Perceptual Learning, Categorical Perception, and Cognitive Permeation

"Mentioning that a human has two legs is useful for differentiating a person from a goat or a toaster, but it is hard to think up further specification that does not degenerate into a long disjunction of special cases. Even if such an expansion were successful, the resulting tome would no longer serve the purposes of efficient communication or […] instruction."
--Lee Brooks and Samuel Hannah

Daniel C. Burnston
dburnsto@tulane.edu
Tulane University

Keywords: Cognitive Penetration, Categorical Perception, Perceptual Learning, Perceptual Content

Abstract: Proponents of cognitive penetration often argue for the thesis on the basis of combined intuitions about *categorical perception* and *perceptual learning*. The claim is that beliefs penetrate perceptions in the course of learning to perceive categories. I argue that this "diachronic" penetration thesis is false. In order to substantiate a robust notion of penetration, the beliefs that enable learning must describe the particular ability that subjects learn. However, they cannot do so, since in order to help with learning they must instruct learners to employ previously existing abilities. I argue that a better approach recognizes that we can have sophisticated *causal precursors* to perceptual learning, but that the learning process itself must operate outside of cognitive influence.

Acknoweldgements: This paper benefitted from discussions with many colleagues over the long course of its development. Early versions of the paper (or aspects thereof) were presented at the Joint Session for the Aristotelian Society in Cardiff, 2016, at Australia National University and the University of Auckland in 2018, and at the Max Planck School of Cognition in 2021. I'm grateful to the audiences at each venue. The paper was also discussed in research groups at Australia National University and Ruhr University Bochum, and I'm grateful to these groups—especially Peter Brössel, Peter Clutton, Justin D'Ambrosio, Colin Klein, Francesco Marchi, Albert Newen, Ross Pain, and Alfredo Vernazzini—for helpful comments and discussion.

Main Body:

If cognitive penetration occurs, then perceptual experience is affected by the content of cognitive states. This entails that perceptual processes are modified by interaction with cognition.[1] Recently, a number of theorists have argued in favor of the *diachronic cognitive penetration thesis* (dCPT), the view that perception is permeated by cognition in the course of perceptual learning (Cecchi, 2014; Siegel, 2013; Stokes, 2021; Stokes & Bergeron, 2015). The argument for the dCPT is abductive, and is based on an *enabling* claim. The idea is that there are certain kinds of contents, particularly those corresponding to kinds of objects, that perception on its own lacks the capacity to represent, but that interaction with cognition eventually enables them to do so.

It has recently been pointed out that the notion of 'penetration' has gendered connotations (Ransom, 2020). In what follows, I will use 'permeation' instead, but I mean to reference the same thesis philosophers have investigated (so, the thesis under consideration is the diachronic *permeation* thesis, with the same acronym). My aims in this paper are to articulate the commitments of the dCPT, to raise problems for the view, and to propose an alternative for the possible role of beliefs in perceptual learning. In particular I argue that, if the dCPT is true, then category-specific beliefs must be held *prior* to perceptual learning and must *specify* the learned perceptual content. If these conditions are not met, then another thesis is more plausible, namely that cognitive states serve as *causal precursors* to a purely internal process of perceptual learning. I will argue for the second position, enlisting current perspectives from the psychology of perceptual learning.

In section 1, I'll flesh out the dCPT in more detail, and articulate the priority and specificity conditions. In section 2 I'll discuss perceptual learning, and argue for the minimal claim that category-relevant perceptual learning *can* occur without cognitive permeation. Section 3 then raises objections against the dCPT, the key move being to question whether the enabling role posited for beliefs posited by the dCPT is in fact incompatible with their fulfilling the specificity condition. In order to guide learning, I argue, beliefs must describe contents that subjects can *already perceive*. But if that is so, then they cannot describe the novel contents learned. Section 4 considers and rejects ways of weakening the dCPT to avoid this argument. Section 5 concludes.

1. The dCPT

1.1. The dCPT.

Cognitive permeation is an *explanatory* thesis. The idea is that there are perceptual differences amongst perceivers, or within a perceiver over time, and the best explanation for those differences is that the contents of cognitive states have modified perceptual

---

[1] There are a variety of ways of describing this relationship: perceptual representations must bear a "logical relationship" to knowledge (Pylyshyn, 1999); perception uses cognition as an "informational resource" (Wu, 2013); there is an "inferential" relationship between cognitive states and the outputs of perception (Brogaard & Chomanski, 2015). These characterizations are highly ambiguous (Burnston, 2017a), but I will take them as read here.

processing. While the thesis itself is internal to the philosophy of psychology, its ramifications are potentially widespread.

One reason for these widespread ramifications is the thought that cognitive permeation might be one way in which perception can be *enriched*. If our percepts come to reflect our beliefs or theoretical assumptions, then they represent more about the world than just what they can glean from sensory input. In turn, cognitive permeation has been proposed as one potential explanation for how perception comes to represent higher-level properties, i.e. categories beyond simple perceptual features like shape and color (Siegel, 2013). It has been used as one way of explaining the kinds of dispositions developed by skilled perfomers (Fridland, 2015), to account for moral perception (Cowan, 2014) and, more recently, for a wider range of expertise effects (Ransom, 2020; Stokes, 2021)[2]. Theorists have gone on to consider the epistemic upshot of enriched perception, arguing both for its potential benefits and its potential detriments for perceptual justification (Siegel, 2012, 2016; Stokes, 2021).

So, the issue of whether cognitive permeation occurs is important for a range of philosophical enterprises. Unfortunately, the extensive debate about cognitive permeation has failed to produce even an agreed-upon definition of the thesis. Theorists disagree on, amongst other things: whether cognitive influence on perception must be direct (Macpherson, 2012; Raftopoulos, 2015); whether causal interactions between them are sufficient, or whether stronger semantic and computational relationships are required (Stokes, 2013; Wu, 2017; Burnston, 2017a); whether cognitive permeation results in representation of higher-level contents or in changes to lower-level properties (Briscoe, 2015; Siegel, 2013; Stokes & Bergeron, 2015); and, importantly, whether attentional effects count as instances of permeation (Gross, 2017; Marchi, 2017; Mole, 2015; Stokes, 2018). Some have even suggested that cognitive permeation should be characterized purely according to its *consequences* for relevant philosophical debates (Stokes, 2015).

While there are different characterizations of cognitive permeation, one thing should not be up for debate, namely that the truth of the cognitive permeation thesis would be a *surprising* and *transformative* result for our understanding of the mind. The idea is that certain empirical and theoretical considerations force us to give up the intuitive view that changing our beliefs does not change what we perceive (Firestone & Scholl, 2016). So, when considering the cognitive permeation thesis, we should ask whether the kind of relations discovered between cognition and perception prompt this sort of foundational change to our understanding, or whether more mundane notions can capture the evidence at hand. It is in this spirit that the present paper attempts to assess the issue.

[2] The papers by Ransom (2020) and Stokes (2021) were published while this paper was in submission. There is considerable commonality between my conclusion and Ransom's. That said, I focus on slightly different phenomena in perceptual learning than Ransom does, and provide distinct (if compatible) arguments against the dCPT.

One thread of the argument that has been present since early discussion of cognitive permeation, is whether it occurs through perceptual *learning*. Churchland (1988) classically argued that, to the extent that perceptual systems were plastic, they were likely to be infiltrated by knowledge, and hence that perception is likely theory laden. Recent interest in learning has picked back up, as it is one potential explanation for the existence of higher-level content and for perceptual expertise. And there is strong reason to focus on perceptual learning as a test case. For one thing, learning often involves changes in one's beliefs, and therefore is one possible case in which a change in belief could eventuate a change in perception. Moreover, expertise often involves *training*, wherein one intentionally focuses on certain features of examples in order to develop one's abilities. If perception is changed during this process, then it seems a likely case for cognitive permeation.

Here as well, however, we find a diversity of views. Perceptual learning itself is defined in different ways, sometimes in terms of generated perceptual *abilities*—e.g. of discrimination or generalization—and sometimes in terms of changes in perceptual *contents* (Connolly, 2014, 2019; Prettyman, 2018). Some have advanced the position that, for certain instances of perceptual learning, learning effects are evidence that cognitive permeation occurs. Stokes and Bergeron (2015), for example cite cases of *categorical perception*, on which learned categories modify perception, as proof that cognitive permeation occurs during perceptual learning, while Firestone and Scholl (2016; cf. Valenti & Firestone, 2019) challenge this view. Others, as mentioned, only take cognitive permeation as one *possible* explanation for perceptual learning (Siegel, 2013), or consider that some instances of learning may be instances of permeation and others not (Stokes, 2021). And some, indeed, propose that perceptual learning is an *alternative* to cognitive permeation (Connolly, 2014; Arstila, 2016), i.e. that learning within the perceptual system is an alternative explanation to the permeation thesis.

