allow that unjustified background beliefs prevent subsequent post-perceptual beliefs from being justified at all, then this would seem to indicate that the background beliefs are playing a mediating role even in cases involving justified background beliefs.

Another consideration in favour of the dependence principle stems from the semantic coherence condition for cognitive permeation, which ensures that there is a non-arbitrary semantic relation between the background belief and the contents of perceptual experience. Since the content of the background belief is connected in a quasi-logical manner to the content of the perceptual experience, rather than a brute causal relation, this renders it plausible to also attribute a relationship of epistemic dependence; it is the content itself that is doing the work of restructuring the perceptual system over time. Those who want to resist the dependence claim bear the

[7] For alternate ways of motivating the claim that learned perception provides only mediate justification see (Vaassen 2016; M. McGrath 2017, pp. 30–34). Though see Siegel (2017, pp. 47–50), who considers various ways in which perceptual experiences may provide immediate justification even if they are cognitively permeated.

[8] For further discussion or endorsement of the dependence principle see (McDowell 1982, p.478; Plantinga 1993, pp.99–101; Markie 2005, p.350).

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

burden of providing some positive reason to believe that etiology does not matter in the case of perceptual expertise, or does not impact immediate justification.

In summary, if it is true that expert perceptual experiences are the result of cognitive permeation then this plausibly renders them only mediately justified. Therefore, perceptual expertise does not vindicate the wider scope thesis: experts do not broaden the scope of their basic perceptual beliefs through learning. In the following section I argue that perceptual expertise is not the result of diachronic cognitive permeation. It is rather the result of perceptual learning, which does not involve cognition in a way that entails cognitive permeation.

4. Perceptual expertise and perceptual learning

In the discussion of expert perceptual experience, people have often assumed cognitive permeation as the default option for explaining how learning can alter the contents of perceptual experience (Churchland 1988; Lyons 2011; Briscoe 2015; Macpherson 2012; Brogaard and Chomanski 2015). It is a tempting option because declarative learning – the learning of verbally conveyed facts – and belief-formation so often occur before, or coincide with, a change in perceptual experience, and so it seems natural to assume that the change depends on these newly acquired beliefs.

However, this inferential leap is unwarranted. The direction of dependence may well run the other way round: changes to perceptual experience may themselves anchor and solidify the learning of new beliefs. Moreover, these changes may depend on repeated exposure to and practice with the perceptual stimuli, rather than declarative learning. That is, what is known as perceptual learning may instead be driving the changes in perceptual experience.

Perceptual learning has emerged as an alternate way of explaining the sort of learning that takes place from novice to expert (Ransom 2020; Connolly 2019; Brogaard and Gatzia 2018; Chomanski and Chudnoff 2018; Chudnoff 2018; Vaassen 2016). It is an enduring functional change in the perceptual system that results in a long-term change in the contents of our perceptual experience (Goldstone 1998; Gibson and Gibson 1955). However, unlike cognitive permeation, the cause of the change is practice with or repeated exposure to a perceptual stimulus. While diachronic cognitive permeation may also involve such practice and exposure, background belief

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

is the main cause of the change. This is not so in the case of perceptual learning, as I shall specify below.

There is considerable evidence that with practice our perceptual experience can become more fine-grained and accurate in terms of low-level properties (Ball and Sekuler 1982). However, perceptual learning has also been implicated in the acquisition of the sorts of high-level properties of interest here. While the psychological study of perceptual expertise (understood as only expertise with subordinate categories) has morphed into its own relatively independent subfield, much of the study of learned categorization more generally has remained within the domain of perceptual learning (Goldstone, Landy, and Brunel 2011).

Perceptual learning involves processes known as attentional weighting and stimulus imprinting (Goldstone and Byrge 2015). In attentional weighting, those features relevant to a given category are weighted in attention more heavily over time, and those that are irrelevant are weighted less heavily. The weights – I argue elsewhere (Ransom 2020) – should be understood as evidential weights, where this means that their weighting corresponds to how diagnostic a given feature is for category membership. For example, eyes and face shape are highly diagnostic of category membership for barn owls, whereas the legs, colour, and beak are diagnostic of category membership for blue herons.

Stimulus imprinting involves grouping the different weighted features together, as well as becoming sensitive to and encoding the relations between these features when they are relevant for diagnosticity. For example, what matters to categorizing faces as faces is not just the detection of eyes, a nose, or a mouth, but detecting these features together, in a specific configuration. The end result is a network of feature and feature-relation detectors that are grouped together and allow for detection and categorization of the object in perception without any input from cognition.

