

## Varieties of perceptual learning

N. J. MACKINTOSH

*Cambridge University, Cambridge, England*

Although most studies of perceptual learning in human participants have concentrated on the changes in *perception* assumed to be occurring, studies of nonhuman animals necessarily measure discrimination learning and generalization and remain agnostic on the question of whether changes in behavior reflect changes in perception. On the other hand, animal studies do make it easier to draw a distinction between supervised and unsupervised learning. Differential reinforcement will surely teach animals to attend to some features of a stimulus array rather than to others. But it is an open question as to whether such changes in attention underlie the enhanced discrimination seen after unreinforced exposure to such an array. I argue that most instances of unsupervised perceptual learning observed in animals (and at least some in human animals) are better explained by appeal to well-established principles and phenomena of associative learning theory: excitatory and inhibitory associations between stimulus elements, latent inhibition, and habituation.

---

According to one definition, perceptual learning “improves discrimination between stimuli that could not be discriminated before the learning; observers may learn to perceive something new that they could not perceive before” (Fahle, 2002a, p. ix). Although Fahle (2002a) used the second part of his definition to exclude (wrongly, in my view) many of the studies I shall be discussing, the first part is sufficiently wide that it can encompass a whole range of different phenomena, reflecting the operation of a wide variety of different psychological processes. If I train rats on a discrimination between a circle and a triangle by reinforcing responses to one stimulus but not the other, they will initially, and for many trials, respond at chance but will eventually learn to choose the reinforced stimulus on 100% of the trials. This seems to be a case of initial failure to discriminate giving way to perfect discrimination. Fahle (2002a) might want to say that my rats have not learned to perceive something new that they could not perceive before, but how does he know that? Their initial behavior is entirely consistent with the inference that they failed to perceive the difference between the circle and the triangle, and their later behavior with the inference that they slowly learned to do so.

### **Supervised Versus Unsupervised Learning**

My own concern with this example is that it is clearly a case of supervised learning; the result is surely unsurprising and would be predicted by any theory of conditioning. It is also surely not the same as what Gibson and Walk (1956) demonstrated in what is widely regarded as a classic study of perceptual learning. Gibson and Walk reared one group of rats in cages with metal circles and triangles hanging on the walls and a control group in ordinary, unadorned cages. When subsequently trained (as above) on

a discrimination between a circle and a triangle, the first group learned substantially more quickly than the second. The apparently unreinforced prior exposure to circles and triangles had made them easier to tell apart. The learning that occurred during this phase of the experiment was unsupervised. Unlike the case in which differential reinforcement teaches an animal (or person) to respond to one stimulus rather than another, the effect observed by Gibson and Walk, widely regarded as a paradigm case of perceptual learning, is not so easily explained.

Although there were a number of follow-up studies published in the next few years (reviewed by Hall, 1980), this perceptual-learning effect attracted rather little attention from most animal learning theorists—largely, I suspect, because it did not fit in with most of their theories. Gibson (1969) herself explained her results by suggesting that mere exposure to two or more similar stimuli would attract animals’ attention to their subtle, differentiating features at the expense of the more salient features they shared in common. A variety of theories of discrimination popular in the 1960s and 1970s appealed to the notion of attention (e.g., Lovejoy, 1968; Sutherland & Mackintosh, 1971), arguing that in order to solve a discrimination problem, animals had to learn both which stimulus dimension to attend to and which particular stimulus was rewarded and which was not. So the notion that animals would learn to attend to the differentiating features of a set of discriminative stimuli was quite widely accepted, as it still is today in more recent theories proposed by Kruschke (2001), Le Pelley (2004), and Mackintosh (1975). The problem is that all such theories assume that attentional learning is supervised learning, requiring differential reinforcement. Where was the differential reinforcement in the exposure phase of Gibson and Walk’s (1956) experiment? It is one

---

N. J. Mackintosh, n.mackintosh@psychol.cam.ac.uk

---

thing to say that animals will learn to attend to some features of a set of stimuli, rather than others, if it pays them to do so; but why should Gibson and Walk's rats have paid any attention at all to the circles and triangles hanging from the walls of their cage, let alone have learned to attend to their differentiating features?