This is a tangled, almost bewildering set of considerations, and I want to remain neutral towards as many of them as possible. I assume that genuine perceptual learning occurs, which modifies perceptual representations and results in novel perceptual abilities. I will further discuss evidence that such changes occur at several "levels" of perception, although I will remain neutral on whether the higher levels constitute higher-level contents. (For an extended discussion of the relationship between these representations and the debate on higher-level contents, see Burnston, in submission.) The question then is whether, in any of these cases, cognitive permeation is the right explanation of effects in perceptual learning. I will use the language of perceptual representation and perceptual content, but I do not commit in this paper to any particular way of typing contents. Instead, I will try to describe the representations at work as directly as possible.

I will assume a broadly semantic conception of cognitive permeation at the outset, and I will consider later whether one can abandon this conception. According to the semantic conception, a specific change occurs within perception and is *explained* by the content of the permeating state. This entails that perception has access to or processes the content of cognitive states (Wu, 2017; Ransom, 2020). Further, it entails that the contents of cognition can explain the changes in the contents of perception. That is, perception operates differently after learning, and the reason for *that specific change* is that it has taken the contents of cognitive states into account in modifying its processing. I further presume that attentional mediation is one good candidate for a mechanism that might bring that change about. That is, cognitive instruction to intend to a stimulus in such-and-such a way is one plausible way in which cognitive permeation could occur. The question I will consider is whether, given the empirical data on perceptual learning, cognitive permeation of this sort is a good explanation for that learning.

This focus on explanation fits well with the abductive nature of arguments that many proponents of cognitive permeation espouse. After looking at a range of effects, these theorists argue, the best account of changes to perceptual experience is permeation (Stokes, 2021; Stokes & Bergeron, 2015). In particular, I am interested in a variety of *enabling* claim. The idea here is that cognitive permeation—i.e. the resources provided by cognitive contents—allows perception to work in a way that it could not on its own. So, for instance, Stokes and Bergeron argue that, while perception may have evolved a capacity to represent faces, "there is no account to be given about the evolution or plasticity of perception for the Pink Panther or the Coca-Cola icon" (p. 16; cf. Stokes & Bergeron, 2015, p. 325). If perception *itself* lacks the resources to discriminate these categories, then perhaps processing cognitive contents is how perception comes to do so. Similarly, Cecchi suggests that, when perceptual learning occurs during intentional practice at a task, it is because "cognitively induced architectural modulations enable […] the visual system to perform the […] task" (2014, p. 91).

So, finally, my construal of the dCPT is this. Perception develops novel abilities during the course of perceptual learning, and the explanation for how it does so is that it processes cognitive contents. In the next section I explore the commitments of this kind of view, and articulate an alternative, namely that the role of cognition in perceptual learning is merely to serve a as more-or-less sophisticated *causal precursor* to a purely internal process of perceptual learning. An effect that can be explained as a causal precursor to a change in perception is not sufficient to compel the transformative consequences that cognitive permeation is supposed to have. The position I will argue for is that, although causal precursors can be quite important and specific, the learning that perception does is based solely on interaction with a stimulus-set, not on processing cognitive contents (cf. Ransom, 2020).

1.2. Candidates and conditions

Given the enabling role in perceptual learning that is posited for beliefs by the dCPT, the first condition that any potential permeator should meet is what I'll call the *priority* condition. Since it is the presence of beliefs that is supposed to enable perceptual learning, those beliefs must be ones that the subject plausibly possesses before the content is learned.

Meeting the priority condition, though, is insufficient, since there are many kinds of beliefs that could meet the condition but fail to be good candidate permeators. Here are three kinds of beliefs that are poor candidates for implementing diachronic permeation.

The first is *essentialist* beliefs. Suppose that you know something about the respective chemical structures of jadeite and nephrite, or the facts about phylogenetic history that distinguish whales from fish. While the propositions that are the contents of these beliefs are (at least if you're an essentialist) definitive of the categories to which they apply, the contents of the beliefs themselves have no upshot for how the categories should be *perceived*. Knowledge of chemical structure doesn't help you perceptually discriminate jadeite from nephrite. Similarly, knowledge about cladistics doesn't suggest modifying your percepts of whales or fish in any particular way. This is true even if you hold the beliefs prior to learning to perceive the kind.

A second poor set of candidates is *demonstrative* beliefs. Suppose I hand you an object of a type you've never seen before, and say "this is a glunk." You might reasonably form the belief that the object you are now holding is a glunk. It is true that the demonstrative 'this' *refers* to a particular glunk, but the simple content of the term doesn't contain the resources to help you learn what's *perceptually* characteristic of glunks. Indeed, the belief would play the same role *no matter what* glunks in fact look like. Hence, the demonstrative belief doesn't have the right kind of content to inform perceptual learning. (I'll discuss this example further in the next section.)

A third kind of beliefs, which we might call *denotational* beliefs, have more content than bare demonstratives, but their primary role is still to *pick out the category to be learned*. So, suppose you're about to walk into a room full of objects, and I tell you "The glunks are on the far left." This belief might help you figure out which are the glunks, by providing a behavioral instruction to look at some objects rather than others. As in the demonstrative case, however, the content of the belief has no resources to inform the actual perceptual category you might learn. Again, the belief will play the same role no matter what perceptual characteristics actually individuate glunks, and hence cannot inform perception how glunks should be represented.

These considerations suggest that another condition is needed, which I will call the *specificity* condition: a candidate permeator must have sufficiently specific content to inform the *particular perceptual content* that is learned. A belief that meets the priority condition but not the specificity condition, I suggest, is best construed as a *causal precursor* to an instance of perceptual learning. A belief or other cognitive state's being a causal precursor to a percept, nearly everyone acknowledges, is not sufficient to make that belief a permeator of the percept. Suppose you know that a particular bird nests only on sheltered alcoves atop very high mountain ranges. This knowledge, along with some sophisticated knowledge about how to climb mountains, might eventuate in your learning to perceive baby birds of that type. But your knowledge of the location of the birds and how to navigate to a place where you can see them does not tell perception anything much about what it should do to recognize that type of baby bird specifically.

This is true even if the knowledge is a *necessary* precursor—i.e. if climbing were the only way you could ever gain access to the birds.

Importantly, we now have an alternative interpretation of the "enabling" effect of cognitive states on some perceptual process. On this view, enabling beliefs are only causal precursors—they might point you in the direction of the objects-to-be-perceived, but do not permeate the eventual learned perception. Only beliefs that meet the specificity condition in addition to the priority condition would force us to read enabling effects in terms of permeation. In section 5, I will consider whether a proponent of the dCPT can reasonably give up on or try to weaken the priority and specificity conditions while still offering an interesting thesis. For now, I will assume that both the priority and specificity conditions must be met by any successful candidate permeator.

Given these considerations, the *prima facie* best candidate for a type of belief that might permeate perceptual learning—and the one that I think most defenders of the dCPT have in mind—is *descriptivist beliefs*. These beliefs have as their content the properties, including the perceptible properties, that members of a kind have. Dachshunds, for instance, are long, brown, and short-legged. Maybe the belief that glunks are (say) *large* and *green* has the right kind of content to permeate perceptual learning, even if demonstrative or denotational beliefs do not. This view has some backing: Leslie (2008) has argued that "generic" beliefs about kinds are fundamental to cognition and learning, and generics often have descriptivist content—e.g. "Tigers have stripes." Reliance on descriptivist beliefs is perhaps the way to interpret Siegel's claim that we learn to recognize pine trees by coming to believe that they have "certain kinds of leaves and structure" (2013, p. 715), or Stokes' (2014) claim that we learn to recognize Mondrian's paintings in virtue of forming beliefs about their "organizational features" (p. 17). The question is, can these kinds of beliefs meet the priority and specificity conditions?

In the next section I will outline the relevant psychological results on perceptual learning. I'll argue that in certain instances, perceptual learning of categorical content occurs without cognitive permeation. This will then provide the framework for asking whether descriptivist beliefs in general have the right kind of content to implement the dCPT.