This picture will strike readers as being very similar to the proposed explanation of diachronic cognitive permeation offered earlier – the belief-driven hypothesis. However, it is one thing to say that perceptual learning involves the attentional weighting of features over time, and another to say that this amounts to cognitive permeation. What is at issue is whether the process of attentional weighting must be guided by cognition in a way that amounts to cognitive permeation.

Against the belief-driven hypothesis I defend the *pattern-driven* hypothesis: perceptual learning is driven by real patterns in the relevant class of training stimuli, and it is the gradual learning and

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

detection of these patterns that drives the restructuring of the perceptual system, rather than declarative learning. Background beliefs are not necessary for perceptual learning to occur, and when cognition is involved in learning, it does not count as cognitive permeation because it fails to satisfy the semantic coherence condition.

Some initial support for the pattern-driven hypothesis comes from the large numbers of exemplars needed for perceptual learning to occur. In the case of perceptual expertise, learning to categorize lab-manufactured objects such as greebles – gremlin-like animals – into subordinate categories requires roughly nine hours of practice with exemplars, broken into one-hour sessions spread out over two weeks (Gauthier and Tarr 1997). In well-studied domains of expertise such as chess, musical performance, and classical composing, the time required to develop 'full blown' expertise – where this outstrips perceptual expertise – is often estimated to be roughly a decade (Hayes 1989).

This exposure to exemplars is not optional to developing perceptual expertise. It counts in favour of the claim that the primary factor in assigning attentional weights is not what background beliefs one has, or one's background theory, but rather what the statistical properties of the environment are; if a white stripe is important to categorizing something as a skunk, then given enough examples this feature will be weighted more heavily. Though this provides some support for the pattern-driven hypothesis over the belief-driven hypothesis, it does not yet rule out possibility that beliefs play an essential role in perceptual learning. The belief-driven hypothesis may also be able to explain the gradual nature of the process – while background beliefs can cause shifts to our attention immediately, the perceptual system perhaps needs time to automate and synchronize the process so that multiple features are selected simultaneously, requiring a certain number of pairings between the features. We therefore require further reasons to favour one hypothesis over the other.

On the belief-driven hypothesis, acquisition of the relevant belief(s) is not a sufficient condition, but it is a necessary condition for perceptual expertise to occur. If the belief is to be the cause of the expertise, then it cannot be absent. In (most) such cases, the following counterfactual will be true:

> If the subject had not learned the background belief(s) B but was seeing and attending to the same distal stimuli, she would not have an experience with content p.[9]

[9] I adapt this counterfactual with some alteration from (Siegel 2012, p.6).

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Attention here should be understood as spatial attention, rather than the feature-based attention involved in perceptual learning. In the latter case, we cannot hold attention fixed between novice and expert in part because synchronous attention to the features relevant for categorization is itself a component of perceptual expertise. In the case of spatial attention we can hold attention fixed. Seeing a distal stimuli here should be understood in terms of low-level properties only, given that what is at issue is whether the expert can come to perceive high-level properties. The counterfactual will not hold in all cases because there may be a closely related set of background beliefs B* that would also allow the subject to have an experience with content p. Nevertheless, the counterfactual is useful because it makes clear that if the attentional weighting involved in perceptual expertise – posited by both the pattern-driven and the belief-driven hypotheses – can occur without being driven by the relevant background belief(s) then it is not an instance of cognitive permeation.

Attentional weighting absent background belief can indeed occur. Goldstone, Landy, and Brunel (2011) characterize this sort of process as 'blind flailing': attentional weights to features are assigned randomly, and those that allow the agent to make important or useful discriminations are selected and made permanent. This process is much like natural selection, where random variation plus a mechanism of selection can lead to strategic adaptations, without anyone having the intent of designing a creature with such an adaptation. The process is 'blind' in that there is no initial background belief on the part of the agent that guides attention.

This sort of process has been well studied in low-level perceptual learning. Such learning can occur absent cognitive guidance, in the sense that the learning occurs despite it's being task-irrelevant – and so is not guided by endogenous attention – or even when the stimuli are not consciously perceived (Watanabe, Náñez, and Sasaki 2001). In both instances this rules out a role for endogenous attention. In the first case, the subject is performing an unrelated and attentionally-demanding central task, and so is not attending to the stimuli that are nevertheless learned. In the second case, the subject does not consciously perceive the stimuli, and so is not plausibly able to direct her attention to them. However, this phenomenon has been less well studied for the perceptual learning involved in perceptual expertise, as category learning is usually part of the assigned lab task and so should guide our attention.

