

Can Basic Perceptual Features Be Learned?

*Forthcoming in *Synthese*. Please cite published version when possible.

Gabriel Siegel

Washington University in St. Louis

Abstract

Perceptual learning is characterized by long-term changes in perception as a result of practice or experience. In this paper, I argue that through perceptual learning we can become newly sensitive to basic perceptual features. First, I provide a novel account of basic perceptual features. Then, I argue that evidence from experience-based plasticity suggests that basic perceptual features can be learned. Lastly, I discuss the common scientific and philosophical view that perceptual learning comes in at least four varieties: differentiation, unitization, attentional weighting and stimulus imprinting (Goldstone 1998, Connolly 2019). Becoming newly sensitive to basic perceptual features, I argue, does not fit into any of these categories. This paper's contribution to the literature is twofold. First, I present a novel view of basic perceptual features which can be used in subsequent theorizing. Second, I show that learning basic perceptual features, since it does not fit into this standardized taxonomy, constitutes an underappreciated form of perceptual learning. This result has important implications for recent discussions in the philosophy of perception and epistemology.

Keywords

perceptual learning; perceptual qualities; experience-based plasticity; perceptual expertise; epistemology of perception

Introduction

Perceptual learning is characterized by long-term changes in perception as a result of practice or experience. For example, through experience we can learn to visually discriminate male and female chickens. We can also learn to discriminate, via olfaction, a Merlot from a Cabernet Sauvignon. Such examples plausibly involve learning to distinguish new arrangements of basic features that perceivers are already sensitive to. But can perceptual learning also involve becoming newly sensitive to basic features? In this paper, I use evidence from experience-based plasticity to argue that, through perceptual learning, we can become newly sensitive to basic perceptual features.

My argument starts with a characterization of basic perceptual features. Few philosophers have discussed what it means for a perceptual feature to be 'basic' or 'fundamental.' Byrne and Hilbert (2008) provide a focused treatment of the issue. The novel account that I develop exhibits important advantages over Byrne and Hilbert's discussion.

After providing this account, I argue that perceptual learning can result in becoming newly sensitive to basic perceptual features. In other words, I argue that basic perceptual features can be

learned. My argument uses empirical research on experience-based plasticity to defend this claim. The upshot of this result is that the standardized taxonomy of perceptual learning is deficient. This taxonomy is: unitization, differentiation, attentional weighting and stimulus imprinting (Goldstone 1998, Connolly 2019). I argue that learning new basic features does not fit into these categories. This result has important implications for recent discussions in the philosophy of perception and epistemology.

The structure of the paper is as follows. First, in Section 1, I provide an account of basic perceptual features. Next, in Section 2, I show that evidence from experience-based plasticity suggests that basic perceptual features can be learned. In Section 3, I outline Goldstone (1998) and Connolly's (2019) standard taxonomy of perceptual learning and argue that becoming newly sensitive to basic perceptual features does not fit into it. To end, in Section 4, I outline philosophical implications of this discussion.

1. Basic Perceptual Features

A perceptual feature is an attribute or quality of individuals to which one's perceptual system is sensitive. When a perceptual system is sensitive to some feature F , the system detects F and can represent F in sub-personal states and/or in perceptual experience. Furthermore, in virtue of being perceptually sensitive to F , subjects can respond to F in certain ways. For example, F becomes available for certain motor and cognitive activities. In perceptual experience, the representation of features conveys to subjects that individuals in the world are a certain way. Individuals can be understood as objects or events. The forthcoming discussion is consistent with a variety of views regarding the ontology of properties, e.g., that they are particulars (tropes) or universals (repeatables).¹

I will define basic perceptual features via two criteria. Taken individually, these two criteria are necessary conditions on a perceptual feature to be basic. Taken together, the conjunction of these two criteria is a sufficient condition for a perceptual feature to be basic. Thus, if a perceptual feature does not meet at least one of the criteria, then it is not basic. Furthermore, if a perceptual feature meets both of the criteria, then it is basic.

First, perceiving complex features *depends* on the capacity to perceive more basic features. By 'depends,' I mean *ontological dependence*. In particular, I adopt Barnes's (2012) account of ontological dependence. For Barnes, an entity x is ontologically dependent on an entity y iff at each moment x exists it is dependent on the existence of y for its own existence. Here we can employ the notion of counterfactual dependency, saying that at each moment y exists, y is counterfactually dependent on x .

Barnes's discussion of ontological dependence is defined with respect to entities, but we can apply the notion just as well to psychological capacities. Consider some perceptual capacity C^1 that is ontologically dependent on C^2 . In this case, at each moment in which the capacity C^2 is

¹ For a discussion of the trope framework, see e.g., Williams (1953), Campbell (1990) and Heil (2021). For a discussion of the universals framework, see e.g., Russell (1912), Armstrong (1993), and for a hybrid framework see Lowe (2005).

present in some psychological system, its thus being instantiated in that system is counterfactually dependent on C^1 . For example, if we performed an *ideal intervention* ('ideal' in a sense that needn't detain us here, see Woodward 2003), and removed C^2 from the system, then C^1 would cease to be instantiated in that system.

Barnes's view can dismiss a related notion of ontological dependence, where a perceptual capacity C^1 depends on C^2 iff C^2 caused, at any point in the past, C^1 . Barnes says that, "rather than merely being counterfactually dependent on the existence of something in its past, the object is dependent at each moment of its existence on the existence of something which exists at that very time" (2012: 880). Analogously, we might think that the presence of C^2 in some perceptual system was paramount to the development of C^1 . But this leaves open the possibility that while C^2 was a necessary condition for the development of C^1 , C^2 could cease to exist in a system while C^1 is still instantiated in that system. Like Barnes's notion of ontological dependence for entities, I do not employ this conception of dependence for perceptual capacities. Rather, C^1 is ontologically dependent on C^2 iff at each moment C^1 is instantiated in a perceptual system, it is counterfactually dependent on C^2 . Thus, C^2 plays a causal *sustaining* role for the existence of C^1 .

To illustrate, consider the feature of squareness. If one's perceptual system could not detect and represent lines and 90° angles, then it could not detect and represent squareness. Here, the psychological capacity to represent squareness is ontologically dependent on the psychological capacities to represent lines and 90° angles. In other words, if we intervened and removed the capacity to represent 90° angles from a visual system, for example, that system would no longer have the capacity to represent squareness. This is the case regardless of whether the capacity to represent squareness depended, in its causal-historical development, on the capacity to represent 90° angles. In this case, the capacity to represent squareness is ontologically dependent on the capacity to represent 90° angles since the former is counterfactually dependent on the latter in a sustaining manner.

Of course, the capacity to represent squareness is also counterfactually dependent, in a sustaining manner, on a variety of capacities that are irrelevant to the present discussion. For example, it also depends on more general perceptual capacities developed in infancy, such as the capacity to discriminate figure from ground. It also depends on properly functioning retinal transduction capacities as well as the heart's capacity to pump blood. However, these capacities are not what we are asking about when defining basic perceptual features. In the present discussion, we are concerned with the ontological dependence of some *featural capacity* with respect to other featural capacities. Roughly, a featural perceptual capacity is a capacity to pick out or refer to some feature of the environment (Schellenberg 2018, O'Callaghan 2019, Burge 2022). When asking whether some feature F is more basic than another feature F^1 , according to the first criterion, we start by asking whether the perceptual system's capacity to pick out F^1 ontologically depends on the capacity to pick out F . Since we are determining basicness relations among perceptual features, we are only concerned with dependence relations among perceptual *featural* capacities.

Certain features will bear the 'more basic' relation to one another and others will not. For example, the capacity to represent squareness in vision depends on the capacity to represent lines

in vision. In this case, visual line attributional capacities are more basic than visual squareness attributional capacities. Thus, the visual feature of being a line is more basic, according to the first criterion, than the feature of being a square. Alternatively, consider some feature F that is picked out via the visual sensory modality and another feature G that is picked out via the auditory sensory modality. For example, F is 90° angles and G is a specific pitch. In this case, F cannot be more basic than G because G is not ontologically dependent on F.²

This view of the more basic relation between perceptual capacities might contrast with certain intuitions about basicness relations among features. For example, certain timbre features in audition might appear more basic than higher-level natural kind features represented in vision, e.g., “dog-body” or “drinkable liquid” (these examples come from Burge 2022: 120-123). However, since the capacity to pick out timbre features does not depend on the capacity to pick out the feature of ‘drinkable liquid,’ the former is not more basic than the latter. The basicness relation as I discuss it here is an ontological dependency relation that simply does not apply in such cases.