2. Perceptual Learning

In this section, I will describe a current perspective on perceptual category learning.[3] The core idea behind the framework is what is called a "morphspace." Perceptual learning, the story goes, forms categories by differentiating and accentuating *dimensions of perceptual difference* between kinds of objects. According to the morphspace view, each

[3] See, e.g. Folstein, Gauthier, & Palmeri, 2010; Folstein, Gauthier, & Palmeri, 2012; Gauthier & Tarr, 2002; Goldstone, 1994; Goldstone & Hendrickson, 2010; Goldstone, Lippa, & Shiffrin, 2001; Goldstone & Steyvers, 2001; Gureckis & Goldstone, 2008.

perceptual category corresponds to a "space" defined along relevant dimensions. Dimensions can either correspond to low- level perceptual features (size, luminance, etc.) or to higher-order relationships between these features. Important higher-order relationships involve configural (arrangement in space) and associational (correlation) relations between lower-order dimension values. Perceptual learning, on this perspective, can both learn novel dimensions and modify extant dimensions. As categories are learned, dimensions can be "morphed" so that intra-category members are seen as closer along the relevant dimensions.[4]

While I will not make any explicit claims about higher-level content, the distinction between lower- and higher-order dimensions does show that perceptual learning operates at several distinct "levels," which correspond to novel categories. I do suggest that this kind of learning underlies the kinds of recognitional dispositions that some take to be indicative of higher-level content. My purpose in this section is to argue that this kind of learning *can* operate to develop and modify morphspaces in category-specific ways without cognitive influence. This will allow us to then question whether the dCPT is the best explanation of perceptual learning in *some* cases.

Studies in categorical perception rely on training with exemplars, either with or without feedback. Studies without feedback show that subjects can form higher-order dimensions through *mere exposure*. Folstein et al. (2010) showed subjects a range of cartoon creatures (see Figure 1), where in the training set there were correlations between different lower-level features—for instance, particular wing shapes and head shapes, as well as particular body and arm shapes, might be correlated with each other, where there was no such relation between (e.g.) wings and legs. Having multiple correlations present in the same stimuli set was done to prevent subjects from forming, unprompted, specific beliefs about category-membership. Subjects also performed a distractor task (judging how centered the stimulus was on the screen), which was intended to prevent them from forming category beliefs.

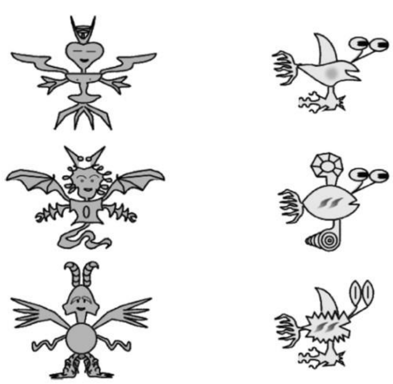

Figure 1. Stimuli from Folstein et al. (2010).

[4] This can sometimes correspond to a *loss* of discriminatory capacity *within* a category. In general, perceptual category learning is a generalization and discrimination tradeoff.

In the experiment, subjects were capable of picking up on the higher-order correlations in the stimuli. This was shown by a secondary task, in which subjects had to categorize novel examples. If categories in the second task matched the correlations in the exposure set, subjects learned them more quickly than if the categories did not match those correlations. If the controls worked, then subjects are capable of this kind of learning even if they form no category-specific beliefs. Folstein et al.'s interpretation is that it is possible for perception to form novel higher-order dimensions purely through "statistical perceptual learning," without influence from beliefs. Similar results have been shown for other kinds of stimuli (Fiser & Aslin, 2001].

Importantly for what follows, there are cases where category-specific feedback is provided to perceivers, and this feedback plays a role in learning, but where, I will argue, the feedback does not meet the specificity condition.

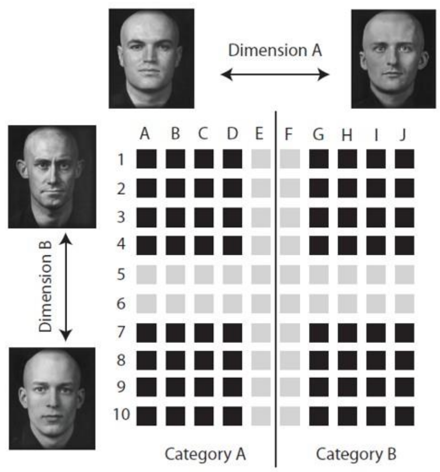

Figure 2. From Gureckis and Goldstone (2008).

Figure 2 is an example from a wide range of studies in which subjects learn to differentiate objects along *arbitrary* dimensions (Goldstone & Steyvers, 2001; Gureckis & Goldstone, 2008; Folstein et al., 2012; Jones & Goldstone, 2013). The experimenters created a morphspace of faces by taking four distinct faces, and creating exemplars (each square in Figure 2) that continuously blended each of their features. Subjects in these studies were shown exemplar pictures, and told whether each was an "A" face or a "B" face, where As and Bs were defined according to the arbitrary vertical line in the center of the space.

A variety of results from this kind of paradigm suggest that subjects learn to differentiate the novel dimensions of the space. For instance, after learning to make the discrimination, the dimensions transfer to new stimuli and categorizations (Goldstone and Steyvers, 2001), such that subsequent discriminations along the previously learned dimension are easier than along other dimensions. Moreover, these learned representation affect *similarity judgments*. Across a range of types of similarity judgment, subjects tend to treat within-category members as more similar to each other after training than they did before training. In the example from figure 2, this involves differentiating and then morphing the category-distinguishing horizontal "Dimension A."

Hence, in this and other cases, perceptual learning both forms novel dimensions and stretches the morphspace along those dimensions to accentuate the difference between categories. But this kind of learning also modifies representations of lower-level features. Consider two cases. On one, subjects might learn to accentuate discriminations made along already differentiable lower-level dimensions. On the other, subjects might learn to differentiate *between* lower-level dimensions that they could not previously tell apart. Goldstone (1994) investigated both types of changes.

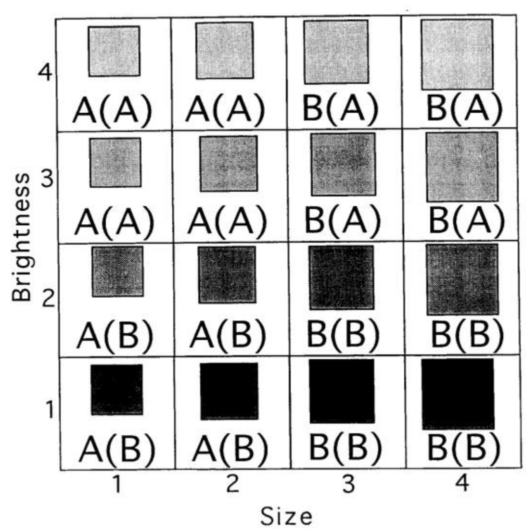

Figure 3. A morphspace for lower-level properties.

As shown in Figure 3, Goldstone constructed a simple morphspace of squares comprising two lower-level dimensions, brightness and size. He then tested a variety of different categories defined in the space. For instance, size might be relevant and brightness irrelevant (i.e. drawing the categorical line between 2 and 3 on the y-axis) or vice versa (i.e. drawing the categorical line between 2 and 3 on the x-axis). Take just the brightness-relevant, size-irrelevant case. In this case subjects learned, through feedback of the type discussed above, which squares belonged to which category. After training, their discriminations along the brightness dimension were *heightened* while their discriminations along the size dimension were *diminished*. That is, they became more sensitive to differences between levels of brightness and less sensitive to differences between sizes.

Goldstone (1994) also created a morphspace of squares based on levels of brightness and *saturation*, rather than size. What is interesting about this case is that, while brightness and saturation can be independently manipulated by an experimenter, subjects do not normally perceive them independently. That is, they cannot selectively attend to one rather than another, or make discriminations along one independently of variation in the other. Goldstone trained subjects in a similar way on these stimuli, separating them into categories based on either brightness or saturation, and training subjects on exemplars with feedback. Somewhat amazingly, subjects *do* in fact begin to differentiate the dimensions, showing similar (although smaller magnitude) learning effects as in the brightness and size case.

So, learning can both differentiate new dimensions, and modify existing dimensions to notice and accentuate category differences. There are now two questions to pursue. The first is whether the subjects' new representations can be considered genuinely perceptual, and then the second is whether they are the result of cognitive permeation.

I suggest that there are two sets of interlocking reasons to consider the learned abilities here as due to changes in perceptual representation. First, notice that it is possible to form these representations *absent* category-specific beliefs. In the mere exposure case of Folstein et al., subjects have no prior beliefs about what will constitute the categories, and the control task and multi-associational structure was set up specifically to prevent subjects from forming those beliefs during learning. If these manipulations worked, then perception can learn categorical content even with no relevant beliefs about the category.

