

Perceptual Learning and Perceptual Recognition: Workshop Report

By Kevin Connolly, Dylan Bianchi, Craig French, Lana Kuhle, and Andy MacGregor

This report highlights and explores five questions that arose from the workshop on perceptual learning and perceptual recognition at the University of York on March 19th and 20th, 2012.

1. What Is Perceptual Learning?

Different speakers at the York workshop defined perceptual learning in different ways (as Fiona Macpherson pointed out during the final panel discussion). Very roughly, Ian McLaren argued that perceptual learning involved gaining a new ability to discriminate. In the final panel discussion, however, some suggested that this might be sufficient but not necessary. For according to some, perceptual learning is exposure that affects subsequent experience. Better discrimination need not follow on changed experience. In what follows, we explain the details of this discussion.

In his workshop talk, McLaren defined perceptual learning as “a type of perceptual expertise that can develop as a consequence of exposure to stimuli that would otherwise be difficult to tell apart.” He provided a theory for how perceptual learning occurs. Suppose you are presented with two similar perceptual stimuli (X and Y) the same number of times. Being similar stimuli, many features are shared between X and Y. When a feature is shared, you are exposed to it twice as much as you are a non-shared feature (once when presented with X, and a second time when presented with Y). According to McLaren’s theory, you become habituated to these shared features, which results in those features becoming less salient. On the other hand, you are exposed to the non-shared features only half as much as the shared features. So, you become less

habituated to the non-shared features. This results in the non-shared features becoming more salient in comparison to the shared features.

McLaren defined perceptual learning in terms of acquiring a new discriminative ability. Fiona Macpherson suggested that McLaren's definition might fall under a broader notion of perceptual learning according to which exposure to stimuli affects our experience in a particular way. On this view, gaining a new discriminative ability would be just one way in which exposure to a stimulus affects our experience, but there might be other ways. Along these same lines, Kevin Connolly pointed to Eleanor Gibson's definition of perceptual learning as "any relatively permanent and consistent change in the perception of a stimulus array, following practice or experience with this array" (1963, p. 29), to which Mohan Matthen noted that Gibson defines perceptual learning not just as learning from perception, but as learning to perceive.

Later on in the workshop, Macpherson offered one reason why defining perceptual learning in terms of exposure affecting one's experience might be too broad. If you stare at a green circle, the edges of the green seem to disappear. That is a case of exposure affecting your experience, but not a case of perceptual learning. In response to Macpherson, Kati Farkas suggested that it is not just exposure, but repeated exposure that makes something a case of perceptual learning. She also pointed out that often the process of perceptual learning involves not just simple exposure, but the use of descriptors. We might learn how to recognize a pinot noir, for instance, because someone tells us which features to look for when tasting or smelling a pinot noir.

2. Can Perceptual Experience Be Modified by Reason?

Is visual experience cognitively penetrable by belief? In their workshop presentation, Kati Farkas and David Bitter acknowledged the influence of prior associations. For example, Delk and Fillenbaum (1965) demonstrated that when subjects are asked to adjust the color of *heart-shaped* cutouts against a grey background they tend to make them redder than when they are assigned the same task with a culturally meaningless shape such as a circle. Farkas and Bitter wanted to know if a *belief*—namely that hearts are red—is responsible for this effect.

Call perceptual experience “cognitively penetrable” just if the phenomenal character of perceptual experience can be altered by higher-level cognitive states (for example, beliefs) in a way that is influenced by the content of the state. Farkas and Bitter doubt that cognitive penetration is at play in the above example. Through perceptual learning (or, as they also put it, “repeated perceptual exposure”) subjects acquire a disposition. The disposition is then responsible for a perceptual effect (in the above case, the effect being that subjects see orange shapes as more red in certain conditions). Farkas and Bitter have no problem with the idea that the effect was genuinely experiential. But they ask: what mental state underlies or constitutes the disposition in the cases? Is it an association or a belief?

The disposition in the color cases is stubborn, that is, difficult to lose. Even if one comes to believe that in a particular country love-hearts are typically green, and *not* typically red, one would nonetheless retain the disposition that causes one to see love-hearts as more red than they are. Farkas and Bitter suggested that losing the disposition is likely to require nothing short of perceptual re-training—a very different kind of process from belief-adjustment. They argue that *belief* is revised in the face of conflicting information and counter-evidence, i.e., through rational persuasion. By contrast, expectations formed by association are *extinguished* by repeated exposure to perceptions that run counter to the expectation. Farkas and Bitter invoke Tamar

Szabo Gendler's (2008) notion of an "alief" here. On Gendler's account, aliefs have behavioral and experiential effects similar to beliefs, but are *not* sensitive to rational persuasion.

Macpherson replied to Farkas and Bitter by arguing that we should distinguish between different belief accounts of the relevant disposition. One view, which we can grant is false, is that the disposition is constituted by just the belief that love-hearts are (typically) red. But another view, which isn't shown to be false by Farkas and Bitter's line of thought, is that the disposition is constituted by a *cluster* of beliefs, all pertaining to love-hearts and redness (e.g., which includes, in the normal case, the belief that *love-hearts are typically red, the love-hearts I've seen have been red, I remember experiencing red hearts*, and so on). Losing a part of the cluster doesn't necessarily bring about loss of the cluster. Belief clusters may be more stubborn than single beliefs.