Barnes develops the account of ontological dependence into the notion of an “ontologically independent” entity. An entity is ontologically independent iff it is not ontologically dependent. This notion is closely related to the notion of ‘capacity-independence’ that I employ here. A featural capacity exhibits capacity-independence iff it is not ontologically dependent on some other featural capacity. The capacity to represent squareness is not ontologically independent, since it’s dependent on the capacity to represent 90° angles. However, the capacity to represent 90° angles is not ontologically dependent on any other featural capacity. Thus, 90° angles meet the first condition for basic perceptual features.

Thus, according to the first necessary condition, a feature F is more basic than a feature G only if perceiving G ontologically depends on the featural capacity to perceive F. If detecting a feature F does not depend on other featural capacities to detect more basic features, then F meets the first condition for basicness. Call this the *capacity-independence* criterion.

The second necessary condition for basic features is what I will call *featural non-decomposability*. The idea here is that certain perceptual features can be decomposed into more basic features and others cannot. For example, to represent squareness is also to represent that four lines are connected at 90° angles. In other words, to represent the feature of squareness is to *simultaneously represent* lines and 90° angles. However, to merely represent that something is a line is not also to represent another set of features. In this way, the feature of being a line cannot be ‘decomposed’ into more simple featural components. The notion of ‘decomposition’ here needs unpacking. It is not physical decomposition. For example, lines can be physically decomposed into smaller points in space. But these smaller points do not help constitute the *feature* of being a line. When a complex feature is featurally decomposed into more basic featural components, those

² However, there might be multisensory cases where some features F and G are both involved in picking out “novel intermodal features” (see O’Callaghan 2019: Ch 3). For example, some perceptual capacity to perceive a certain taste (the capacity ‘T’) is ontologically dependent on both a gustatory (‘G’) and olfactory (‘O’) capacity for perceptual sensitivity. In such cases, perceiving a taste feature might involve the deployment of a featural capacity T that ontologically depends on featural capacities in distinct sensory modalities, i.e., the capacities G and O.

components often correspond with featural constituents of how we might define the feature. For example, the definition of a square plausibly involves the components of four lines and 90° angles.³

What is important here is that a set of more basic features F^n that featurally compose a higher-level feature F^1 , and thus are simultaneously represented when representing F^1 , are *potential answers* to the question of “what is F^1 ”? For example, in asking “what is squareness?” we are not asking about color, shading or figure-ground relations. These cannot be implicated as potential answers to the question of what defines squareness. On the other hand, four lines being connected at 90° angles are appropriate answers.

We can thus talk about more basic features, according to the second criterion, insofar as some feature is a building block in composing another feature, and basic features, insofar as some feature cannot be featurally decomposed into more basic featural components. For example, while squareness can be featurally decomposed into the features of lines and 90° angles, lines, as I’ve suggested, cannot be decomposed in a similar way. Thus, squareness is not a basic feature according to featural non-decomposability. On the other hand, 90° angles and lines are good candidates to meet this condition.

In what follows, I assume that basic perceptual features meet these two conditions: capacity-independence and featural non-decomposability. These are necessary conditions on a perceptual feature to be basic. Taken together, their conjunct is a sufficient condition on a perceptual feature to be basic. Thus, if a perceptual feature F meets both of these criteria, then F is basic. Furthermore, if F does not meet at least one of these two criteria, then F is not basic.

There is the possibility that certain features will meet one condition but not the other. While such features would not be ‘basic’ on the proposed view, they would fit into distinct categories of perceptual feature types regarding capacity-independence and featural non-decomposability respectively. This possibility exists because these two conditions can, in principle, come apart. For example, consider two distinct perceptual representations of distinct features that are processed in different brain areas. Also consider that a series of computations over these two distinct feature representations outputs a representation of an entirely *distinct* feature. We can assume that, due to these computational processes, the two features represented in earlier stages of processing are not simultaneously represented in the new feature. Thus, the new featural representation does not

³ As one reviewer suggests, a series of points in space oriented in a particular way can be a way to define the feature of ‘lineness.’ However, this fact does not necessarily threaten a line’s status as being a basic perceptual feature. While a set of points that physically decompose a line, and those points being oriented in a particular way, can be used to define lineness, this does not entail featural decomposability since featural decomposability requires simultaneous representation. In the visual representation of lineness, we likely do not also perceptually represent the attribution of a large set of orientational features to each physical point along the line. While we might necessarily perceptually represent each point along the line, we do not necessarily represent an *orientational feature* attributed to most of these points (which would plausibly be needed to define lineness). Again, we are looking for featural constituents of representations as candidates for featural decomposition. Perceptually representing all of these orientational features would pose unnecessary demands on visual working memory. Rather, it is much more likely that the visual system represents the feature of lineness as a single featural unit. If this empirical claim turns out to be false, then perhaps lineness is not the best example of a basic perceptual feature (90° angles might be a better one). Since I reject below that the set of ‘low-level’ features is equivalent to the set of ‘basic’ features, this result is acceptable even if it goes against common intuitions about what constitutes basic or primitive features.

simultaneously represent these two features. In this case, this new feature would fail to meet capacity-independence but could meet featural non-decomposability. This illustrates that these two conditions can, in principle, come apart. I leave it as an open empirical question whether there are perceptual features that meet one condition but not the other.

Byrne and Hilbert (2008) provide a focused treatment of basic perceptual features, or as they call them, “basic sensible qualities.” However, they provide a *phenomenological* account of basic perceptual features. In other words, they are exclusively concerned with basic perceptual features that are constituents of phenomenally conscious perceptual episodes. Phenomenally conscious episodes are episodes where there is ‘something it’s like’ to be the subject of those episodes. While my account can apply to phenomenally conscious basic features, it is more broad. My view is concerned with a broader category of features of the environment, i.e., those features to which we are perceptually sensitive. We might represent those features in unconscious perceptual processes and/or in perceptual phenomenal consciousness. So, at the start, the range of our explanandum phenomenon are, at least in principle, different. Nevertheless, my account is broad enough to capture basic features in both perceptual phenomenal consciousness and unconscious perceptual states.

Byrne and Hilbert also defend a conception of basic features that appeals to a notion of non-decomposability. However, they don’t explicitly appeal to *featural* non-decomposability. In the case of taste perception, the authors suggest that to perceptually represent some taste is to represent a conjunction of four of the basic taste features: bitterness, sweetness, sourness and saltiness. For example, some taste might be composed of bitterness and sweetness, but not sourness nor saltiness. These basic taste constituents cannot likewise be broken down into representations of further conjuncts. Rather, they are fundamental non-decomposable constituents of taste perception’s contents.

The authors suggest that while this is a plausible view of basic features in taste, this kind of predicational structure is not the way to think about basic features in color. For color, they understand basic perceptual features as basic hue magnitudes. For example, the color purple might be represented as the total color hue being 60% bluish and 40% reddish (2008: 398). These different hue magnitudes, which might come in four basic hue types, e.g., red, yellow, blue and green, are the basic features implicated in color perception. Importantly, these color hues, unlike basic taste features, are represented as being present in certain degrees.

Thus, Byrne and Hilbert’s view makes use of a notion of decomposability. Tastes contents can be decomposed into basic taste constituents and color contents can be decomposed into basic color hue magnitudes. It is unclear what the necessary and sufficient conditions are for basic features on Byrne and Hilbert’s account. For this reason, it is unclear to what extent their view overlaps with the one developed here. However, on a plausible reading, my account functions as a way to capture how both examples constitute basic perceptual features. For example, take their analysis of taste perception. The capacity to perceive certain tastes appears ontologically dependent on capacities to perceive the aforementioned basic tastes such as sourness and saltiness. Likewise, the capacity to perceive certain colors requires the capacities to be sensitive to various

color hues in various magnitudes. Furthermore, it is plausible that the capacities to perceive basic tastes and basic color hue magnitudes don't ontologically depend on other featural capacities. For present purposes, I just stipulate that this claim, which is an empirical one, is plausible. Thus, at first glance, the examples of basic features Byrne and Hilbert discuss for taste and color can be captured by capacity-independence. Furthermore, we might think that basic tastes and color hue magnitudes are possible answers to the question of what defines some particular taste or some particular color. For this reason, it is also plausible that their examples are captured by featural non-decomposability.