**Self-supervised learning in humans.** The problem seemed to become even more acute in studies of perceptual learning in human participants (although the authors of those studies did not, I am sure, lose any sleep over the failure of animal learning theory to explain their data). In experiments on *hyperacuity* (Fahle, 2002b), for example, Vernier acuity is measured by the participant's ability to say whether the top or bottom of two lines shown one above the other is displaced to the left. After several hundred trials, participants can detect displacements very much smaller than the diameter of foveal cones. Although this final level of performance may indeed be remarkable (but see Mollon & Danilova, 1996, for an acute analysis of how this is probably achieved), if participants are told on each trial whether their guess is right or wrong, it is not really very surprising that their performance should improve as they learn to attend to the very subtle cues that give the answer. The problem is that such feedback is not necessary. Although people learn more slowly in the absence of feedback, and although random feedback (being told you were right or wrong, regardless of whether your guess *was* actually right or wrong) prevents learning, the same final level of performance can be achieved whether feedback is given or not.

The question is whether the absence of feedback turns this into a case of genuinely unsupervised learning. The fact is surely that participants in these experiments are not *simply* being exposed to the stimuli. They know that on each trial, they are looking at two lines, one of which is displaced to the left or right of the other, and they can be reasonably confident that, about half the time, it will be the top line, and half the bottom. Their task is to decide which it is on any trial. So they will examine the stimulus array on every trial, trying to detect any difference between the array seen on one trial and that seen on the one before or the one after that. Since there is, of course, a difference, albeit very, very hard to detect, it is not so surprising that they eventually discover it. This could reasonably be called *self-supervised learning*.

Other experiments with people, more similar in design to Gibson and Walk's (1956), do involve separate exposure and test phases. But in order to engage the participants' attention in the exposure phase, they are often instructed to look carefully at all the stimuli, which will not all be the same. Here is one example:

At the start of the experiment, participants were seated approximately 60 cm from the computer monitor and were presented with instructions on the screen. They were told to pay attention to the stimuli, that any stimulus differences they detected would be useful later in the experiment.

(Mitchell, Kadib, Nash, Lavis, & Hall, 2008, p. 477)

In another case, the instructions required participants to keep track of the number of times each different stimulus

appeared (Mundy, Honey, & Dwyer, 2007). This is not genuinely unsupervised learning. Once again, participants are being told to notice the differences between the stimuli they are exposed to.

Whether such instructions are explicit or not, it seems plausible to suppose that most people acting as participants in psychological experiments will assume that, regardless of what the experimenter says, it will pay to pay attention to anything going on. It is not at all obvious that the same can be said about Gibson and Walk's (1956) rats. In experiments with nonhuman animals, there does seem to be a chance that exposure to a set of stimuli will indeed be mere exposure and that any learning that occurs during this phase will be unsupervised. That inference is supported by the two most obvious outcomes of such exposure to a single stimulus: The animal's initial response to the stimulus habituates with repeated exposure to it, and subsequent conditioning to the stimulus is seriously retarded (latent inhibition). Both effects suggest, if anything, a loss of attention to the stimulus, rather than any gain.

### Explanations of Unsupervised Perceptual Learning in Animals

The procedure we have mostly used in our studies of perceptual learning is to give thirsty rats the opportunity to drink two or more flavored solutions over the course of a number of exposure trials. In the simplest case, one compound flavor is a mixture of saline and lemon, and the other a mixture of sucrose and lemon. These compound solutions can be made as similar or dissimilar as needed by increasing the concentration of the unique (saline/sucrose) or common (lemon) flavor. The preparation has the further advantage that thirsty rats will certainly consume the solutions offered them, and so we can guarantee that they are exposed to the various flavors during the exposure phase, but without any requirement to discriminate between them. That comes in the next phase of the typical experiment, when an aversion is conditioned to one solution by pairing its consumption with lithium chloride and the generalization of that aversion to the other solution is measured.