Moreover, even when subjects were asked to reflect on their judgments and describe them, their explanations were often coarse grained and mapped poorly to the representations that guided their judgments. For example, Goldstone and Steyvers (2001) report that subjects sometimes used abstract language to describe the categories, including such statements as "Faces in [category A] were happier" (Goldstone & Steyvers, 2001, p. 135). But this is clearly not specific enough to have informed their judgments. Subjects presumably can already discriminate happy from unhappy looking faces, but this is not, prior to training, detailed enough for them to discriminate *these* categories of faces from each other. Moreover, Goldstone (personal communication) notes that different subjects sometimes use similar descriptive language even if they have learned different categorizations, thus proving that their beliefs do not discriminate the categories, or at least that they are not required to do so in order for subjects to learn the categories.

The second set of reasons stresses the *structure* of these learned representations, i.e. their dimensional structure. In these studies, there is a continuous pattern of variation amongst the conjunction of features that comprise the examples. What subjects do is learn to represent this pattern of variation, suggesting that the representations learned have a kind of *metric* structure (Burnston, 2017). While it is true that, according to the morphspace framework, the dimensions can be morphed to accentuate category difference, this morphing is done within the metric structure—that is, what is modified is the distance metric between examples along the relevant dimensions, and this is what underlies the changes in similarity judgments.

Arguably, representations of this type do not meet some standard conditions on concept possession, such as the generality constraint (Beck, 2014). If one thinks that a way of distinguishing the conceptual from the non-conceptual is in terms of the structure of representations, and thinks further that the way to distinguish the perceptual from the cognitive is in terms of the conceptual/non-conceptual divide, then one will be strongly motivated to view these representations as perceptual.[5]

Moreover, this way of thinking corresponds with some traditional motivations in the non-conceptual content literature, which has historically focused on perceptual content. Notice that subjects develop the ability to apply new categorical concepts demonstratively—e.g. "that's an A face". But, this ability requires a *previously existing* representation which serves as the ground for that demonstrative reference (see further discussion below). This kind of point has been used by defenders of non-conceptual content to combat the idea that all perceptual content requires demonstrative concepts (Roskies, 2010).

I consider these reasons to be mutually supportive rather than decisive. But recall the dialectic here—proponents of the dCPT posit genuinely *perceptual* contents which are the result of cognitive permeation. Hence, there should be strong motivation for the dCPT proponent to accept these learned representations as perceptual. The question, then, is whether these cases are instances of cognitive permeation. I suggest that whatever beliefs subjects have are inadequate to meet the specificity and priority conditions.

Consider the mere exposure cases first. If the manipulations and controls worked, then in this case subjects simply had no category-relevant beliefs prior to learning, and the resulting representations could not be due to permeation. Next, consider the face case, in which subjects do in fact form beliefs during the learning process. In particular, based on the feedback, they are in a position to form a series of demonstrative beliefs, such as that a particular example was an 'A' face. In the last section, I suggested that demonstrative beliefs of this sort do not have the right kind of content to meet the specificity condition. Let's consider this a little further.

There are two ways of individuating belief contents, narrow and wide. Speaking *very* loosely: narrow contents correspond to what the subject is prepared to *do* in virtue of a belief—i.e. the effect that the belief has on other psychological processes and behavior. Wide contents correspond to the extension of the belief. Neither way of individuating demonstrative belief contents supports a reading in terms of cognitive permeation.

The narrow content that the demonstrative belief has is very sparse. It can only convey something like "treat this as an A face." Even a *series* of demonstrative beliefs of this sort, corresponding to a set of A faces, can only result in something like "treat these all

[5] Beck (2014) argues that representations of this type can be *cognitive* but *non-conceptual*. He bases this argument on analogue magnitude representations which, he argues, can be abstracted from any particular instance of magnitude judgment. The kinds of representations I am discussing here, conversely, cannot be so dissociated from their instantiations. I discuss this at length in (Burnston, in submission).

as A faces." But this content has no upshot for the *perceptual* recognition of A faces. Knowing that a set of objects should be treated as belonging to the same category doesn't say anything about the perceptual space that they share. Moreover, simply forming a series of demonstrative beliefs about already-seen exemplars does not, on its own, say anything about how the category should be morphed or extended to novel exemplars, but this is precisely what the learned perceptual representation does.[6]

So, narrow individuation for demonstrative beliefs won't secure a result of permeation. Wide individuation fares no better. On the wide individuation, we could construe the demonstrative belief as conveying the content that the particular face is a member of the set of all A faces, where perhaps one defers to the experimenter to determine the extension of the belief. That's nice enough, but again simply referring to the set does nothing to inform the ability that subjects actually learn, which is to *recognize* that a novel example is a member of the set. So, the demonstrative beliefs that subjects might form on the basis of these instructions cannot meet the specificity condition. As discussed in the previous section, the best way to describe the role of the beliefs is as causal precursors. The demonstrative beliefs provide a behavioral instruction to look for commonalities among a set of objects. Perceptual learning then does the work of actually forming the discriminating representations.

One objection to this view would argue that subjects rely on *tacit* knowledge. One might suggest that, in mere exposure cases, the controls were insufficient to rule out the forming of tacit beliefs about what individuates the stimuli. And, in the feedback case, perhaps subjects develop tacit knowledge about how to apply their new concept of an "A face," that they fail to articulate when asked, but which shapes perceptual learning nonetheless. On this objection, subjects' tacit beliefs might permeate perceptual learning in these cases.

I will consider this objection further in section 5. For now, there are two points to be made about it. First, it is non-trivial to articulate the tacit knowledge objection in a way that is not question begging. Recall that proponents of the dCPT *must* admit that there are perceptual learning processes that result in novel perceptual content. I have suggested here that, at least in important cases, these processes can happen without cognitive permeation. Insisting that the process *must* be due to tacit beliefs in spite of the arguments above sounds suspiciously like a definitional claim that learning must be due to permeation. A definitional claim is out-of-keeping with the kinds of empirical causal arguments put forward by the proponents of the dCPT.

Second, invocation of tacit beliefs is often motivated by dispositionalist concerns, for

---

[6] Now, I am not trying to deny the importance of labels in general – it has been shown in many instances that labels can provide powerful perceptual/attentional cues (although they do not have to be semantically specific to the object to do so; Lupyan & Spivey, 2010). Simply labeling something, however, does nothing to say how that object should be perceived, hence the purely demonstrative function of these beliefs. *After* one can perceive the object a label can provide a powerful attentional cue, but the label itself does not instruct perception to modify its processes in any particular way (Burnston, 2017).

instance the fact that people are inclined to assent to many more propositions than those for which they are likely to have explicit, stored propositional representations (e.g. "Neither cats nor dogs are numbers"; see Schwitzgebel, 2015). But dispositionalists are *not* committed to a particular underlying psychological nature of the mechanism that produces the disposition, and thus are not antagonistic to the possibility that the dispositions are underlain by perceptual states. Given that neo-empiricist accounts of knowledge are at least on the table in cognitive science, it is illicit to simply *assert* that the presence of tacit knowledge means there is an influence of a propositional state on a perceptual process.

Neither of these points is decisive. All I hope to have suggested here is that an appeal to tacit knowledge can't just be a trump card in this debate. It has to be accompanied by specific claims about the contents of those beliefs and how they affect perceptual learning. I will give reasons in the following sections to suggest that the dCPT proponent has no way of formulating this kind of proposal that will meet the priority and specificity conditions.

Even if you grant me all this, I have only established that cognitive permeation doesn't occur *in these cases*. Proponents of cognitive permeation, however, generally don't insist that perception is *always* permeated, only that it is *permeable*. That is, they suggest that in *some* cases perception is cognitively permeated. Perhaps, even if I am right about these cases, cognition permeates perceptual learning in other cases. Indeed, proponents of the dCPT often invoke particular kinds or particular types of learning/expertise for which permeation is a likely explanation of developed perceptual content. The perceptual abilities of arborists or art experts might fall into this category (Siegel, 2012, 2013; Stokes, 2014) propose. And Stokes and Bergeron (2015), as discussed above, suggest that perception cannot, on its own, learn to represent particularly novel kinds such as cultural icons.

In the remainder of the paper I will question whether, even in cases of expertise and highly novel kinds, the dCPT is the correct account of how perception comes to represent categorical content.