A limitation of Byrne and Hilbert's view is that their account is not applicable to different features across distinct sense modalities. They provide a specific characterization for taste and another distinct account for color. On the other hand, I've shown that my account captures both of these cases. In addition, my account can capture basicness relations among features across a wide variety of sense modalities, as well as features represented in multisensory contexts. It thus exhibits greater explanatory power than Byrne and Hilbert's discussion, while using a single framework.

To further illustrate this point, consider how my account of basic features might be applied to smell perception. It has been shown that, while perceiving an odor stimulus, perceivers can only discriminate 3-4 components of that stimulus (Livermore and Laing 1996). For example, in perceiving a chocolate odor, the perception of this odor can only be broken down into roughly 3-4 discriminable components.⁴ On my framework, these results show that a chocolate odor feature itself is not basic, since the olfactory capacity to pick out that chocolate feature is dependent on distinct featural capacities to pick out these other 3-4 odor components. Furthermore, in representing the chocolate odor, it might be said that we simultaneously represent those 3-4 components. Thus, chocolate odor features appear featurally decomposable. In this way, a chocolate odor feature stimulus is not itself a basic feature. However, the 3-4 discriminable components of odor stimuli are good candidates for basic features.

To sum up, I suggest that basic perceptual features meet both the capacity-independence and featural non-decomposability criteria. These features are represented either sub-personally and/or in perceptual experience. Furthermore, I've shown that my account can determine which features are basic in the context of visual shape, visual color, taste and smell perception. This result demonstrates the widespread applicability of the proposed account, in contrast to Byrne and Hilbert's discussion. My hope is that this new view can play a role in subsequent theorizing in philosophy and psychology.

Before moving to the next section, it is worth investigating the extent to which this novel view of basic perceptual features accords with empirical discussions of early visual processing. In other words, what, if any, are the differences between 'basic' perceptual features and 'low-level' features processed at early stages? I will propose that many such 'low-level' features likely meet the conditions for basicness outlined above, but some do not. Marr's (1982) seminal view of visual processing, which continues to be the basis for contemporary computational models (e.g., see Zhu and Wu 2023), included a discussion of 'low-level' features. Marr proposed that early stages of

⁴ I am grateful to Clare Batty for bringing this case to my attention.

visual processing started with what he called a *primal sketch* which is a primitive description of the intensity changes of an image on the retina. This includes information about, among other things, the image's position, edge-shadings, size, blobs, contrast and orientations. The low-level components, or in Marr's language "primitives," of the primal sketch are not yet equivalent to basic features as discussed above. Basic features are detectable features of *worldly individuals* to which the perceptual system is sensitive. Many of Marr's components of the primal sketch are not features of worldly entities, but only features of the *retinal image* (which may or may not also represent actual features of the environment).

On the basis of the primal sketch, another representation called the *2 1/2-d sketch* is constructed via a series of computational processes. According to Marr, it is only with the *2 1/2-d sketch* that the visual system begins to represent actual physical features of the environment. The *2 1/2-d sketch* represents the orientation and depth of visual surfaces, in addition to certain refined components of the primal sketch. After a series of computations over the *2 1/2-d sketch* representation, the *3-d model* representation is constructed. Based on this 3-d model, an object's shape can be estimated by comparing the model's components with a stored catalogue.

There are cases in which the relationship between the components of a *2 1/2-d sketch* bear the appropriate basicness relations with respect to a 3-d model. For example, as discussed above, a square is a complex feature that depends on picking out lines and is, in part, featurally decomposable into lines. Likewise, computing an object's 3-d model (which later determines its shape) depends on representing lines in the *2 1/2-d sketch*. In this instance, the relationship between a 'higher-level' feature representation at later stages of visual processing is both ontologically dependent on capacities to pick out, and featurally decomposable into, certain 'low-level' features represented at earlier stages. However, given that certain features of early visual representations are not of worldly individuals, but of retinal images, e.g., the image's 'greyiness' and 'blobiness,' this shows that being a low-level feature at the earliest stages of processing does not entail that it is basic. Furthermore, while certain 'low-level' features can be important for constructing more complex representations downstream, many such features cannot themselves be responded to behaviorally or cognitively. Since we are interested in which basic features we are perceptually sensitive to, which requires such response capacities, this provides another reason for why certain low-level features will not be basic in the sense that motivates the present discussion.

This same point applies to contemporary discussions of visual processing, which largely follow Marr's idea of a hierarchical structure of perceptual processing.⁵ These discussions often denote 'low,' 'mid,' and 'high' levels of processing. Low-level processing is associated with areas in the brain that process visual information first, such as V1 and V2, and mid-level and high-level processing with areas that process information further downstream. To describe some paradigmatic examples, low-level processing often involves representing features such as orientation, contrast and color. Mid-level processing often involves representing features such as texture and shape and high-level processing involves representing categorical features such as an

⁵ Although, many contemporary models indicate that the hierarchy is much more complicated, with feedback occurring between certain regions as well several areas processing information in parallel.

object being a face or a human body (for a recent review, see Groen et al. 2017). Again, the set of ‘low-level’ features is not equivalent to the set of ‘basic’ features. For example, color is a paradigmatic low-level feature in these discussions but, as discussed above, is not basic. The reason is that perceptual sensitivity to color ontologically depends on capacities to pick out, and is featurally decomposable into, hue magnitudes.

In short, the set of low-level features at early stages of visual processing, as often discussed in the empirical literature, is not equivalent to the set of basic perceptual features. The proposed view of basic perceptual features allows us to focus on what features of the world are the basic featural units of sensitivity by the perceptual system. By looking at capacity-independence and featural-decomposability, as opposed to the different hierarchical stages at which features are processed, we can focus our attention on this question specifically.

2. Experience-Based Plasticity and Learning Basic Features

In this section, I show that evidence for experience-based plasticity suggests that perceptual learning can involve becoming newly sensitive to basic perceptual features. First, I outline what experience-based plasticity is and how researchers have provided evidence for it. Then I show that, using the case of orientational features, experience-based plasticity suggests that basic perceptual features can be learned.

Before proceeding, it will be helpful to provide an explicit definition of perceptual feature learning. When some perceptual feature *F* is learned for some organism *O*, this means that *O* has gone from never having perceptual sensitivity to *F* to having perceptual sensitivity to *F*. Thus, when *O* *perceptually learns* *F*, *O* becomes newly able to detect *F* and, in virtue of that detection, can respond to *F* in various ways.

It has been shown that particular visual cortex neurons respond to particular orientational features, e.g., bars that are vertical as opposed to horizontal (Hubel and Wiesel 1965, Haynes and Rees 2005). In *experienced-based plasticity*, perceptual experience alters how these neurons respond to stimuli from the environment. The idea is that, as the visual system is exposed to different physical environments, the neurons responding to certain orientations remain active, while those responding to other orientations do not (Karmarkar and Dan 2006). For this reason, altering the prevalence of some feature in one’s physical environment will alter how responsive visual cortex neurons will be to that feature. Neural *responsiveness* to some feature *F* can be defined as when a neuron (at least sometimes) fires when sensory systems are presented with *F*. Relatedly, experience-based plasticity can result in certain neurons having strong ‘preference’ for the stimuli to which they are exposed. Neural *preference* can be understood as a neuron being more likely to fire when presented with *F* over other features or firing when presented with *F* at a significantly greater strength than with others. Certain neurons could stay responsive to *F*, i.e., sometimes fire when presented with *F*, but lack preference for *F*. Responsiveness and preference are ways in which neurons can be ‘tuned’ to particular features. Altering the responsiveness or preference of neurons to specific features, where those neurons are involved in processing visual information at early stages, can alter one’s perceptual sensitivity to that feature. If enough neurons

do not respond to or prefer some feature, then the perceptual system will not go on to represent that feature downstream such that it can be responded to in behavior or cognition. In short, alterations to neural responsiveness or preference are mechanisms by which perceptual sensitivity to features can be altered.

In a classic study, Blakemore and Cooper (1970) demonstrated experience-based plasticity by putting kittens, from 2 weeks to 5 months old, into an enclosed tube with stripes 5 hours a day. Through this design, kittens were only presented with either horizontal or vertical lines (stripes in the tubes) in their visual environment for their first five months of development. Cell recordings in the visual cortex later revealed selective tuning of neurons to the orientations the kittens were exposed to in the tubes. Kittens in the horizontal tube condition demonstrated visual cortex neurons responsive to stimuli with horizontal orientations. Furthermore, neurons in the visual cortex appeared *not responsive* to vertical stimuli. Kittens in the vertical tube condition showed the inverse effect. Stryker et al. (1978) performed a similar study with more stringent deprivation techniques. In their design, they had similarly aged kittens wear goggles that only presented one type of orientation. Their study provided corroborating results.