The dispute between Farkas and Bitter on one side and Macpherson on the other is, thus, *not* about whether acquired dispositions can affect perceptual experience, but rather about whether rational inference can affect perceptual experience.

3. How Does Perceptual Learning Alter Perceptual Phenomenology?

Some philosophers argue that perceptual learning enriches perceptual phenomenology. Charles Siewert, for instance, writes that after we learn to recognize some general type, such as a sunflower, certain features "'stand out for us as significant' and 'go together.'" Such cases, he argues, show the "wealth" or "richness" of visual experience (Siewert, 1998, pp. 255, 259). In the workshop, however, an alternative account arose for how perceptual learning alters perceptual phenomenology. In his commentary on Ian McLaren, Kevin Connolly argued that one implication of McLaren's account is that perceptual learning in fact *impoverishes* one's perceptual phenomenology.

Ian McLaren has explored the influence of “latent inhibition” on perceptual learning: subjects pre-exposed to a stimulus have more difficulty being conditioned on that stimulus than on a novel stimulus. For example, R.E. Lubow and A.U. Moore (1959) found that if first shown a light ten times, sheep were slower to associate a shock with the light than with a novel stimulus—a turning rotor. McLaren argues that this is because pre-exposed stimuli tend to be less salient than novel stimuli by latent inhibition. Here, perceptual learning decreases the salience of features that a stimulus shares with other stimuli. Thus, the *unique* features of the stimulus become more salient by comparison. For example, according to McLaren’s account of face recognition, exposure to a variety of faces renders common features less salient and unique features relatively more salient.

In his commentary on McLaren’s talk, Connolly pointed out that philosophers often focus on *increased* salience in perceptual learning. Siewert, for example, writes: “Notice how different your neighborhood looks to you now that you have lived there for a while, than it did on the day you first arrived” (1998, p. 257). Is this a difference in increased or reduced salience? Arguably (and McLaren later agreed with this), the familiar features of your neighborhood become less salient over time and this accounts for the difference in phenomenology.

Further, if McLaren is right, features that are salient at one time might not be salient at another time. Susanna Siegel argues that for someone who has recently learned to distinguish pine trees from non-pine trees, pine trees look visually salient (2010, p. 100), and become a part of visual phenomenology. On McLaren’s account, Connolly argued, unique features can become familiar over time. But if pine tree salience doesn’t hold constant, why think that pine trees remain part of your visual phenomenology?

Finally, on McLaren's account, the features common to pine trees and non-pine trees *decrease* in salience, thereby creating a differential between those common features and the features unique to pine trees. So, why think that perceptual learning is enriching your perceptual phenomenology at all? It is relevant, though, *what* is salient. Arguably, McLaren's focus is on feature salience, while Siegel's is on object or kind salience. Kind salience could be mediated by unique features.

4. How Does Perceptual Learning Alter the Contents of Perception?

When you learn to discriminate pine trees from non-pine trees, do pine trees start to stand out and engage your attention? In the closing panel, Mohan Matthen argued that there is no reason to say that *pine trees* engage your attention as such. Perhaps it is simply the "pine tree appearance." Is a natural kind represented (as Siegel argues), or is it just a complex appearance (as Matthen suggests). Is there—in addition to the sensuous pine tree appearance—non-sensuous perceptual content which goes beyond it (as Paul Noordhof suggests)? More generally, in perceptual learning cases, how are the contents altered?

In Ian McLaren's model, recognition is mediated by altered salience (see question three). As against this, Tom Stoneham described Berkeley's argument that perceptual learning yields awareness of something that we cannot strictly be said to see, namely distance or depth. According to Berkeley, one does not *see* any feature about which an image is ambiguous or indeterminate. These limitations restrict what is strictly seen to light [i.e. luminance], colors, and figures. Awareness of distance and depth are added to awareness of these simple things.

While Berkeley's argument might appear to be an account of the limitations of vision, Stoneham argued that it is properly understood as a positive explanation of how experience, despite the limits imposed by optics and physiology, is nonetheless able to represent distance and

depth. As such, it offers an account of the relationship – and indeed the gap – between the content of visual experience (what is represented in experience) and what is made available to the eyes when we see (what is strictly seen).

Pascal Belin described more up-to-date findings from psychology that also bear on our understanding of perceptual learning and perceptual content, in this case concerning perception of speech and other vocal sounds. He described recent evidence that vocal and non-vocal sounds are processed quite differently in the brain and that, within the category of vocalizations, it is possible to discern different processing pathways for various aspects of vocal sounds (e.g. meaning, affect, and speaker's identity). As Åsa Wikforss noted, this might bear on how we view the listener's auditory content – that his or her experience in listening to a speaker represents special kinds of objects and properties, such as voices, emotions, and so on.