Nevertheless, there is some evidence that perceptual expertise can be acquired despite its being irrelevant to a training task (Wong, Folstein, and Gauthier 2011). In the study in question, subjects

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

trained to perform an orientation detection task with a class of novel objects called 'ziggerins' – where the subtle variation in the shape of the individual ziggerins was irrelevant – were able to increase their performance in a post-training test that measured shape discrimination performance. This is a common marker of perceptual expertise with subordinate categories: the ability to make fine discriminations between similarly shaped individuals of the same category. The significant improvement of subjects in the post-test is therefore taken as evidence for task-irrelevant perceptual expertise.

In addition, the sort of background belief that is posited by the belief-driven hypothesis is not present in this case. Experimenters did not provide any prior instruction to subjects as to which intrinsic features of the ziggerins were relevant to distinguishing between them, as only rotation was relevant to completing the task. So in this case no background beliefs or endogenous attention to category-distinguishing features are necessary for perceptual expertise to develop. The causal condition for cognitive permeation is violated.

While one counterexample is sufficient to falsify the claim that background beliefs are necessary for perceptual expertise to develop, we may wonder how often blind flailing actually occurs in practice. Much perceptual learning for expertise takes place with some sort of cognitive guidance, or the adoption of a strategy on the learner's part. In light of this, then perhaps we ought to make distinctions between different types of perceptual expertise, and evaluate the epistemic consequences separately, given their different etiologies. Perhaps background belief is necessary for some varieties of perceptual expertise.

Indeed, there are studies that suggest that cognitive guidance is required. For example, Stokes (2020) cites a study on perceptual expertise with birds as demonstrating that mere exposure to perceptual stimuli is insufficient to produce perceptual expertise. In the study, experimenters trained twenty-one bird novices who performed an initial subordinate-level matching task prior to training (Tanaka, Curran, and Sheinberg 2005). Ten subjects then received training in categorizing owls at the subordinate level ('Eastern screech owl') and wading birds at the basic level, where basic-level categories are composed of several subordinate-level categories, and so are less fine-grained and more inclusive ('wading bird'). For the group of eleven, the training was reversed; they learned to categorize wading birds at the subordinate level ('blue heron') and owls at the basic level ('owl'). After this training period, the subjects performed the same pre-training matching task, along with two new matching tasks of the same kind as the first but with different images, either

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

of novel instances of the same subordinate bird categories used in training, or novel instances of new subordinate categories.

In these post-tests, both groups did better at categorizing birds into the subordinate levels they had been trained in. However, what is of interest is whether they did better at making subordinate-level categorizations for the category on which they did not receive subordinate-level training, but which they were nevertheless exposed to during training. Here, subjects did make some significant improvement in the first post-test, though it was below the improvements seen for those who received subordinate-level training. In the second and third post-tests, there is no pre-training baseline to compare the results to – only the results of the other group. In both these post-tests, those untrained in the subordinate categories did worse than those with training, though the difference between the groups was smaller than in the case of the first pre-test.

In deciding whether or not this case counts against the pattern-driven hypothesis, we must be careful to properly characterize the nature of the cognitive influence. Not all guidance of perceptual learning by cognition will count as cognitive permeation. Goldstone, Landy, and Brunel (2011) call the strategic manipulation of the perceptual learning process 'myopic flailing' to characterize a process that is more efficient than random variation but that falls short of cognitive permeation. They provide a list of ways in which we can 'hack' or tune our perceptual systems to accelerate perceptual learning. For example, we might put two images of an almost-identical viceroy and a monarch next to each other in order to compare them and spot the differences. Or, in wine tasting we might move the wine around our mouths to get a fuller and longer taste, and suck in some air to enhance its olfactory profile.

Such strategies serve merely as catalysts – rather than cognitive permeation – ensuring that our perceptual equipment gets the right input, or enough practice with the right exemplars. Catalysts serve to speed a process that would have occurred anyways given enough time and the right conditions, so in such cases the guiding beliefs are not necessary for learning. It is not the case that if the subject had not learned background belief(s) B then she would not have perceptual experience with content p. In such cases, the subject would still come to have perceptual experience with content p, albeit after a longer time period.

That the guiding beliefs in this sort of myopic flailing do not count as cognitive permeation can be seen by noting that they do not observe the semantic coherence condition. We may adopt a given strategy in the belief that it will be helpful in learning to perceptually

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

distinguish a given category of objects without possessing a background belief that the category is distinguished in virtue of certain characteristics, or even a belief as to which category is which. In such cases, there is a semantically arbitrary relationship between the belief and the resulting perceptual content. For example, I may put a picture of a viceroy and monarch side by side to facilitate perceptual learning without knowing in advance what the perceptual difference is between them, or which is the viceroy and which the monarch – the training sets might be simply labeled 'x' and 'y'. It happens to be the case that they are distinguishable partly in virtue of a black stripe that cuts the bottom of only the viceroy's wings. But using the same strategy I might have learned to distinguish them in another way if they possessed different distinguishing features.