In a more recent study, Sengpiel et al. (1999) raised kittens in single orientation striped cylinders. After 1-2 months of age, they used optical imaging (considered a more reliable technique than single-cell recordings) to examine the kitten's visual cortexes. This imaging revealed a significant neural preference for the orientation exposed to in the striped cylinders. This effect of so-called "stripe rearing" was also demonstrated by Kreille et al. (2011). They used skull-mounted goggles onto recently born mice that only allowed one orientation to be present. Unlike the previously mentioned studies, they allowed mice to briefly mature in a natural visual environment and mounted the goggles three weeks after birth. Furthermore, they used a newer and more precise imaging technique called "two-photon calcium imaging."⁶

After imaging, the visual cortex neurons again showed significant preference for the orientations to which the mice were exposed (as opposed to the orthogonal orientation). They found that responsiveness dropped significantly in the upper layer of the visual cortex but remained relatively unchanged in lower layers. The authors reported that while neurons in the lower layer stayed somewhat responsive to the deprived orientations, they demonstrated a substantial shift in preference to particular orientations. This result shows that, even three weeks after birth, what orientations are exposed to by the visual system can alter both neuronal responsiveness and, for the neurons where responsiveness is intact, neural preference. For these reasons, this study further demonstrates that exposure to different visual environments can alter perceptual sensitivity to orientational features.

While confirming behavioral evidence is unavailable, it is reasonable to infer that many of the organisms in these studies failed to develop, or would have failed to develop if they underwent extended deprivation, perceptual sensitivity to the orientational features to which they were not exposed. Again, substantial loss of neuronal responsiveness or preference to particular orientations

⁶In their use of this technique, they injected Calcium sensitive dye into the mice's visual cortexes and then, after exciting the dye using light, they took images of the visual cortex using a two-photon microscope.

can lead to a loss of perceptual sensitivity. A certain threshold of responsiveness and preference to some orientation is needed in order for that feature to be processed at later stages and responded to in various ways. Thus, differential responsiveness and preference are neuronal tuning mechanisms by which perceptual sensitivity to some feature can fail to develop, be lost, be regained or be developed for the first time.

A similar effect, where exposure to certain stimuli alters perceptual sensitivity to those stimuli, has been shown in humans. For example, Mitchell and Wilkinson (1974) studied subjects with severe uncorrected astigmatism and showed that they could not see horizontal lines as clearly as vertical lines. The authors speculated that the differences in visual acuity for horizontal versus vertical was a result of distorted visual input due to the astigmatism. The astigmatism deprived the subjects from receiving clear images of horizontal stimuli. Like the kittens, this is a case where what one is exposed to in the environment alters how one is perceptually sensitive to certain stimuli. Mitchell et al. (1973) showed that these effects of reduced discriminatory capacities for certain orientational stimuli has a neural basis and that the reduction of acuity to the orientational stimuli correlated with the severity of the astigmatism.

Another relevant result of experience-based plasticity is called the *oblique effect*. This is the effect that visual systems are more responsive to vertical and horizontal lines than slanted ones. To put it crudely, people cannot see slanted lines as well as they can see vertical and horizontal lines. Like the explanation for the kittens, the idea is that vertical and horizontal lines are more prevalent in our visual environment than slanted ones. Coppola et al. (1998) demonstrated the oblique effect in ferrets and Furmanksi and Engel (2000) showed an analogous effect in humans. The latter used fMRI to show that the primary visual cortex exhibits significantly greater responses to horizontal and vertical stimuli than oblique stimuli. It has been shown that the oblique effect grows in strength from childhood to adulthood (Mayer 1977).

Evidence has shown that experience-based plasticity is not limited to early developmental stages (Gauthier et al., 1999, Bukach et al., 2006, May 2011, Carcea and Froemke 2013). On this topic May notes that, “contrary to assumptions that changes in brain networks are possible only during crucial periods of development, research in the past decade has supported the idea of a permanently plastic brain. Novel experience, altered afferent input due to environmental changes and learning new skills are now recognized as modulators of brain function and underlying neuroanatomic circuitry” (2011: 475). Perhaps even more relevant to the present discussion, Carcea and Froemke (2013), after reviewing different mechanisms of neuromodulation relevant to perceptual performance, say that, “plasticity in the adult sensory cortex is often transient...neurons return to their original tuning in time and, as a consequence, sensory maps recover their initial representation” (2013: 77). These latter authors suggest that, due to experience-based plasticity, it is possible for a neural tuning map structure, i.e., where particular regions of the brain are tuned to particular features, to be lost and subsequently recovered. This means that for some perceptual feature F , experience-based plasticity can result in neural tuning to F at time t^1 , then losing that tuning to F at t^2 , then regaining tuning to F at t^3 . Importantly, neuroscientific evidence suggests that such alterations can occur in adults (i.e., beyond the ‘critical stages’ of early development).

Examples of this sort of phenomenon are widespread. For example, Tolias et al. (2005) demonstrated that V4 neurons in adult macaques are not initially tuned to stimuli involving the direction of motion. Nevertheless, they showed that through adaptation via repeated exposure, these neurons became selectively tuned to direction of motion. When the monkeys were briefly shown moving stimuli, V4 neurons were not initially selective to its different directions of movement. However, in another condition, if they showed an adapting stimulus, where dots were moving upward for one second before the test moving stimulus, V4 neurons became sensitive to different directions of movement.

Froemke et al. (2013) looked at neuromodulation in the auditory cortex of rats. By exposing the rat auditory system to different tone stimuli, they showed that neurons initially preferring high-intensity tones could come to prefer low-intensity tones. As the authors noted, this resulted in increased perception of low-intensity tones as demonstrated by behavioral performance. The authors claimed that these mechanisms of alterations to neuronal tuning can “increase detection...[of] *previously imperceptible stimuli*” (my italics, 2013: 79). These studies suggest that experience-based plasticity, including subsequent alterations to perceptual sensitivity, extends beyond early stages of visual development.

Recent empirical work has indicated that experience-based plasticity can alter perceptual sensitivity to orientational features after early developmental stages. Since orientational features are the principal example of learning basic features to which I will turn shortly, these studies are worth mentioning. Schoups et al. (2001) demonstrated that training of monkeys on an orientation discrimination task altered the tuning properties of primary visual cortex neurons tuned to the orientations targeted during the training. In humans, Jehee et al. (2012) trained participants over several weeks to discriminate between minor changes in the orientation of sinusoidal gratings (multiple lines of different spatial frequencies within a shape). The training was shown to increase the responsiveness of neurons in the primary visual cortex to the orientations involved in the training. The authors note that their results, “suggest that the functional plasticity of early visual areas is important for realizing the benefits of extended perceptual training...The improved reliability of orientation-selective activity found here may reflect the enhanced gain or sharpening of orientation-tuned responses” (2012: 16752). Thus, experience-based plasticity in human adulthood, resulting in differential sensitivity to perceptual features, is found to be present for orientational features.⁷

Due to clear ethical limitations, stripe rearing on humans, or other types of extreme deprivation from birth to certain perceptual features, has not been conducted.⁸ However, given the principles of experience-based plasticity discussed above, we can make empirically-informed inferences about what would happen were humans to grow up in an environment completely devoid of specific orientational features. Differential responsiveness and differential preference for specific orientations in the adult human visual cortex has been shown to be extremely plastic

⁷ For even more evidence that the response properties of orientation selective neurons can be shaped by differential visual inputs, see Tagawa et al. (2005) and Hofer et al. (2006).

⁸ Although, in their experiment on opossums, Dooley et al. (2017) showed that this kind of stripe rearing effect on the visual cortex is present in our earliest mammalian ancestors.

and alterable given differential exposure to visual stimuli. We can thus consider an extreme case of sensory deprivation to some feature, where a perceiver is *never* exposed to that feature, and where this results in major to complete loss of neural responsiveness or preference for that feature. The empirical evidence cited above shows that this would likely result from such deprivation. Again, if a certain number of neurons are not tuned to certain featural stimuli then the perceptual system will not be able to process information about that feature such that the organism can respond to it. Thus, while we cannot rely on specific empirical studies to directly show that feature deprivation can lead to a loss of perceptual sensitivity in humans, which can subsequently be regained, we can argue for its empirical plausibility given the evidence currently available. This empirically-informed inference can add to our theoretical understanding of perceptual learning, while moving beyond the limitations of currently available empirical methods.