Later in the workshop, Randy Flanagan's talk brought out various ways in which perception aids and interacts with motor control and adaptation. A natural way to think of perceptual learning is as a process in which experience alters perceptual content and *thereby* promotes adjustments to behavior. However, as was noted in discussion, there is also the enactivist view (e.g. Noë, 2004) that perceptual content is partly constituted by possibilities for action. If that were so, certain forms of perceptual learning could be understood as influencing content *via* their effects on behavior.

The closing panel discussed the case of professional chicken-sexers, who are trained to segregate newly born chicks according to their gender. It was suggested that these expert subjects may be unable to explain the basis upon which they tell male from female chicks, and the reason for this is that their experiences represent the chicks' being male or being female rather than their having a specific sensible feature that marks the difference. In short, even where

training enhances the salience of some features of a stimulus, this does not necessarily result in the subject's representing any hitherto unnoticed sensible features. The increased salience might not show that the subject is representing the sensible feature, but rather just what that sensible feature indicates.

5. How Is Perceptual Learning Coordinated with Action?


In the final presentation of the workshop, Randy Flanagan showed how the motor system displays a high degree of cognitive sophistication. During the course of executing tasks and learning new ones, the motor system employs a variety of information. Although much of the information contained within it is inaccessible to consciousness and higher cognition, the motor system contains rich bodies of information about task-relevant features of the environment and the motor system itself. This information is not innate or hardwired, but is learned through perceptual and motor engagement with a particular task or environment. The phenomena of motor learning raise at least two philosophically interesting questions: (1) what kinds of representations are involved in motor learning? and (2) Can the contents of these representations be shared among sense modalities, and if so, to what extent? On this second question, Flanagan showed some ways in which the contents are shared, while Malika Auvray, his commenter, showed some ways in which the sharing is limited.

In answering the first question, Flanagan identified four main kinds of representations that store the information utilized by the motor system: (i) probability distributions encoding statistical regularities in the environment; (ii) models used to generate predictions about the perceptual consequences of a motor command; (iii) uncertainties associated with estimated features of the environment; and (iv) evaluations of possible courses of action on the basis of decision-relevant information such as energy cost. Notably, these representations are intrinsically

probabilistic, and are manipulated according to Bayesian operations during the execution of motor skills and motor learning.

The second question is prompted by the observation that motor learning is sometimes facilitated by information contributed from a different sense modality. Experiments so far have focused on vision; two noteworthy results are as follows. First, subjects compute a Bayes-optimal prediction from noisy information from the visual and motor systems when extracting task-relevant information from the environment (Ernst and Banks, 2002). Second, subjects who had watched others moving in an unusual force field were quicker to learn to move in that force field than subjects who had not (Diedrichsen, 2007).

These results provide clear evidence that there is some sharing of information between visual and motor systems. But other well-known findings suggest that this sharing is limited. In her comments on Flanagan's talk, Malika Auvray discussed a study (Aglioti, DeSouza, and Goodale, 1995) in which subjects were shown a pair of disks embedded in a Titchener-like context (see figure below) and asked to grasp them. Despite reports to the effect that one of the disks appeared larger (an illusion), reaching behavior remained constant between disks; both disks were grasped with an accurate grip size. Apparently the visual system is duped by the illusion while the motor system is not. Results such as these have motivated the distinction in cognitive science between dorsal and ventral processing streams.



In his reply to Auvray, Flanagan pointed out that studies in which motor behavior remains insensitive to visual illusions have tended to focus on reaching. When grasping, object manipulation, and abnormal force fields are involved, Flanagan hypothesized, there is more potential for information sharing. These operations are more computationally intensive and display sensitivity to cognitive load, even where that load draws on higher cognitive processes. If Flanagan is correct in this speculation, the distinction between dorsal and ventral streams may need to be reconceived when the full range of cognitive operations involved in motor learning is taken into account. The sophistication of motor learning may suggest a closer connection between the motor system and core cognition than theorists have previously acknowledged.

References:

- Aglioti, S., DeSouza, J. F. X., and Goodale, M. A. (1995). "Size-contrast illusions deceive the eye but not the hand." *Current Biology*. 5, 679-685.
- Delk, J.L. and Fillenbaum, S. (1965). "Differences in Perceived Colour as a Function of Characteristic Color," *The American Journal of Psychology*, 78 (2): 290–93.
- Diedrichsen, J. (2007). "Optimal task-dependent changes of bimanual feedback control and adaptation." *Current Biology*. 17, 1675–1679.
- Ernst M.O., and Banks, M.S. "Humans integrate visual and haptic information in a statistically optimal fashion." *Nature*. 415: 429–433, 2002.
- Gendler, Tamar Szabo (2008). "Alief and Belief." *Journal of Philosophy*, 105 (10): 634–63.
- Gibson, E. J. (1963) "Perceptual learning." *Annu. Rev. Psychol.* 14, 29–56.
- Lubow, R.E., and Moore, A.U. (1959). "Latent inhibition: The effect of non-reinforced preexposure to the conditional stimulus." *Journal of Comparative and Physiological Psychology*, 52, 415-419.
- Noë, Alva (2004). *Action in Perception*. Cambridge, Mass.: MIT Press.
- Siegel, Susanna (2010). *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siewert, Charles (1998). *The Significance of Consciousness*. Princeton: Princeton University Press.