The case of bird expertise discussed above may be explained in this way. First, the results are suggestive of the catalyst interpretation. While greater post-test gains were seen for the subordinate categories that subjects had been trained in (subjects trained to distinguish various wading birds were better at performing the same/different task for various wading birds), there was nevertheless significant improvement for the subordinate categories that the subjects had only received basic-level training for. Subjects trained to distinguish various wading birds were better at performing the same/different task for various owl species that they had been previously exposed to during their training (Tanaka, Curran, and Sheinberg 2005, pp.148–49). This improvement suggests that there is some level of perceptual expertise that can be acquired with passive exposure, or task-irrelevance. While the degree of expertise is comparatively lower than that of those who received subordinate-level training, it may nevertheless be the case that subjects would gain a higher degree of perceptual expertise with more exposure. This interpretation is also plausible with respect to the way people often form perceptual expertise in real life. When I first moved to Vancouver, I was not able to recognize Northern flickers (a kind of woodpecker) on sight, but I soon learned to do so as they are abundant in my neighbourhood. I was able to reliably recognize the bird long before I looked up the name or read a description of their identifying characteristics.

Second, the sort of training that was given to subjects does not satisfy the semantic coherence condition. Subjects were not given explicit instruction on the defining or typical perceptual characteristics of different basic-level or subordinate categories. Instead, they were given mere semantic labels, and feedback in terms of correct or

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

incorrect semantic labeling during the training process. Subjects had to learn on their own, in the face of the many exemplars of birds, which features were relevant for categorization.

Here one might be tempted to characterize the relevant background belief along the lines of 'that is a blue heron' and so claim that it does meet the requirement. After all, it is a belief that is virtually identical to the resulting content of perceptual experience. However, the belief concerns a mere semantic label. The person might have learned 'that is a blue-footed booby' instead, and still come to develop a perceptual category for blue herons on this basis. No content concerning the relevant perceptual characteristics of blue herons is conveyed in the belief involving the semantic label (except of course in this case that the heron is blue – though really it is more grey than blue).

A more difficult case for the pattern-driven hypothesis is that of X-ray technicians, where passive exposure does not seem to result in any perceptual expertise (Stokes 2020). They are held in contrast with radiologists, who are physicians trained to interpret medical images. Radiological expertise is often thought to be partly perceptual in nature, due to the speed and accuracy with which radiologists are able to make diagnoses from a single glance (Bilalic et al. 2016). Unlike radiologists, X-ray technicians are not physicians and are not typically trained in how to diagnose medical images. Instead, their training consists in how to properly use the machines in order to get reliable images. They therefore plausibly lack the background beliefs that would provide attentional guidance posited by the belief-driven hypothesis in order to form the perceptual category 'tumour' or 'nodule', amongst other diagnostic categories. If they indeed lack perceptual expertise with these categories after a long period of exposure to images that contain them, then it looks like vindication that at least some forms of perceptual expertise require cognitive permeation.

A recent study provides some support for this (Manning et al. 2006). The experimenters did not find a statistically significant difference between the performance of eight novices at detecting pulmonary nodules and that of five experienced X-ray technicians without training in chest image interpretation. This is in comparison to the superior performance of the same five X-ray technicians after completing six months of training in chest image interpretation, whose accuracy was then on par with the performance of eight experienced radiologists.

However, cases such as these can be accommodated with a fuller understanding of the role of goal-directedness in perceptual learning.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Just as with the indeterminacy of translation there are many candidate interpretations of a foreign phrase or word, so too with perceptual learning there are many candidate patterns available for learning in a given set of images. In image sets as complex as X-rays, there are many features that remain invariant across the set. When we select one pattern for learning – perhaps through myopic flailing – this may be at the expense of another pattern because we may not be able to synchronously direct our attention to all features relevant to both patterns. So when myopic flailing is initiated or guided by the goal or the desire of the novice to improve in detecting a certain pattern, this may diminish or eliminate her ability to learn another pattern.