I now provide an argument that basic perceptual features can be learned. In other words, I argue that it is possible, due to alterations to neural tuning to some basic feature F, that a creature can go from never being perceptually sensitive to F to becoming newly perceptually sensitive to F. To make the point clear, consider the following scenario. Susan grows up in an environment without vertical or horizontal lines. This physical environment only has slanted lines. Due to the principles of neuronal tuning, her visual system becomes exclusively tuned to these slanted orientations. At this point, vertical and horizontal orientations are not something to which her perceptual system is sensitive. Again, as I understand perceptual sensitivity, this means that Susan, at this moment, cannot respond to these features of the environment behaviorally or cognitively in virtue of perceptual detection. Now consider that Susan moves to Manhattan. Initially, she will not be perceptually sensitive to vertical and horizontal orientations of the skyscrapers in her visual field. However, research on experience-based plasticity suggests that, through experience and time, she will. Her visual cortex will ultimately become tuned to these new orientational features such that perceptual sensitivity can be obtained.⁹

The next step in the argument will be to show that verticality and horizontalness are basic features as defined in Section 1. An orientational feature is a perceptual feature because it represents objects as being a particular way. These features can be construed as *relational features*. When an orientation feature is instantiated in some object, that object is positioned with respect to/directed toward other objects, the ground or gravitational pull, in a particular kind of way. For example, when an object instantiates the feature of verticality, it's longest points along its medial

⁹ I recognize that certain cases of sensory deprivation, in early developmental stages, might permanently inhibit sensitivity to certain features. However, the evidence surveyed above suggests that the case of Susan is empirically plausible. The highly plastic nature of the brain suggests that, at the level of processing in visual cortex areas, perceptual sensitivity can develop for features that were never previously exposed to. It is plausible that, in many purported cases of permanent loss of perceptual sensitivity due to a lack of exposure, that this loss is due to an incapacity for the sensory transduction system to properly develop, e.g., from wearing an eyepatch over one's eye during childhood. In other words, in many cases of permanent sensitivity loss, this loss might be due to stunted development of the sensory transduction system as opposed to perceptual processing areas. In any case, all that is needed to demonstrate the current point is a case where some perceiver at some time t^1 is not perceptually sensitive to some feature F. Then, due to sustained exposure to F, the perceiver becomes newly sensitive to F at another time t^2 . The evidence reviewed above suggests that such a case is possible.

axis structure is often positioned in aligned with gravitational pull on the object. In this case, the object is directed toward and positioned with respect to the ground, and the center of the earth, in a particular way.

The feature of verticality meets both the capacity-independence and featural non-decomposability criteria. First, the capacity to represent verticality does not depend on other *featural* sensitivity capacities. While sensitivity is needed to at least one geometrical pattern, it is not needed for any *specific* pattern. For example, one can represent verticality without being sensitive to lines or squares.¹⁰ Furthermore, it might be supposed that to be sensitive to the feature of verticality requires sensitivity to its determinable feature of orientation. However, it is implausible that the perceptual system exhibits the capacity to pick out orientation without a determinate dimension value. In other words, the perceptual system does not represent that some object has an orientation *simpliciter*.¹¹ Thus, picking out verticality does not ontologically depend on picking out its determinable feature of orientation. For these reasons, orientational features like verticality meet the capacity-independence condition for basic perceptual features. In Barnes's (2012) language, they are picked out via "ontologically independent" perceptual capacities. At each moment the capacity to pick out and represent verticality is instantiated in the visual system, there is no other featural capacity onto which it is counterfactually dependent.

Second, when you represent verticality, you do not also represent simpler orientational features that featurally compose verticality. Verticality is a non-featurally-decomposable fundamental positional feature of objects. To attribute verticality to an object is to attribute a positional relation to the object. Representing geometrical patterns, e.g., lines, shapes, boundaries, etc., are necessary conditions for representing orientations. However, this does not mean that orientational features are decomposable into lines, curves, shapes, etc. Rather, orientational features are basic features *of* geometrical patterns. Hence, they are not *featurally* composed out of them.

When Susan moves to New York, her visual system ultimately becomes sensitive to new orientational features. This means that, due to experience, Susan's perceptual system has learned to detect and respond to new basic features. This capacity had to be developed through repeated exposure to novel orientational stimuli. Given the principles of neuronal tuning, I've shown that this thought experiment, where there is movement from a lack of perceptual sensitivity to verticality, to perceptual sensitivity to verticality, is empirically plausible. In other words, the empirical evidence surveyed earlier motivates that such a case is nomologically possible. This result demonstrates that learning new basic features is both possible and an instance of perceptual learning.

¹⁰ It is unlikely that perceivers will only be sensitive to one orientation. In other words, it is unlikely that some perceiver is perceptually sensitive to verticality but not sensitive to at least some other orientational features. However, the point here is that there is not a dependence relation between the featural capacity to pick out verticality and to pick out any other *specific* orientational feature(s).

¹¹ We can also be skeptical that attributing orientation broadly to some object is really attributing a feature to it. In order for there to be feature attribution, the perceptual system must specify that the object is a particular way. In this vein, it is hard to imagine what perceptually attributing 'bare orientation' would be like.

It might be objected that, in the case of Susan, it is not a new basic *feature* that is learned but a new basic feature *dimensional value*.¹² In other words, Susan only becomes newly sensitive to some of the many values on the dimension of orientation, i.e., verticality and horizontalness. In response, it will be helpful to make an analogy with color perception. Perceivable color features can be modeled in terms of a three-dimensional *mental quality space* where there is an axis for hue, saturation and luminosity (Churchland 2005, Rosenthal 2010). A value on this space corresponds with a distinct color feature to which organisms are perceptually sensitive. It should be relatively uncontroversial that distinct values on this color space are distinct color features, e.g., one value could model the feature of ‘crimson’ and another could model ‘scarlett.’ There is a sense in which crimson and scarlett are different values of the same feature dimension, i.e., the color red. However, there is another sense in which crimson and scarlett are clearly distinct features that can be attributed to objects. Crimson and scarlett being determinates of the determinable red shouldn’t discount them from being conceptualized as distinct color features. Analogously, the feature ‘red’ being a determinate of the determinable feature ‘color’ doesn’t discount redness from being a feature as distinct from other colors.

If we accept this story in the case of color, then we should accept an analogous one for orientation. Like color, different orientations can be modeled in a quality space and different values within that space will correspond with different orientations, e.g., 1 – 359°. Each value within that space can, just like the case of color, be construed as a distinct orientational feature that can be attributed to objects. Such orientational features like horizontalness are the determinates of the determinable feature of having orientation. Thus, Susan, before moving to Manhattan, is sensitive to certain values within the orientation quality space but not yet to others. Once Susan moves to Manhattan, she becomes newly sensitive to orientational values within that space to which she was not previously sensitive. We can construe Susan becoming newly sensitive to these orientational values as her becoming newly sensitive to a set of basic perceptual features.

Finally, it might be thought that learning novel orientational features, in the way described above, is only a case of perceptual *development* or *maturation*, not perceptual learning. However, this would be a mistake. According to Connolly, perceptual maturation involves genetically predisposed developmental trajectories for our perceptual system. For example, this might involve the developmental capacity for infants to separate their visual array into bounded segments and perceive three-dimensional space. The distinguishing feature of perceptual learning, for Connolly, is that perceptual differences result from interactions with the environment (2019: 15). I largely follow Connolly in this characterization. However, there are forms of maturation where an environmental stimulus is needed to ‘trigger’ a genetically pre-wired process to unfold. This triggering importantly involves the environment and is how maturation is generally understood in a variety of cases, e.g., for certain innate linguistic capacities (see Chomsky 1981).

Neuronal tuning to orientational features is the result of a visual system being exposed to such features over time. Thus, learning to represent new orientational features involves interaction with the environment, which is characteristic of perceptual learning. Furthermore, it is also the

¹² For some discussion of feature dimensions, see e.g., Rosenthal (2010) and Green (2020).

case that, in neuronal tuning, the role of the environment is not merely to ‘trigger’ a genetically pre-determined process. Rather, neuronal tuning can *continuously change* perceptual sensitivity across an organism’s lifetime due to *sustained exposure* to different environmental stimuli. Thus, we can dismiss the view that learning novel orientational features is merely a case of perceptual maturation.

I’ve shown that perceptual learning can involve becoming newly sensitive to basic perceptual features. My argument first demonstrated that we can become newly sensitive to orientational features such as verticality and horizontalness in vision. Then, I argued that these orientational features are basic perceptual features given the account developed in Section 1. The conclusion of this argument is that basic perceptual features can be learned. Next, I’ll show that this species of perceptual learning does not fit into the standardized taxonomy.