3. Against the dCPT.

3.1. Setup

The dCPT posits that certain instances of perceptual modification cannot be explained by citing purely internal processes of perceptual learning. This is what underlies the abductive inference that, in certain instances, cases of perceptual learning must be due to cognitive permeation. As mentioned, one of the motivations for this view is that learning processes are often highly mediated—they involve intentional, knowledge-based learning that requires explicit belief formation and practice. In the previous sections, I suggested that there *is* empirical evidence that perception can learn to distinguish categories of objects, and that this involves changes to perceptual representations at multiple "levels," but questioned whether this process *must* be a result of cognitive permeation. I also offered an alternative, on which beliefs are merely causal precursors for an independent process of perceptual learning.

In this section, I consider whether the dCPT is likely to be the best explanation of perceptual learning, even in cases of novel, socially mediated kinds, or kinds requiring expert training. Importantly, the morphspace framework developed above has been extended to artefactual kinds such as cars (Folstein et al., 2012) as well as to kinds involving developed expertise, such as subspecies of birds (Tanaka, Curran, & Scheinberg, 2005). We can thus ask whether, for these kinds of cases, the processes discussed above require cognitive permeation. In section 3.2., I seek to loosen the supposedly close connection between category knowledge and categorical perception, by suggesting that our descriptivist beliefs regarding kinds are very frequently *equivocal* with regards to perceptual categories. In section 3.3., I make a stronger argument that the enabling role posited for descriptivist beliefs by the dCPT is in fact *incompatible* with their meeting the specificity condition. If I am right, then the causal precursor view is the better interpretation of how prior beliefs interact with perceptual learning.

3.2. Equivocal Descriptions

The dCPT suggests that perceptually learned representations are due to permeation of perception by descriptivist beliefs. Let's start to assess this claim by considering the category 'Smurfs', a cultural artefact if ever there was one, and one that Stokes and Bergeron list as a good case to be explained by the dCPT. What kind of beliefs might one have about Smurfs *prior* to learning how to perceive the category? Here's one candidate list: 'Smurfs are small cartoon people'; 'Smurfs wear red hats'; 'Some Smurfs have beards'. (Perhaps one forms these beliefs by talking to a neighbor about their kids' favorite cartoons, or something.) Alas, this descriptivist content won't discriminate between Smurfs and Gnomes of the sort pictured below.

Figure 4. Small cartoon people in red hats.

Anyone with a modicum of experience with these two different cartoons will be able to discriminate Smurfs from Gnomes. But, patently, the list above does not make a discrimination between these two kinds, since it applies equally to either case. Hence, this set of descriptivist beliefs, even if it met priority, would fail specificity— the descriptivist belief that the learner has is equivocal between the two perceptual categories, but the perceiver precisely learns to distinguish between those categories.

The natural response here is to posit that learners have more detailed beliefs about the categories. One might need the belief that Smurfs are *blue* (although one would then need the further clarification that it is their *skin*, and not their shirt, that is blue) to pick out the Smurfs rather than the Gnomes. Here is the problem with this. The dCPT suggests that perception on its own *cannot* come to discriminate the relevant kinds. As such, the view is committed to the idea that subjects who do learn the perceptual discrimination have prior beliefs that *do* distinguish the categories, since these are what enable the subsequently developed perceptual ability. As cases become more fine-grained, this requires that the grain of subjects' prior belief become comparatively more fine-grained. The idea that all learners have prior beliefs at the requisite level of grain before learning to recognize kinds is, empirically speaking, just unlikely to be correct. Consider a more-fine grained discrimination between "Gnomes" and "Littl' Bits":

Figure 5. Gnomes (left) and Littl' Bits (right).

Despite the significant similarity in terms of their features (they both wear pointy hats, the girls wear red dresses, the boys blue shirtsleeves, they both have small noses and big cheeks, etc., etc.), anyone who has watched a lot of both "David the Gnome" and "The Littl' Bits" can easily make the perceptual discrimination between one and the other. The proponent of the dCPT is forced into a pretty awkward stance regarding these kinds. They must insist that anyone who learns to make this discrimination, *prior to learning to do so*, has sufficiently fine-grained beliefs to inform the perceptual categories. I submit that, in my own case, this is not what happened. At least, it certainly wasn't the case that, when I was four years old watching these cartoons, anyone sat me down and gave me a thorough list of things to look for before I learned how to recognize the kinds. The proponent of the dCPT is forced into trying to articulate a process by which novice perceivers come by very fine-grained perceptual beliefs prior to learning, or they must admit that perceptual learning forms the discriminating representations on its own, at least in many cases.

If perception *can* learn to discriminate these categories, at least in many cases, absent permeation, the inference to the best explanation posited by the dCPT is strongly questioned. On the other hand, the view that beliefs are important causal precursors to perceptual learning, I suggest, is fully compatible with the datum that our descriptive beliefs are often equivocal between kinds that we can easily discriminate perceptually, at least after some learning. Consider the beliefs one is actually likely to have prior to category learning in everyday contexts. Sure, this set may include some beliefs describing general perceptual features. But it is also likely to include beliefs about when and where to find the objects. One might know, for instance, that *The Smurfs* is on Nickelodeon at 4, whereas *David the Gnome* is on at 5. This could help you learn to discriminate the objects without requiring fine-grained descriptivist beliefs that meet the specificity condition.

What I am ultimately suggesting is that descriptivist beliefs are just *another variety* of denotational beliefs. Just like I might say, "Glunks are the objects on the left," I might say, "Smurfs are the little cartoon people in red hats that are on Nickelodeon at 4." What each set of beliefs does is help you *locate the set of objects to be learned*, so that these can be treated as exemplars for the category. But to play this role, all descriptivist beliefs have to do is enable you to sort the Xs from the non-Xs. And so long as the beliefs are descriptive *enough* to sort the exemplars appropriately, they will do the job. That is, they can do the job without specifically describing how Xs should be represented perceptually. As with the birds-on-top-of-mountains example discussed in section 2, this prior knowledge can play an important role in learning, but doesn't need to do so via describing to perception the content it should learn. A similar point goes for guided attention. All that one has to do attentively is focus on the right objects so that perceptual learning can go to work, and if I am right, then that's all that they do. (This is the point of the quote in the epigraph from Brooks and Hannah, 2006.)

This is true even when descriptivist beliefs are very specific. Return again to figure 5. You may or may not have noticed that Littl' Bits, but not Gnomes, have little red dots on their cheeks. My telling you this might indeed help you look at them and say "Ah, ok, these are the Gnomes and these are the Littl' Bits". But notice how far short the content comes of describing the *perceptual* category learned, at least if the morphspace view is correct. If the morphspace view is correct, then what is definitive of the *perceptual* category of Little Bits is not just their red cheeks—it's a complex set of correlations and configural relationships between lower-order properties (the shape and spacing of facial and bodily features, etc.). *This* content, however, is not described by your knowledge of red cheeks. Again, it has served as a (albeit important) causal precursor to perceptual learning.

All I have established so far, however, is that in many quotidian cases we can expect descriptive beliefs to fail to be fine-grained enough to describe learned perceptual content, and hence to deny the view that perception on its own cannot learn to discriminate categories. In many ways, expertise-through-training is the best case for the dCPT theorist. In these kinds of examples, learners are often specifically encouraged to look for certain features of the objects that fall within categories. In the next subsection, I offer an argument that even this apparently obvious case is misleading. Indeed, I will suggest that the enabling role posited by the dCPT is actually *incompatible* with descriptivist beliefs meeting the specificity condition.

3.3. An Incompatibility Argument.

The kind of argument I take to most strongly speak against the dCPT suggests that, precisely *because* of the enabling role posited for descriptivist beliefs by the dCPT, they cannot meet the specificity condition. Put informally, the concern is this. Learning requires leveraging *extant* abilities in service of developing new ones. Training and expertise indeed involves describing the objects to be recognized, but in order to help the trainee, these descriptions must tell subjects what to *do*. That is, they must invoke them to focus on certain objects or properties they can *already* perceive, on pain of being unhelpful for learning. But if descriptions name already-perceivable content, and what subjects learn is *novel* perceptual content, then the specificity condition cannot be met. (This is, basically, a variety of Meno's paradox for the cognition-perception interaction.) Here is the argument in more formal gloss.

1. If the dCPT is true, then prior descriptivist beliefs that meet the specificity condition enable perceptual learning.
2. In order to enable learning, descriptivist beliefs must have as their content perceptual features that subjects can already perceive.
3. Perceptual learning results in novel perceptual content.
4. Novel perceptual content is distinct from content that subjects can already perceive.