In the case of radiologists and X-ray technicians, it is likely that their divergent goals have led them to focus on developing perceptual expertise with different categories by selecting different patterns. Whereas the goal of radiologists is to detect and diagnose potential abnormalities such as lesions and tumours, the goal of X-ray technicians is to make sure the X-rays are good quality medical images. Considerable training is required in order to discern whether an image is of good quality. While there is no exact list of characteristics for what makes a quality image, the image properties X-ray technicians are attuned to include image sharpness, density (degree of blackening on the film), and contrast (difference in density between two adjacent structures) (Easton 2009). So X-ray technicians are likely perceptual experts with respect to whether a given X-ray is of good quality, though no studies have been done to test for such perceptual expertise. If it is the case that X-ray technicians do have such perceptual expertise, then it would go towards explaining their lack of expertise with respect to tumours and other abnormalities – during their exposure to images over the course of their work, they are tuned to different patterns.[10]

Finally, the most difficult cases of perceptual learning to accommodate without cognitive permeation will be those where instruction does take the form of providing an explicitly articulated rule along the lines of 'category C is characterized by perceptual features x, y,

[10] A further source of support for understanding human expertise with radiological images as perceptual may be found in machine learning, where neural networks are becoming increasingly adept at categorizing images as containing benign or malignant tumours only on the basis of having received training with labeled images (Shen et al. 2019). While this might be proof of concept, a full discussion of this line of support would take us far afield of the main task here, as there are several challenges to this line of argument.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

z'. However, in such cases while the semantic coherence condition is satisfied, it is only luckily so. While in normal conditions semantic coherence is a good way of determining that the causal condition is satisfied, in these cases, I argue that there is only lucky semantic coherence: the background belief and the perceptual experience are semantically coherent only because the belief happens to be true. This lucky coincidence of belief and perceptual experience does not establish that it is the belief that is doing the causal work of reshaping experience.

To illustrate, our previous counterfactual analysis will not help us here. Instead, we need to consider what would happen in cases where the background beliefs are false. For example, suppose I erroneously believe that the difference between viceroys and monarchs is the specific shade of orange – monarchs are dark orange, and viceroys are slightly lighter. This background belief causes me to spend a lot of time focusing on attending to the colour of these two categories of butterflies. In fact, colour is not a reliable way of telling the two apart. In the absence of any actual studies of perceptual expertise and false background beliefs, there are three distinct hypothetical scenarios to consider.

The first is that I come to develop perceptual expertise, and am able to rapidly, automatically, and reliably identify and distinguish between viceroys and monarchs. This does not support the belief-driven hypothesis, but is rather consistent with the pattern-driven hypothesis. It is implausible to causally attribute my ability to the background belief in this case. My perceptual expertise is achieved in spite of, rather than due to, the background belief. Rather, what is likely happening is that with enough exposure to reliable training sets of monarchs and viceroys I am able to begin to be able to pick out and attentionally weight the characteristics that are actually distinctive of these two categories. That is, I engage in perceptual learning. In this scenario, the semantic content of the false background belief and the perceptual experience are not coherent, and if they do become so it is in virtue of the acquired perceptual expertise driving a change in belief. This is one sense in which the semantic coherence condition may be 'luckily' satisfied.

Moreover, given that the background belief is causally inert in bringing about perceptual expertise in cases where it is false, we can reasonably conclude that it is also causally inert in cases where the belief is true.[11] It indicates that perceptual learning privileges

[11] One might object here that if the background belief does slow down perceptual learning in the case where it is false, then we should think that it likewise has an effect in cases where it is true (thanks to Rob Goldstone for

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

the statistical properties of the environment over whatever background beliefs the novice holds. And if the statistical properties of the environment are doing the work in cases where the background belief is false, then we should also hold that they are doing the work in cases where the background belief happens to be true. This is another sense in which the semantic coherence condition is 'lucky': it is in virtue of the truth of the belief in the good cases, rather than the causal efficacy of the belief in altering perceptual experience, that semantic coherence is obtained.

In the second hypothetical scenario, I do not develop perceptual expertise as a result of my training with monarch and viceroy images. My performance is by chance when given an unlabeled image, assuming there is an equal probability of viewing an image of either category. This is not accompanied by any change in perceptual experience. I do not represent the butterfly as a monarch or a viceroy. Instead, I have the phenomenology of guessing. Again, this case does not support the belief-driven hypothesis. Much like the first scenario, the belief is inert in driving perceptual expertise, and so suggests that it is not the belief that is doing the work in the cases where the background belief and perceptual experience happen to cohere.

The final hypothetical scenario to consider is one in which I do not develop perceptual expertise, but where my perceptual experience does change. While I am distinguishing between monarchs and viceroys by chance, I do come to represent these properties in perceptual experience, arbitrarily representing one or the other property. This case would support cognitive permeation, in that it would reveal that background beliefs, even when false, could alter the contents of perceptual experience. In these cases, the semantic coherence condition would be robustly satisfied, in that the content of the belief would semantically cohere with the perceptual experience whether true or false. This in turn would support the claim that the causal condition is satisfied.