3. Goldstone and Connolly on Perceptual Learning

According to Goldstone (1998) and Connolly (2019), perceptual learning comes in at least four varieties: unitization, differentiation, attentional weighting and stimulus imprinting. In this section, I argue that learning new basic perceptual features does not fit into any of these categories.

Unitization is when perceivers learn to discriminate a property, object or event as a single individual, where they previously perceived them as multiple (Connolly 2019: 22). This capacity to unitize components of stimuli can be quite helpful in performing certain tasks. As Goldstone says, “via unitization, a task that originally required detection of several parts can be accomplished by detecting a single unit” (1998: 602). For example, research on word perception suggest that words can come to be perceived as single units as opposed to a conjunction of individual letters (O’Hara 1980). Learning to group chunks of letters together into words is crucial for reading and is developed through exposure to the linguistic stimuli. A string of letters might look like nonsense to an English speaker but be seen as a word to a German speaker. Through experience and practice, English speakers can learn to unitize letters into units of German words. This phenomenon is a paradigmatic case of unitization.

Learning new basic perceptual features is not the same as unitization. A salient reason is that learning a new basic feature does not necessitate that the feature was previously perceived to be multiple distinct features. For example, we can become newly sensitive to verticality without having previously perceived verticality as distinct components of a set of orientational features.

In the cases discussed above, its apparent that learning new basic features is a distinct phenomenon from unitization. Unitization implies that one’s perceptual system is sensitive to distinct features, and then learns to ‘chunk’ these distinct features together into a single featural unit. For example, in the case of word perception, this requires pre-established perceptual sensitivity to each letter stimulus. Then, at later stages of processing, each letter stimulus is chunked together into a cohesive word representation. This chunking phenomenon is clearly different from newly developed sensitivity to each single letter stimulus. There is a clear difference between being able to detect and respond to each letter stimulus and fusing those letters together

at later stages of processing. It is precisely this type of distinction that shows that unitization and novel basic feature sensitivity are distinct phenomenon.¹³

Analogously to the letter case, we can distinguish novel sensitivity to basic features, on the one hand, and learning to chunk those basic features together at later processing stages on the other. For example, consider again the case of Susan. Susan develops novel sensitivity to orientational features to which she was not previously sensitive. Susan was sensitive to a range of slanted orientational stimuli, but it was not the case that, upon moving to Manhattan, she started to chunk these distinct slanted orientational stimuli together to represent verticality. That is not a plausible explanation for Susan's new perceptual capacity to pick out verticality. Verticality, as noted above, is a distinct value within the orientation feature space. It is not composed of multiple dimension values. Thus, it would not make sense for novel sensitivity to verticality to be the result of chunking together multiple orientations. For these reasons, we can dismiss the view that unitization is the same as learning new basic features.

Differentiation is when perceivers learn to discriminate among multiple objects, features or events that they previously thought were the same (Connolly 2019: 20). As Goldstone says, "by differentiation, stimuli that were once psychologically fused together become separated. Once separated, discriminations can be made between percepts that were originally indistinguishable" (1998: 596). For example, before learning French, you might be unable to discriminate between the nasal phonemes /ã/ and /õ/. However, after learning French and repeated exposure to the language, you can make the appropriate auditory discriminations. Lively et al. (1993) demonstrated that speakers of Japanese could learn to differentiate between the phonemes /r/ and /l/, which are not present in Japanese. Tanaka and Taylor (1991) showed different human differentiation capacities for distinct dog and bird species, given different levels of prior perceptual exposure to these categories.

It might be thought that becoming newly sensitive to verticality implies that you come to differentiate verticality from other orientational features that you previously perceived as the same. For example, in the case of Susan, perhaps the orientational features of 90° and 75° were previously, using Goldstone's language, "psychologically fused together" at some stage of perceptual processing, but through repeated exposure to verticality, the perceptual system learned to treat them as different. To show that this type of explanation is different from learning new basic features, consider again the distinction between the phonemes /ã/ and /õ/ that a learner of French comes to appreciate. In order for differentiation to occur, one's auditory system needs to be sensitive to both phonemes and at later processing stages the phonemes needed to be fused together. After differentiation, these fused together phonemes are treated as different. However, it might also be that one's auditory system is simply not yet sensitive to one of the phonemes (given that they had never been exposed to it before). In that case, it is not that the perceiver was previously sensitive to both phonemes and then fused them together at later processing stages. Rather, the perceiver could not detect and respond to one of the phonemes to begin with. In short, there are multiple potential reasons for why a beginning student of French cannot discriminate

¹³ This case is meant to illustrate the distinction under consideration. Such letter stimuli are not 'basic.'

between /ã/ and /õ/. First, it might be that they are sensitive to both stimuli but cannot yet differentiate between them. Second, it might be that, due to a lack of sensitivity, information about one phoneme is not processed to begin with. These are importantly distinct explanations for differences in perceptual discriminability.

Analogously, instead of saying that Susan was sensitive to both 90° and 75° orientations, and previously fused them together, we can say that Susan lacked sensitivity to one of these orientational features altogether. In other words, the feature of verticality, before spending sufficient time in Manhattan, was not processed by her visual system to begin with. This should be the preferred explanation for Susan given that she was never exposed to this feature in the slanted world. Her deprivation to verticality and horizontalness is analogous to that of the kittens who were stripe reared in single orientation cylinders. In the studies mentioned in Section 2, kittens stripe reared to one orientation (e.g., vertical lines) exhibited visual cortex neurons that appeared not to fire, or not to fire often, when presented with the unexposed orthogonal orientation (e.g., horizontal lines).

In short, two importantly distinct explanations can be implicated in the phenomes, kittens and Susan cases. These distinct explanations illustrate the difference between learning new basic features and differentiation. Differentiation requires previous sensitivity to multiple features and a previous fusing of those features together. On the other hand, learning new basic features involves starting with no sensitivity to some feature and then becoming newly sensitive to it. These distinct types of explanations show that learning basic features and differentiation are distinct phenomena.

Attentional weighting results from learning to attend to specific individuals, specific features of individuals or specific spatial regions (Connolly 2019: 24). Learning to attend to these different perceptual aspects can allow you to better discriminate individuals or features and become newly sensitive to certain complex features. This result is often demonstrated in studies where certain features of a stimulus will be more important, and others less important, for completing particular discrimination tasks (Goldstone 1998: 588). For example, subjects learn to attend to certain features, even the less salient ones, when those features help discriminate between stimuli that fall into different categories (Livinston and Andrews 1995).

Becoming newly sensitive to basic features is also not captured by attentional weighting. Attentional weighting to distinct basic features presupposes that the perceptual system is already sensitive to those features. In other words, if we learn to attend to a basic feature F, this already assumes perceptual sensitivity to F. For example, the case of Susan is explained above by newly acquired sensitivity to certain orientational features. It's not that Susan begins to attend to a feature in Manhattan that she could have attended to in the past. Since Susan was not perceptually sensitive to verticality, she could not have attended to that feature. Thus, in the case of Susan, learning to perceive verticality didn't involve starting to preferentially attend to features to which Susan was already sensitive. This explanation for Susan learning to perceive verticality is clearly distinct from

one that implicates attentional weighting. These two types of explanations demonstrate that learning new basic features and attentional weighting are distinct phenomena.¹⁴

Stimulus imprinting involves the development of specialized detectors (or ‘receptors’) for particular stimuli or parts of stimuli (Goldstone 1998). Goldstone uses the word ‘imprinting’ since the receptors are “shaped” by the stimulus from the environment. The basic idea is that when visual systems are repeatedly exposed to certain stimuli, whether at the level of the entire visual image, features, or the stimuli’s topographical structure, specialized receptors develop to enhance processing of that stimuli. Of particular importance here is *feature imprinting* where the visual system develops specialized detectors for certain perceptual features. Extended exposure to a featural stimuli, e.g., certain curvatures within an object, can later enhance processing of that feature when the visual system is subsequently exposed to it (Schyns and Rodet 1997). Sometimes stimulus imprinting can happen via neuronal tuning (see Weinberger 1993).

While learning novel basic features and featural stimulus imprinting might, under certain circumstances, both involve neuronal tuning, these are distinct phenomena. In order for stimulus imprinting to occur, the perceptual system must already be sensitive to the feature that is imprinted. Stimulus imprinting involves the development of enhanced processing for some feature, but not new sensitivity. In order to develop stimulus imprinting of some feature F, there must be a previously established capacity for perceptual sensitivity to F. Thus, learning basic features and stimulus imprinting constitute distinct phenomena.