Therefore,

5. The content of the beliefs that enable learning is distinct from the perceptual content that is learned. (From 2 and 4).

Therefore,

6. It is not the case that prior descriptivist beliefs that meet the specificity condition enable learning particular perceptual categories. (From 5)

Therefore,

7. It is not the case that the dCPT is true. (From 1 and 6 via Modus Tollens)

The argument hangs on premise (2), and the move from (5) to (6). Premise (3) is granted by all parties, and premise (4) is trivial. Premise (1) is true so long as the dCPT theorist accepts the priority and specificity conditions. Step (5) follows from the lack of identity from learned to novel content, (6) from that claim plus a strong version of specificity, and then (7) is a simple deduction. This section will focus on premise (2). I will then consider in the next section whether a dCPT theorist might attempt to challenge (1), or the step from (5) to (6), by abandoning or weakening the conditions.

Premise (2) is intended to drive a wedge between the enabling thesis and specificity, and show that the two cannot be maintained together. The idea is simply that any instruction- or belief-based learning must leverage our extant abilities in the service of generating new abilities. So, if contents of beliefs are to enable learning, then they must name and enlist already extant perceptual abilities—my telling you that Smurfs have red hats, big shoes, beards, etc., will avail you not at all if you can't already perceptually recognize those features. But if prior beliefs must name features that a subject can already perceive, and the content they learn is distinct from that content, then the content of the beliefs cannot be specific to the content that is learned. Let's take novel dimensions first, and then modifications to already represented dimensions.

Consider the category of A faces from Gureckis and Goldstone's study. There are a number of features of these faces that subjects can already perceive—noses, ears, eyes, etc. But naming any of these features, or even a conjunction of them, is not the same as naming the dimensions that subjects actually learn, because these dimensions are higher-order ones that capture the configural and correlational structure of the space. One cannot describe these dimensions in terms of simple feature descriptions, of the type that subjects are likely to already be able to perceive. Indeed, it is hard to describe them in simple terms at all, as evident by the poor job that subjects do in describing the dimensions they've learned. Thus, descriptivist beliefs that might actually help learning aren't going to do so by describing the novel contents that perceivers learn.

Similarly, there is evidence that learned perceptual categories *outstrip* descriptive beliefs. Brooks & Hannah (2006) had subjects learn to recognize a set of cartoon creatures on the basis of a description. They then had them perform a transfer task on objects that *equally met the description*, but varied in their overall similarity with the training set. Subjects performed better when the similarity was high, which showed that their learning outstripped the descriptions they had been given.

This dynamic is seen even more clearly in the Goldstone case which distinguished brightness from saturation. Given that subjects cannot perceive these dimensions independently, prior to training, simply telling them, for instance, that category "A" squares are distinguished by their saturation, cannot help them learn how to discriminate the squares. What should they look for to see the difference in saturation? Given that saturation, for them, is bound up perceptually with the orthogonally varying brightness, the instruction doesn't help. However, with the demonstrative feedback and training over exemplars, they can learn to differentiate this dimension. The process doesn't *require*, and indeed would not be *helped by*, descriptions of the category-relevant features.

This leaves us with the best-case scenario for the dCPT theorist, on which descriptions name already-perceivable features, *and* these features are definitive of the category. So, in the brightness versus size case from Goldstone (1994), subjects could easily be told that category As are bigger and Bs smaller, or that As are bright while Bs dark, etc. But, as I suggested at the time, this seeming content specificity is misleading. The behaviors that are *novel* are the ability for increased discrimination along these dimensions. But what the descriptivist beliefs describe—i.e. to sort by size or brightness, is an ability that the subjects *already had* before that learning. So, while what is learned is semantically consistent with the instruction, the instruction doesn't tell perception *how* to represent the stimulus—the novel ability arises due to the repeated interaction between perception and the stimulus (cf., Ransom, 2020). Hence, the content of the descriptivist belief in fact "runs out" before perceptual learning takes over. (I mean this semantically, not temporally. It is likely that we continue to use our beliefs to sort during learning.)

Again, there are empirical cases in which this exact dynamic plays out. Sowden et al. (Sowden, Davies, & Roling, 2000) had inexperienced subjects study radiography images, with the instruction that abnormalities in these images show up as dots. It is well-established that expert radiographers have more fine-grained perceptual sensitivities than novices in these kinds of stimuli. After training with the images, subjects in fact showed increased sensitivity—they could perceive dots at lower levels of contrast than they had before. However, *failure of transfer* shows that this ability clearly outstripped the descriptive belief about dots. Subjects who were trained on positive contrast (brighter than background) dots did not improve on discriminating negative contrast (darker than background) dots, and vice-versa. But this is just to say that the ability they learn is not specified by the beliefs they had, since the same belief resulted in *distinct* abilities (positive versus negative contrast sensitivity) depending on the stimulus. Again, the content of the descriptivist belief named an already extant perceptual ability (recognizing dots), and it was perceptual engagement with the training set that actually produced the learning effect. These kinds of effects have been posited to be relevant to expertise in general (Brooks & Hannah, 2006).

A last, and famous, example is that of chicken sexing. Biederman and Shiffrar (1987) showed that one could short-circuit the extensive exemplar training usually required by chicken sexers by (i) showing subjects *where* to look for the "genital bulb," which is the distinguishing feature of males and females, and (ii) telling them that male genital bulbs were *convex* and female ones *concave*. The fact that performance improves almost immediately has been taken as a way of arguing that no perceptual learning *at all* occurs in this case (Pylyshyn, 2003). Indeed, as Biederman and Shiffrar note, it is the fact that the visual system is *already* well-attuned to convexity and concavity that allows this immediate improvement. What is generally glossed over in discussion of this case, however, is that performance improved, but not fully to the level of experts. This is because there are range of specific concave or convex shapes that experts can discriminate. Rather than suggesting finer-grained descriptions, however, Biederman and Shiffrar suggest that the instructions would have to be *combined* with extensive training on exemplars. That is, the content of the description, and its aid in learning, is exhausted by describing features subjects can already perceive.

So, I suggest that the enabling role posited for descriptivist beliefs by the dCPT is in fact incompatible with their meeting the specificity condition. And if so, then the dCPT misdescribes the learning process—the better account is one that restricts beliefs to causal precursors. I wish to emphasize that the arguments in the last two sections have been about the possible roles of descriptivist belief contents, and therefore don't rely on whether the beliefs are explicit or tacit. In the next section, I will consider several objections. First, I will consider whether a more sophisticated view of tacit knowledge could save the dCPT here. Then I will consider whether there is a substantive version of the dCPT that could weaken or abandon the priority or specificity conditions.

4. Objections

4.1. Sophisticated Tacit Belief

One might complain that I have oversimplified the contents of the beliefs at play here, by taking them as analogous to simple linguistic descriptions. There is an informal and a formal way to cash out this objection. The informal way involves noting that, at some level, there *is* a match in content between the prior belief and the resulting perceptual state – both represent the category 'gnomes'. One might suggest that simply by suggesting that the set of objects can be grouped together, the specificity condition is met. Or, one might say that there is *more* content to the belief than simply to label a set of objects as 'gnomes'. Perhaps the 'gnome' content carries with it a range of deeper connotations that perception can use in learning the category.

These informal responses fail because, along the lines given in section 2, they fail to explain how the belief's content could instruct perception how to represent the category, and it is just this kind of informational relation that is posited in the enabling claim. The first version gives no account of how the grouping label informs the specific content that constitutes the perceptual category – that is, the morphspace. The second version, which posits a richer content to the 'gnome' belief than the ones I've listed, is obscure. One would have to theorize about what these richer connotations might be, and, if what I have said so far is correct, they cannot consist in any of essentialist, demonstrative, or descriptivist beliefs.

The more formal way of pushing the objection would appeal to Bayesian and other hierarchical generative approaches to perception to push against my rejection of tacit knowledge in section 3. While proponents of such views don't agree on their upshot for cognitive permeation (Brössel, 2017; Hohwy, 2013; Stokes & Vance, 2017), they do suggest that both perception and perceptual learning are kinds of abductive inferences that take top-down information into account. Hence, someone might be tempted to claim that, rather than the lexicalized beliefs I've been discussing, diachronic permeation comes about due to the role of top-down knowledge in model-based learning.