Notice, however, that this scenario does not support the belief-driven hypothesis, which I have taken in this paper to be the most plausible account of how diachronic cognitive permeation can lead to perceptual expertise. On this hypothesis, the background belief guides overt, spatial, endogenous attention to features one takes to

raising this point). I think this is right, but consistent with denying that the effect is due to its semantic content. It may instead be understood in terms of myopic flailing, as discussed above, where the false background belief hinders the process.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

be relevant to diagnosis, and this repeated attentional guidance eventually leads to the restructuring of the perceptual system such that the relevant features are picked out together. This perceptual grouping is then thought to correspond to the representation of the high-level properties in perceptual experience. However, in this case the properties one attends to in order to identify a monarch and those one attends to in order to identify a viceroy will be identical, given that we are supposing one is at chance in identification. This suggests that, contrary to our initial description of the scenario, one does not in fact represent the high-level property 'monarch' or 'viceroy' in perceptual experience, but something more akin to a conjunction of both categories, and then arbitrarily attaches one or the other semantic label. In order to make it the case that background belief, regardless of falsity, can alter perceptual experience such that we represent 'viceroy' or 'monarch' in perceptual experience, we must adopt the synchronic view of cognitive permeation. This diverges from the diachronic process of cognitive permeation that is hypothesized to result in perceptual expertise, and so does not provide support for this account.

In summary, whereas at first glance the examples of perceptual expertise I have discussed above seem to support the belief-driven hypothesis (cognitive permeation), they can be accommodated by the pattern-driven hypothesis (perceptual learning). The route by which the contents of perceptual experience become enriched thus seems to involve an increased sensitivity to the environment. This is as James and Eleanor Gibson characterized perceptual learning. As a result of perceptual learning, perception 'is progressively in *greater* correspondence with stimulation, not less' (Gibson and Gibson 1955). Moreover, background beliefs are not required for this enhanced sensitivity to occur. The main take-away point from the discussion is this: just because a learning process involves cognition does not mean that it counts as cognitive permeation. A challenge also now faces proponents of diachronic cognitive permeation. While there may be other cases of perceptual expertise that do support the belief-driven hypothesis, the burden on those who favour cognitive permeation is to provide such examples, and convincingly argue they cannot be accommodated by perceptual learning in the manner I have outlined.

5. Conclusion

In this paper I have argued that perceptual expertise supports the wider scope thesis. On the wider scope thesis the class of immediately

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

justified perceptual beliefs can be enlarged or broadened through learning. Perceptual expertise involves an enrichment of the contents of perceptual experience, and so appears to support the wider scope thesis. However, if this enrichment occurs via cognitive permeation, then – as I have argued here – the post-perceptual beliefs of experts are only mediately justified because such beliefs are epistemically dependent on the permeating background beliefs. In addition, while several theorists have suggested diachronic cognitive permeation as the means by which perceptual expertise is achieved, this suggestion has remained underspecified. I have thus provided a detailed account, in terms of covert, feature-based, endogenous attention – what I call the belief-driven hypothesis – that fills in the gaps.

However, I have argued against the belief-driven hypothesis, defending instead the pattern-driven hypothesis: the change from novice to expert should be understood in terms of perceptual learning, where this does not involve the permeation of background beliefs. I have considered three cases of the acquisition of perceptual expertise, two actual and one hypothetical, arguing in each case that background beliefs are not essential to the learning process, and that when they are involved they play a role that falls short of cognitive permeation, first because they do not fulfill the semantic coherence condition, and second because they are only 'luckily' semantically coherent, failing to fulfill the causal condition. A defence of the pattern-driven hypothesis thus furthers our understanding of the nature of the acquisition of perceptual expertise.[12]

References

William P. Alston, 'What's Wrong with Immediate Knowledge?', *Synthese*, 55 (1983): 73–95.

K.M. Armstrong, 'Covert Spatial Attention and Saccade Planning', In *Attention: Philosophical and Psychological Essays*, edited by Christopher Mole, Declan Smithies, and Wayne Wu, (New York: OUP, 2011, 78–96).

Karlene Ball and Robert Sekuler, 'A Specific and Enduring Improvement in Visual Motion Discrimination', *Science* 218 (1982), 697–98.

[12] I would like to thank Dom Lopes, Chris Mole, Susanna Siegel, Rob Goldstone, Murat Aydede, Evan Thompson, Eric Margolis, Lawrence Ward, and two anonymous reviewers for this journal for their helpful feedback on various versions of this paper. I would also like to acknowledge the generous funding provided to me by the American Society for Aesthetics dissertation fellowship, and the SSHRC Bombardier doctoral scholarship.