To sum up, perceiving novel basic features does not fit into the categories of unitization, differentiation, attentional weighting or stimulus imprinting. If this is right, then there is an important lesson to be learned. This lesson is that perceptual learning can involve more than coming to distinguish new arrangements of, attend to, or enhance processing of, basic features that perceivers are already sensitive to. In this next section, I outline implications of this result for some recent philosophical discussions.¹⁵

¹⁴ Unitization, differentiation and attentional weighting can all be implicated in what is called acquired category perception. In *acquired category perception*, the perceptual system becomes newly sensitive to certain categorical distinctions which, in turn, modulate perceptual representations. For example, being a speaker of languages with terms for more fine-grained color shades, e.g., of yellow, has been shown to correlate with better discriminatory capacities of yellow shades (Roberson et al. 2005). Categorical perception can involve encoding perceptual stimuli of the same category type as more similar or of stimuli of distinct category types as more dissimilar (Goldstone and Hendrickson 2010, Dubova and Goldstone 2021). These instances of acquired categorical perception are instances of unitization and differentiation respectively. Others have suggested that it can also involve attentional weighting. For example, Connolly (2019) suggests that learning to discriminate different types of trees, e.g., pine trees from fir trees, involves attending to certain low-level features that are distinct between the tree types.

¹⁵ In a recent paper, Jenkin (2023c) critiques the ‘offloading view’ of perceptual learning. According to this view, the function of perceptual learning is to free up cognitive resources by ‘offloading’ cognitive tasks onto perceptual systems (Connolly 2019). In lieu of this view, Jenkin argues in favor of the ‘perceptual view’ where perceptual learning functions to improve perceptual capacities. The argument proceeds by showing that the perceptual view can accommodate studies often cited in support of the offloading view. Furthermore, it can capture others, such as studies involving infants and animals, that the offloading view has trouble accommodating since cognition is underdeveloped or absent. It is plausible that adding learning new basic features to our taxonomy further motivates Jenkin’s perceptual view. If perceptual learning can result in novel perceptual sensitivity to some basic feature, this appears to be an exclusively perceptual phenomenon. In other words, it is hard to see how this form of perceptual learning would have

4. Philosophical Implications

Perceptual learning has notable implications for various debates in the philosophy of perception, as well as epistemology, aesthetics and moral philosophy (see Jenkin 2023a). For example, several theorists have suggested that perceptual learning can alter what contents are represented to perceivers in perceptual experience (e.g., Siegel 2006, 2010, Chudnoff 2021). If basic perceptual features can be learned, and learning these features alters what perceptual contents can be represented, then novel sensitivity to basic features constitutes another way in which perceptual learning alters perceptual experience. Since the differential presence of basic features in the environment will alter perceptual phenomenology, perceivers from different physical environments will perceive the world in different ways. I've shown that this result does not only apply to acquired categorical perception, which has often been the focus of studies involving cross-cultural variation in perception. Rather, this result also applies to differential sensitivity to basic perceptual features present in one's environment.

Another relevant issue is that of perceptual expertise. This topic is often discussed in terms of how perceptual learning enables perceivers to represent certain contents that others cannot. The capacity to represent these contents makes certain individuals "expert perceivers." In particular, theorists have focused on the interaction between perceptual learning and coming to represent new high-level contents like an object being a pine tree (Siegel 2006, 2010). Chudnoff (2021) suggests that through perceptual learning experts perceptually represent "expertise-specific contents." For example, radiologists perceptually represent contents where the property of 'abnormality' is attributed to certain regions of x-ray images (2021: 25-26). The feature of 'abnormality' is a high-level property. This property plausibly involves the combination of a variety of distinct basic features. However, we can also imagine cases of perceptual expertise where such expertise is a function of unique or specialized sensitivity to basic features. For example, consider that Susan returns to her slanted homeland. Among her native slanted-world inhabitants, she would exhibit perceptual expertise in visual orientation sensitivity that other perceivers among her would not (even if this capacity is never deployed in this environment). Furthermore, there might be complex features that she would be uniquely sensitive to as a result of unitization, differentiation and attentional weighting to basic orientational features to which she is uniquely sensitive. Being uniquely sensitive to these basic orientational features thus opens up a variety of further perceptual learning opportunities, and thus opportunities for variation in perceptual expertise.

Finally, consider implications of learning new basic features for the epistemology of perception. Perception is often thought to play a role in justifying beliefs. In particular, this includes beliefs that are based on perceptual experience (e.g., see Chisholm 1977, Pryor 2000, Bonjour 2003, Siegel 2017). If learning basic perceptual features alters perceptual experiences, then it would also alter what beliefs perception is capable of justifying.

Several theorists have discussed how perceptual learning interacts with the epistemology of perception (e.g., Brogaard and Gatzia 2018, Chudnoff 2021, Jenkin 2023a, 2023b). For

the function of 'offloading' a cognitive task onto a perceptual one, given that the process is exclusively concerned with perceptual sensitivity.

example, Chudnoff (2021) defends *presentational conservatism*, the thesis that perceptual experiences justify beliefs in virtue of their *presentational phenomenology* (2021: 101-103). Presentational phenomenology is characterized by a “felt awareness” of the truth-makers for the propositional content of some experience (2021: 106-107). Given presentational conservatism, Chudnoff argues that perceptual experts have perceptual experiences with greater justificatory power than novices. He says that, “in virtue of having presentational phenomenology associated with expertise-specific representational content...experts have a source of immediate justification that novices lack” (2021: 144). An analogous point can be made for perceivers who have sensitivity to basic features that others do not. Within Chudnoff’s framework, differential sensitivity to basic features will alter the presentational phenomenology available to perceivers, which in turn will alter their epistemic capacities for justification.

Another example is provided by Jenkin (2023b) who argues that in certain cases of perceptual learning, perceptual states can be based on reasons and thus are epistemically evaluable. More specifically, perceptual learning can result in new information being stored in the perceptual system. When this information is used to form certain perceptual states, that information, according to Jenkin, can be construed as epistemic reasons onto which perceptual states are based. For example, expert chess players learn to unitize chess pieces in specific locations into “chunks” which indicate possible moves, e.g., castling. According to Jenkin, this perceptual unitization is based on visual inputs and a unitization rule stored in perceptual memory, which should be construed as reasons. Thus, perceptual states that represent these chunks are based on reasons and thus are epistemically evaluable. These states are evaluable as justified or unjustified in virtue of the way in which they respond to reasons.

Jenkin’s discussion of perceptual expertise in chess is one way in which the process of unitization enables agents to have justified perceptual states that non-experts would not. In principle, the same sort of epistemic advantage of perceptual experts can happen for differentiation, attentional weighting and stimulus imprinting. By adding learning novel basic features to our taxonomy, we can emphasize a previously unacknowledged way in which perceivers can have different perceptual states with different epistemic statuses. As noted above, all cases of unitization, as well as the other varieties of perceptual learning found in Goldstone and Connelly’s taxonomy, presuppose sensitivity to basic perceptual features. Thus, differential perceptual sensitivity to basic features implies differential capacities for unitization, differentiation, etc. If cases of unitization enables new epistemic statuses for perceptual states, and such unitization depends on sensitivity to certain basic perceptual features, then differential sensitivity to basic features will necessarily modulate the epistemic capacities of perceivers. Thus, adding learning new basic features to our taxonomy of perceptual learning is also important for understanding the ways in which perceptual learning interacts with the epistemology of perception.

Conclusion

In this paper, I first defended a novel view of basic perceptual features. Second, I argued that perceptual learning can result in becoming newly sensitive to basic perceptual features. By

surveying research on experience-based plasticity, I showed that the mechanism of neuronal tuning likely underpins this kind of learning. Third, I suggested that Goldstone and Connolly's four varieties of perceptual learning cannot capture this phenomenon. To end, I discussed implications of this result for issues in the philosophy of perception and epistemology.

Acknowledgements

I am grateful to Casey O'Callaghan, Clare Batty, Aleksandra Kuciel, Madeleine Ransom, Carl Craver, audience members at the 2023 ISPSM online conference, the 2024 Central APA in New Orleans, the 2024 SSPP in Cincinnati and two anonymous reviewers for helpful comments.