Still, however, appealing to tacit knowledge of Bayesian priors in support of the dCPT runs into problems specifying the content of the tacit beliefs that are supposed to permeate perception. Consider two possibilities. First, the beliefs involved in perceptual learning might be descriptive beliefs with probabilities attached to them (although these kinds of contents have also been attributed to perception itself; see Morrison, 2020). So, one might believe that if an object is a glunk then it is green and round with probability P. This kind of belief pretty clearly will not solve the problem, since the attached probability does not add any *perceptual* content to the belief. If "green and round" is not sufficient to describe novel perceptual contents without the probabilistic modifier, then having the probability attached does not change anything.

Second, the priors might be couched in a representation that *directly describes the feature space*. In Tenenbaum and colleagues' (Yildirim, Kulkarni, Freiwald, & Tanenbaum, 2015) model of face recognition, for instance, the priors are encoded in a feature-space of lower-level features such as face shape, pose, and lighting conditions, and conditional probabilities defined over those parameters constitute a probabilistic representation in the face space. But the problem is now apparent: the feature spaces that define categories in the examples above are not complete until learning has occurred. So, we can't represent the relevant category via prior knowledge of the categorically relevant dimensions of the feature space because, prior to learning, the feature space does not distinguish the relevant categories. And given result that perceptual learning *does* generate novel features and dimensions, the prior knowledge will not describe those dimensions.

It is thus telling that, when Bayesians model perceptual learning, they often combine generative Bayesian models with more bottom-up deep learning ones (Salakhutdinov, Tenenbaum, & Torralba, 2013; Yildirim et al., 2015; cf. Buckner, 2018). According to Tenenbaum and colleagues, this allows for "a bottom-up latent variable recognition pipeline for our generative model" (Yildirim et al., 2015, p. 2). And this is for good reason—Bayesian models are limited by their need to enlist "a priori" (read, already known) variables to describe the domain, whereas bottom-up networks are not. And "committing to the a-priori defined feature representations, instead of learning them from data, can be detrimental" for novel stimuli or tasks (Salakhutdinov et al., 2013, p. 1)

4.2. Weakening Priority?

We can generalize the discussion of Bayesian models above to assess whether a proponent of the dCPT could attempt to abandon or weaken the priority condition. I don't believe that abandonment is an option, since if X enables Y, then it seems obvious that X must precede Y. One might attempt to weaken the condition by positing that beliefs and percepts are developed *in tandem*, for instance by generating new descriptivist beliefs and checking them against the data in an iterative hypothesis-and-test method.

This can't work as a way of defending the dCPT, because the same situation described in the previous sections would arise in terms of the *generation* and *confirmation* of the hypotheses. Suppose that values along some dimension X are definitive of a perceptual category. If one can already perceive X, then one is in a position to both generate and confirm a hypothesis about category membership. But in this case, both the generation and confirmation of the hypothesis are being based on already existent perceptual capabilities, and not on the generation of novel representations.

On the other hand, suppose you can't already perceive X. If you cannot perceive it, according to the morphspace framework, that means you can't differentiate it from other dimensions. So it would not be available to you as a distinct hypothesis from your experience. Now, you might know independently that there is a dimension X, or someone might tell you about it, or you might guess that there is one. In this case, you could form the hypothesis "I wonder if it is X that distinguishes these exemplars," but since you cannot discriminate X, you are in no position to tell if it is really X that determines between examples. The only way you could *perceptually* confirm this hypothesis is by coming to discriminate the dimension. And as suggested above, this learning is not informed by the descriptivist belief.

4.3. Weakening specificity?

One might be tempted to argue that I've foisted too strong a specificity condition on the dCPT. In particular, the step from (5) to (6) in the incompatibility argument seems to imply a *very* strong notion of specificity. One might contend that the dCPT theorist can reject the argument by abandoning specificity (thus rejecting premise (1)) or weakening it (thus denying the move from (5) to (6)).

Proponents of cognitive permeation are often non-committal about how close a semantic relationship, in addition to a causal relationship, is required for an interaction between cognition and perception to count as permeation. Siegel is satisfied with the idea that cognitive and perceptual states might have "close" contents. Stokes (2012, 2013, 2015) has offered a definition of cognitive permeation that doesn't define the notion in terms of content relationships at all, but instead in terms of whether a causal relationship between a cognitive state and a perceptual process is "internal" and "mental." Recently, he has argued (Stokes, 2018) that these kinds of internal connections can be mediated by attention.

So, the question on the table is whether the specificity condition can be weakened or simply abandoned. The main worry about this move is that it risks trivializing the dCPT. We often must have *some* relevant beliefs about a category prior to learning to perceive it—consider the birds-on-a-mountaintop case again. Without something like a specificity condition, all of the essentialist, demonstrative, and denotational beliefs discussed in section 2 will count. If one really wants to posit that my knowledge that "glunks are on the left" permeates my eventual learned perceptual category for glunks, it's hard to legislate against it—it's certainly a logically possible move. However, this kind of influence on perception is neither surprising nor particularly informative about cognitive architecture.

The situation looks worse when we think about more general knowledge of categories. Consider my belief that ostriches are flightless. If I have this belief before I can perceptually discriminate ostriches, it certainly will provide some general constraints on how I come to learn to perceive them. I will only bother, for instance, looking at objects on the ground. (And things can get worse than that; consider "ostriches are objects".) If *these* instances count as permeation, then no one would have disagreed with the dCPT in the first place.

This suggests that some degree of semantic relevance or coherence is required for the dCPT to hold. If one wanted to weaken the specificity condition without abandoning it, one would have to posit some degree of semantic coherence more strict than the cases above, but more permissive than the specificity condition. One could, for instance, suggest that the prior beliefs must describe particular perceptible properties. However, it is highly unlikely that there is a principled way of drawing this distinction. All of the beliefs I've mentioned have *some* upshot for perceiving categories—for instance, the belief that baby birds of a certain type live on mountain tops means something about the kind of perceptual surroundings they're likely to be found in.

One might go further with this, perhaps by saying that descriptivist beliefs can only name lower-level perceptual properties that adhere to the bodies of category members. But the only possible justification for restricting the candidates this way would be an intuition that these contents are *more specific* to the perceptual content that's learned. And, given that weakening specificity is what the dCPT proponent is purportedly trying to do here, that is an odd move at best.

Lastly, consider purely causal construals, and in particular causal construals that posit attentional mediation between cognition and perception. I suggest that the trivialization worry holds in these cases as well. In particular, the worry about semantic specificity is replaced by a worry about causal specificity. Suppose that you are going to sit stock-still while I present you a lineup of objects, and I tell you that the glunks are on the left. You will likely covertly attend to the objects on the left. And this cognitively-mediated process will contribute to your ability, with enough practice, to recognize glunks. But if this is not a sufficiently close relationship in the semantic case, it's hard to see why it is on a purely causal story either. Similarly, many loosely connected beliefs might be causally prior to perceptual learning. The causal version of the dCPT would, similar to the semantic version, have to posit a way of constricting candidate permeators, or risk trivializing the thesis.

On the contrary, the position that I've defended, on which cognitive states are causal precursors to purely perceptual learning, need make no arbitrary distinctions of this sort. There is nothing wrong with causal precursors being more or less specific, and thus focusing us on more or less particular characteristics and more or less constricted sets of objects. Indeed, our doing so might play an important role in perceptual learning.

It just needn't be done via permeation.

5. Conclusion

A correct view of perceptual learning should recognize that we do in fact sometimes have descriptive, demonstrative, and denotational beliefs about objects prior to learning to perceive them. I'm not entirely sure, as I sit here, what broccoli rabe looks like. But I am pretty sure that it's green, that it's the kind of thing I can find at the grocery store, and moreover that there will be a label there to help me fix the demonstrative belief that a particular object is an exemplar of that vegetable. However, rather than implementing cognitive permeation, employing these beliefs *puts us in a position* to acquire certain perceptual abilities, by getting us to focus on the right objects, and thus provides causal precursors to perceptual learning. Recognizing the capabilities of perceptual learning, independent of cognitive influence, shows us that we don't need anything more than these precursors to explain the role of cognition in in generating new perceptual representations.

References

Arstila, V. (2016). Perceptual Learning Explains Two Candidates for Cognitive Permeation. *Erkenntnis*, 1-22.

Beck, J. (2014). Analogue magnitudes, the generality constraint, and nonconceptual thought. *Mind, 123*(492), 1155-1165.

Beck, J. (2018). Marking the perception–cognition boundary: The criterion of stimulus-dependence. *Australasian Journal of Philosophy, 96*(2), 319-334.

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.

Briscoe, R. (2015). Cognitive permeation and the reach of phenomenal content. In *Cognitive Influences on Perception: Implications for Philosophy of Mind, Epistemology, and Philosophy of Action.* Oxford: Oxford University Press.