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Merim Bilalic, Thomas Grottenthaler, Thomas Nagele, and Tobias Lindig, 'The Faces in Radiological Images: Fusiform Face Area Supports Radiological Expertise', *Cerebral Cortex* 26 (2016).
Laurence Bonjour, 'Can Empirical Knowledge Have a Foundation?' *American Philosophical Quarterly* 15 (1978), 1–13.
Bill Brewer, *Perception and Reason* (Oxford: Oxford University Press, 1999).
Robert Briscoe, 'Cognitive Penetration and the Reach of Phenomenal Content', In *The Cognitive Penetrability of Perception*, edited by John Zeimbekis and Athanassios Raftopoulos, (Oxford: Oxford University Press, 2015, 174–203).
Berit Brogaard and Bartek Chomanski, 'Cognitive Penetrability and High-Level Properties in Perception: Unrelated Phenomena?', *Pacific Philosophical Quarterly*, 96 (2015), 469–86.
Berit Brogaard and Dimitria Electra Gatzia, 'The Real Epistemic Significance of Perceptual Learning', *Inquiry* 61 (2018), 543–58.
Marisa Carrasco, 'Visual Attention: The Past 25 Years', *Vision Research*, 51 (2011), 1484–1525.
Bartek Chomanski and Elijah Chudnoff, 'How Perception Generates, Preserves, and Mediates Justification', *Inquiry* 61 (2018), 559–68.
Elijah Chudnoff, 'The Epistemic Significance of Perceptual Learning', *Inquiry* 61 (2018), 520–42.
Paul M. Churchland, *Scientific Realism and the Plasticity of Mind*. (Cambridge, Massachusetts: Cambridge University Press, 1979).
Paul M. Churchland, 1988. "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor." *Philosophy of Science* 55 (2): 167–87.
Kevin Connolly, *Perceptual Learning: The Flexibility of the Senses*, (Oxford: Oxford University Press, 2019).
Robert Desimone, 'Visual Attention Mediated by Biased Competition in Extrastriate Visual Cortex', Edited by G. W. Humphreys, J. Duncan, and A. Treisman. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 353 (1998), 1245–55.
Fred I Dretske, *Knowledge and the Flow of Information* (MIT Press, 1981).
Suzanne Easton, *An Introduction to Radiography*, (Churchill Livingstone/Elsevier, 2009).
Jerry A. Fodor, 'A Reply to Churchland's 'Perceptual Plasticity and Theoretical Neutrality', *Philosophy of Science*, 55 (1988), 188–98.
Isabel Gauthier and Michael J. Tarr, 'Becoming a 'Greeble' Expert: Exploring Mechanisms for Face Recognition', *Vision Research* 37 (1997), 1673–82.
Isabel Gauthier, Pepper Williams, Michael J. Tarr, and James Tanaka. 1998. "Training 'Greeble' Experts: A Framework for Studying Expert Object Recognition Processes." *Vision Research* 38 (15–16): 2401–28.
James Gibson and Eleanor J. Gibson, 'Perceptual Learning: Differentiation or Enrichment?', *Psychological Review*, 62 (1955), 32–41.
Alvin I Goldman, 'Immediate Justification and Process Reliabilism', In *Epistemology: New Essays*, edited by Quentin Smith, (Oxford: OUP, 2008, 63–82).