References

- Armstrong, D. M. (1993). A world of states of affairs. *Philosophical Perspectives*, 7, 429-440.
- Barnes, E. (2012). Emergence and fundamentality. *Mind*, 121(484), 873-901.
- Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228(5270), 477-478.
- Bonjour, L. (2003). A version of internalist foundationalism. In L. Bonjour, & E. Sosa, (Eds.), *Epistemic Justification: Internalism vs. Externalism, Foundationalism vs. Virtues* (pp. 3–96). Malden, MA: Blackwell.
- Brogaard, B., & Gatzia, D. E. (2018). The real epistemic significance of perceptual learning. *Inquiry*, 61(5–6), 543–558.
- Bukach, C. M., Gauthier, I., & Tarr, M. J. (2006). Beyond faces and modularity: the power of an expertise framework. *Trends in cognitive sciences*, 10(4), 159-166.
- Burge, T. (2022). *Perception: First form of mind*. Oxford University Press.
- Byrne, A., & Hilbert, D. (2008). Basic sensible qualities and the structure of appearance. *Philosophical Issues*, 18, 385-405.
- Campbell, Keith. (1990). *Abstract Particulars*. Oxford: Basil Blackwell.
- Carcea, I., & Froemke, R. C. (2013). Cortical plasticity, excitatory–inhibitory balance, and sensory perception. *Progress in brain research*, 207, 65-90.
- Chisholm, R. (1977). *Theory of knowledge*. Prentice Hall.
- Chomsky, N. (1981). Knowledge of language: Its elements and origins. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 295(1077), 223-234.
- Chudnoff, E. (2021). *Forming impressions: Expertise in perception and intuition*. Oxford University Press.
- Churchland, P. (2005). Chimerical colors: Some phenomenological predictions from cognitive neuroscience. *Philosophical psychology*, 18(5), 527-560.
- Connolly, K. (2019). *Perceptual learning: The flexibility of the senses*. Oxford University Press.
- Coppola, D. M., White, L. E., Fitzpatrick, D., & Purves, D. (1998). Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proceedings of the National Academy of Sciences*, 95(5), 2621-2623.
- Dubova, M., & Goldstone, R. L. (2021). The influences of category learning on perceptual reconstructions. *Cognitive Science*, 45(5), e12981.

- Froemke, R. C., Carcea, I., Barker, A. J., Yuan, K., Seybold, B. A., Martins, A. R. O., ... & Schreiner, C. E. (2013). Long-term modification of cortical synapses improves sensory perception. *Nature neuroscience*, 16(1), 79-88.
- Furmanski, C. S., & Engel, S. A. (2000). An oblique effect in human primary visual cortex. *Nature neuroscience*, 3(6), 535-536.
- Gauthier, I., & Tarr, M. J. (1997). Becoming a “Greeble” expert: Exploring mechanisms for face recognition. *Vision research*, 37(12), 1673-1682.
- Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature neuroscience*, 2(6), 568-573.
- Goldstone, R. L. (1998). Perceptual learning. *Annual review of psychology*, 49(1), 585-612.
- Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1), 69-78.
- Green, E. J. (2020). The perception-cognition border: A case for architectural division. *Philosophical Review*, 129(3), 323-393.
- Groen, I. I., Silson, E. H., & Baker, C. I. (2017). Contributions of low-and high-level properties to neural processing of visual scenes in the human brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1714), 20160102.
- Heil, J. (2021). *Appearance in reality*. Oxford University Press.
- Hofer, S. B., Mrcic-Flogel, T. D., Bonhoeffer, T., & Hübener, M. (2006). Prior experience enhances plasticity in adult visual cortex. *Nature neuroscience*, 9(1), 127-132.
- Hubel, D. H., & Wiesel, T. N. (1965). Receptive fields and functional architecture in two nonstriated visual areas (18 and 19) of the cat. *Journal of neurophysiology*, 28(2), 229-289.
- Haynes, J. D., & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature neuroscience*, 8(5), 686-691.
- Jehee, J. F., Ling, S., Swisher, J. D., van Bergen, R. S., & Tong, F. (2012). Perceptual learning selectively refines orientation representations in early visual cortex. *Journal of Neuroscience*, 32(47), 16747-16753.
- Jenkin, Z. (2023a). Perceptual learning. *Philosophy Compass*, e12932.
- Jenkin, Z. (2023b). Perceptual learning and reasons-responsiveness. *Noûs*, 57(2), 481-508.
- Jenkin, Z. (2023c). The function of perceptual learning. *Philosophical Perspectives*, 37(1), 172-186.
- Karmarkar, U. R., & Dan, Y. (2006). Experience-dependent plasticity in adult visual cortex. *Neuron*, 52(4), 577-585.
- Lewis, D. K. (1986). Causal explanation. In *Philosophical Papers: Volume II*. Oxford: Oxford University Press.
- Lively SE, Logan JS, Pisoni DB. (1993). Training Japanese listeners to identify English /r/ and /l/. II. The role of phonetic environment and talker variability in learning new perceptual categories. *J. Acoust. Soc. Am.* 94:1242–55.

- Livermore, A., & Laing, D. G. (1996). Influence of training and experience on the perception of multicomponent odor mixtures. *Journal of Experimental Psychology: Human Perception and Performance*, 22(2), 267.
- Livingston KR, Andrews JK. (1995). On the interaction of prior knowledge and stimulus structure in category learning. *Q. J. Exp. Psychol. Hum. Exp. Psychol.* 48A:208–36.
- Lowe, E. J. (2005). *The four-category ontology: A metaphysical foundation for natural science*. Clarendon Press.
- May, A. (2011). Experience-dependent structural plasticity in the adult human brain. *Trends in cognitive sciences*, 15(10), 475-482.
- Mayer, M.J. (1977). Development of anisotropy in late childhood. *Vision Research*. 17, 703–710.
- Mitchell, D. E., Freeman, R. D., Millodot, M., & Haegerstrom, G. (1973). Meridional amblyopia: evidence for modification of the human visual system by early visual experience. *Vision research*, 13(3), 535-I.
- Mitchell, D. E., & Wilkinson, F. (1974). The effect of early astigmatism on the visual resolution of gratings. *The Journal of physiology*, 243(3), 739-756.
- O’Callaghan, Casey. (2019). *A multisensory philosophy of perception*. Oxford University Press.
- O’Hara W. 1980. Evidence in support of word unitization. *Percept. Psychophys.* 27: 390–402
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549.
- Roberson, D., Davidoff, J., Davies, I. R., & Shapiro, L. R. (2005). Color categories: Evidence for the cultural relativity hypothesis. *Cognitive psychology*, 50(4), 378-411.
- Rosenthal, D. (2010). How to think about mental qualities. *Philosophical Issues*, 20, 368-393.
- Russell, B. (1912). *The problems of philosophy*. OUP.
- Schellenberg, S. (2018). *The unity of perception*. Oxford University Press.
- Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, 412(6846), 549-553.
- Schyns PG, Rodet L. (1997). Categorization creates functional features. *J. Exp. Psychol.: Learn. Mem. Cogn.* 23:681–96.
- Siegel, S. (2006). Which properties are represented in perception? *Perceptual experience*, 1, 481-503.
- Siegel, S. (2010). *The contents of visual experience*. Oxford University Press.
- Siegel, S. (2017). *The rationality of perception*. Oxford University Press.
- Stryker, M. P., Sherk, H., Leventhal, A. G., & Hirsch, H. V. (1978). Physiological consequences for the cat's visual cortex of effectively restricting early visual experience with oriented contours. *Journal of Neurophysiology*, 41(4), 896-909.
- Tagawa, Y., Kanold, P. O., Majdan, M., & Shatz, C. J. (2005). Multiple periods of functional ocular dominance plasticity in mouse visual cortex. *Nature neuroscience*, 8(3), 380-388.
- Tanaka J, Taylor M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cogn. Psychol.* 23:457–82.

- Tolias, A. S., Keliris, G. A., Smirnakis, S. M., & Logothetis, N. K. (2005). Neurons in macaque area V4 acquire directional tuning after adaptation to motion stimuli. *Nature neuroscience*, 8(5), 591-593.
- Weinberger NM. (1993). Learning-induced changes of auditory receptive fields. *Curr. Opin. Neurobiol.* 3:570–77.
- Williams, D. C. (1953). The Elements of Being: I. *Review of Metaphysics*, vol. 7, pp. 3-18. Reprinted in Mellor and Oliver 1997.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford university press.
- Zhu, S. C., & Wu, Y. N. (2023). *Computer vision: Statistical models for Marr's paradigm*. Springer Nature.