Brogaard, B., & Chomanski, B. (2015). Cognitive Penetrability and High-Level Properties in Perception: Unrelated Phenomena? *Pacific philosophical quarterly, 96*(4), 469-486.

Brooks, L. R., & Hannah, S. D. (2006). Instantiated features and the use of" rules.". *Journal of Experimental Psychology: General, 135*(2), 133.

Brössel, P. (2017). Rational relations between perception and belief: The case of color. *Review of Philosophy and Psychology, 8*(4), 721-741.

Buckner, C. (2018). Empiricism without magic: transformational abstraction in deep convolutional neural networks. *Synthese, 195*(12), 5339-5372. doi:10.1007/s11229-018-01949-1

Burnston, D. C. (2017a). Cognitive permeation and the cognition–perception interface.

Synthese, 194(9), 3645-3668.
Burnston, D. C. (2017b). Is aesthetic experience evidence for cognitive permeation? *New Ideas in Psychology, 47*, 145-156.
Burnston, D. C. (2020). Fodor on imagistic mental representations. *Rivista internazionale di Filosofia e Psicologia, 11*(1), 71-94.
Burnston, D. C., & Cohen, J. (2015). Perceptual integration, modularity, and cognitive permeation. In *Cognitive Influences on Perception: Implications for Philosophy of Mind, Epistemology, and Philosophy of Action.* Oxford: Oxford University Press.
Cecchi, A. (2014). Cognitive Permeation, Perceptual Learning and Neural Plasticity, *Dialectica, 68*(1), 63-95.
Churchland, P. M. (1988). Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science, 55*(2), 167-187.
Connolly, K. (2014). Perceptual learning and the contents of perception. *Erkenntnis, 79*(6), 1407-1418.
Connolly, K. (2019). *Perceptual learning: the flexibility of the senses*: Oxford University Press.
Cowan, R. (2014). Cognitive Penetrability and Ethical Perception. *Review of Philosophy and Psychology, 6*(4), 665-682. doi:10.1007/s13164-014-0185-4
Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences, 39*, 1-77.
Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science, 12*(6), 499-504.
Folstein, J. R., Gauthier, I., & Palmeri, T. J. (2010). Mere exposure alters category learning of novel objects. *Frontiers in psychology, 1*, 40.
Folstein, J. R., Gauthier, I., & Palmeri, T. J. (2012). How category learning affects object, representations: Not all morph spaces stretch alike. *Journal of Experimental, Learning, Memory & Cognition, 38*(4), 807-820.
Fridland, E. R. (2015). Skill, Nonpropositional Thought, and the Cognitive Penetrability of Perception. *Journal for General Philosophy of Science, 46*(1), 105-120.
Gauthier, I., & Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 28*(2), 431.
Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123*(2), 178.
Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*(1), 69-78. doi:10.1002/wcs.26
Goldstone, R. L., Lippa, Y., & Shiffrin, R. M. (2001). Altering object representations through category learning. *Cognition, 78*(1), 27-43.
Goldstone, R. L., & Steyvers, M. (2001). The sensitization and differentiation of dimensions during category learning. *Journal of Experimental Psychology: General, 130*(1), 116.
Gross, S. (2017). Cognitive permeation and attention. *Frontiers in psychology, 8*, 221.

Gureckis, T. M., & Goldstone, R. L. (2008). *The effect of the internal structure of categories on perception.* Paper presented at the Proceedings of the 30th Annual Conference of the Cognitive Science Society.

Hohwy, J. (2013). *The predictive mind*: Oxford University Press.

Jones, M., & Goldstone, R. L. (2013). The structure of integral dimensions: Contrasting topological and Cartesian representations. *Journal of Experimental Psychology: Human Perception and Performance, 39*(1), 111.

Leslie, S.-J. (2008). Generics: Cognition and acquisition. Philosophical Review, 117(1), 1-47.

Lupyan, G., & Spivey, M. J. (2010). Making the invisible visible: verbal but not visual cues enhance visual detection. *PloS one, 5*(7), e11452. doi:10.1371/journal.pone.0011452

Macpherson, F. (2012). Cognitive permeation of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*(1), 24-62.

Marchi, F. (2017). Attention and cognitive penetrability: The epistemic consequences of attention as a form of metacognitive regulation. *Conscious Cogn, 47*, 48-62. doi:10.1016/j.concog.2016.06.014

Mole, C. (2015). Attention and Cognitive Permeation. In J. Zeimbekis & A. Raftopoulos (Eds.), *The cognitive penetrability of perception: New philosophical perspectives* (pp. 218-238). Oxford: Oxford University Press.

Morrison, J. (2015). Perceptual Structuralism.

Prettyman, A. (2018). Perceptual learning. *Wiley Interdiscip Rev Cogn Sci*, e1489. doi:10.1002/wcs.1489

Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *The Behavioral and brain sciences, 22*(3), 341-365; discussion 366-423.

Pylyshyn, Z. W. (2003). *Seeing and visualizing: It's not what you think*: mit Press.

Raftopoulos, A. (2009). *Cognition and perception: How do psychology and neural science inform philosophy?* : Mit Press.

Raftopoulos, A. (2015). Reframing the Problem of Cognitive Penetrability. In *Philosophy and Cognitive Science II* (pp. 3-20): Springer International Publishing.

Ransom, M. (2020). Expert knowledge by perception. *Philosophy*, *95*(3), 309-335.

Roskies, A. L. (2010). 'That'response doesn't work: Against a demonstrative defense of conceptualism. *Noûs, 44*(1), 112-134.

Salakhutdinov, R., Tenenbaum, J. B., & Torralba, A. (2013). Learning with hierarchical-deep models. *IEEE Trans Pattern Anal Mach Intell, 35*(8), 1958-1971. doi:10.1109/TPAMI.2012.269

Schwitzgebel, E. (2015). Belief. https://plato.stanford.edu/archives/sum2015/entries/belief/

Siegel, S. (2012). Cognitive Penetrability and Perceptual Justification. *Noûs, 46*(2), 201-222.

Siegel, S. (2013). The epistemic impact of the etiology of experience. *Philosophical Studies, 162*(3), 697-722.

Siegel, S. (2016). *The rationality of perception*: Oxford University Press.
Sowden, P. T., Davies, I. R. L., & Roling, P. (2000). Perceptual learning of the detection of features in X-ray images: A functional role for improvements in adults' visual sensitivity? *Journal of Experimental Psychology: Human Perception and Performance, 26*(1), 379-390. doi:10.1037/0096-1523.26.1.379
Stokes, D. (2012). Perceiving and desiring: a new look at the cognitive penetrability of experience. *Philosophical Studies*, 1-16.
Stokes, D. (2013). Cognitive penetrability of perception. *Philosophy Compass, 8*(7), 646-663.
Stokes, D. (2014). Cognitive permeation and the perception of art. *Dialectica, 68*(1), 1-34.
Stokes, D. (2015). Towards a Consequentialist Understanding of Cognitive Permeation. In J. Zeimbekis & A. Raftopoulos (Eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives* (pp. 75-100). Oxford: Oxford University Press.
Stokes, D. (2018). Attention and the cognitive penetrability of perception. *Australasian Journal of Philosophy*, *96*(2), 303-318.
Stokes, D. (2020). On perceptual expertise. *Mind & Language*.
Stokes, D., & Bergeron, V. (2015). Modular architectures and informational encapsulation: A dilemma. *European Journal for Philosophy of Science, 5*(3), 315-338.
Tanaka, J. W., Curran, T., & Sheinberg, D. L. (2005). The training and transfer of real-world perceptual expertise. *Psychological science, 16*(2), 145-151.
Valenti, J. J., & Firestone, C. (2019). Finding the "odd one out": Memory color effects and the logic of appearance. *Cognition, 191*, 103934.
Vance, J., & Stokes, D. (2017). Noise, uncertainty, and interest: Predictive coding and cognitive permeation. *Consciousness and cognition, 47*, 86-98.
Yildirim, I., Kulkarni, T. D., Freiwald, W. A., & Tenenbaum, J. B. (2015). *Efficient and robust analysis-by-synthesis in vision: A computational framework, behavioral tests, and modeling neuronal representations.* Paper presented at the Annual conference of the cognitive science society.
Wu, W. (2013). Visual spatial constancy and modularity: Does intention permeate vision? *Philosophical Studies, 165*(2), 647-669.
Wu, W. (2017). Shaking up the mind's ground floor: the cognitive permeation of visual attention. *The Journal of Philosophy, 114*(1), 5-32.