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Robert L. Goldstone, 'Perceptual Learning', *Annual Review of Psychology* 49 (1998), 585–612.
Robert L. Goldstone and Lisa A. Byrge, 'Perceptual Learning' In *The Oxford Handbook of Philosophy of Perception*, edited by Mohan Matthen, (Oxford: OUP, 2015, 812–32).
Robert L. Goldstone, David Landy, and Lionel C. Brunel, 'Improving Perception to Make Distant Connections Closer', *Frontiers in Psychology* 2 (2011), 385.
J.R. Hayes, 1989. *The Complete Problem Solver*. 2nd ed. Hillsdale, NJ: Erlbaum.
Michael Huemer, *Skepticism and the Veil of Perception* (Rowman & Littlefield Publishers, 2001).
Thomas S. Kuhn, *The Structure of Scientific Revolutions*, (Chicago: University of Chicago Press, 1962).
Jonathan L. Kvanvig and Wayne D. Riggs, 'Can a Coherence Theory Appeal to Appearance States?', *Philosophical Studies* 67 (1992), 197–217.
Jack Lyons, 'Circularity, Reliability, and the Cognitive Penetrability of Perception', *Philosophical Issues*, 21 (2011), 289–311.
Fiona Macpherson, 'Cognitive Penetration of Colour Experience: Rethinking the Issue in Light of an Indirect Mechanism', *Philosophy and Phenomenological Research*, 84 (2012), 24–62.
David Manning, Susan Ethell, Tim Donovan, and Trevor Crawford, 'How Do Radiologists Do It? The Influence of Experience and Training on Searching for Chest Nodules.' *Radiography* 12 (2006), 134–42.
Peter J. Markie, 'The Mystery of Direct Perceptual Justification', *Philosophical Studies* 126 (2005), 347–73.
John McDowell, 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy* 68 (1982), 455–79.
Matthew McGrath, 'Knowing What Things Look Like', *Philosophical Review* 126 (2017), 1–41.
Sarah McGrath, 'Moral Knowledge by Perception', *Philosophical Perspectives* 18 (2004), 209–28.
Christopher Mole, 'Attention and Cognitive Penetration', In *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, edited by John Zeimbekis and Athanassios Raftopoulos, (Oxford: OUP, 2015, 218–37).
Alvin Plantinga, *Warrant and Proper Function*, (Oxford: Oxford University Press, 1993).
Zenon W. Pylyshyn, 'Is Vision Continuous with Cognition?: The Case for Cognitive Impenetrability of Visual Perception', *Behavioral and Brain Sciences* 22 (1999), 341–65.
Madeleine Ransom, 'Attentional Weighting in Perceptual Learning', *Journal of Consciousness Studies* 27 (2020).
Li Shen, Laurie R. Margolies, Joseph H. Rothstein, Eugene Fluder, Russell McBride, and Weiva Sieh. 'Deep Learning to Improve Breast

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press

Cancer Detection on Screening Mammograph', *Nature: Scientific Reports*, 9 (2019), 12495.
Susanna Siegel, 'Subject and Object in the Contents of Visual Experience', *Philosophical Review*, 115 (2006), 355–88.
Susanna Siegel, *The Contents of Visual Experience*, (New York: Oxford University Press, 2010).
Susanna Siegel, 'Cognitive Penetrability and Perceptual Justification' *Noûs* 46 (2012), 201–22.
Susanna Siegel, *The Rationality of Perception* (Oxford: Oxford University Press, 2017).
C.P. Siewert, *The Significance of Consciousness* (Princeton: Princeton University Press, 1998).
Nicholas Silins, 'The Significance of High-Level Content' *Philosophical Studies* 162 (2013), 13–33.
Dustin Stokes, 'Cognitive Penetrability of Perception' *Philosophy Compass* 8 (2013), 646–63.
Dustin Stokes, 'Cognitive Penetration and the Perception of Art', *Dialectica* 68 (2014), 1–34.
Dustin Stokes, 'Attention and the Cognitive Penetrability of Perception', *Australasian Journal of Philosophy*, 96 (2018), 303–18.
Dustin Stokes, 'On Perceptual Expertise', *Mind & Language*, mila.12270, 2020.
James W. Tanaka, Tim Curran, and David L. Sheinberg, 'The Training and Transfer of Real-World Perceptual Expertise', *Psychological Science* 16 (2005), 145–51.
Michael Tye, *Ten Problems of Consciousness* (MIT Press, 1995).
Bram M.K. Vaassen, 'Basic Beliefs and the Perceptual Learning Problem: A Substantial Challenge for Moderate Foundationalism', *Episteme* 13 (2016), 133–49.
T Watanabe, J E Náñez, and Y Sasaki, 'Perceptual Learning without Perception', *Nature* 413 (2001), 844–48.
Yetta K. Wong, J. R. Folstein, and I. Gauthier, 'Task-Irrelevant Perceptual Expertise', *Journal of Vision*, 11 (2011), 3–3.
Wayne Wu, 'Shaking Up the Mind's Ground Floor: The Cognitive Penetration of Visual Attention', *Journal of Philosophy*, 114 (2017), 5–32.
John Zeimbekis and Athanassios Raftopoulos, eds. *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. (USA: OUP, 2015).

MADELEINE RANSOM (madeleineransom@gmail.com) *is a postdoctoral researcher at the Percepts and Concepts lab at the University of Indiana, Bloomington, where she works on bias in perceptual learning. She has published in Philosophy and Phenomenological Research, and her forthcoming paper, 'Attentional Weighting and Perceptual Learning' – winner of the 8th annual essay prize at the Centre for Philosophical Psychology, University of Antwerp – will appear in the Journal of Consciousness Studies.*

https://doi.org/10.1017/S0031819120000157 Published online by Cambridge University Press