**Differential latent inhibition.** If it is implausible to suppose that animals learn to attend to the differentiating features of these compound flavors during the exposure phase, what then is the explanation of the enhanced discrimination between them that has standardly been seen as a result of such exposure? Perhaps surprisingly, part of the answer is provided by the two phenomena just referred to: latent inhibition and habituation. The role of latent inhibition is readily understood. If an aversion is conditioned to a solution of saline and lemon, it will generalize to sucrose-lemon to the extent that the aversion was conditioned to the lemon flavor. Conversely, discrimination will be good to the extent that the aversion was conditioned to saline rather than to lemon. Now, if animals receive 10 trials of exposure to each of the two compound solutions, they will receive 10 trials of exposure to saline and sucrose but 20 trials of exposure to the lemon. But exposure to a stimulus causes latent inhibition—that is, retards

**Table 1**  
**Experiment by Scahill and Mackintosh (2004)**

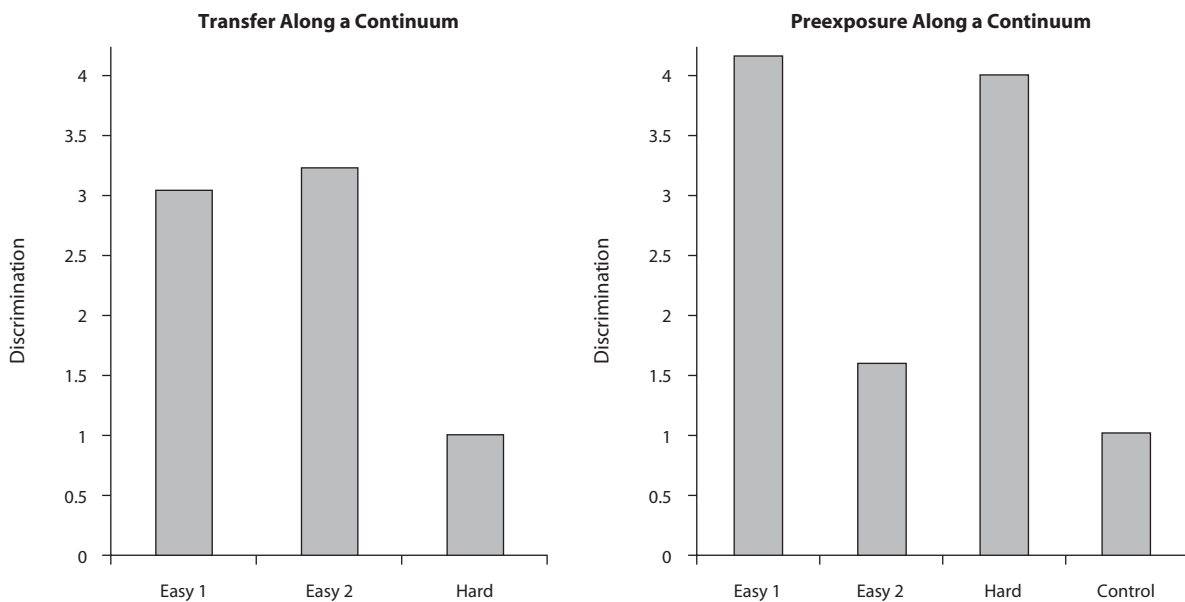
Group	Stage 1	Stage 2
Easy 1	AX-BX	
Easy 2	A-B	AX-BX
Hard	AX-BX	

Note—A and B are saline and sucrose, X is lemon. The size of the letters indicates the concentration of the solution.

subsequent conditioning. So, this exposure regime will have resulted in 10 trials worth of latent inhibition to saline and sucrose but 20 trials worth of latent inhibition to lemon. When the saline–lemon solution is paired with the lithium injection, the aversion conditioned to saline will be stronger than that conditioned to lemon, and generalization to sucrose–lemon will be reduced. This differential latent inhibition of common and unique elements will thus have much the same effect as if Gibson and Walk (1956) had been right and animals had learned to attend to the unique elements and ignore the common elements of the two compound solutions. And there is abundant evidence that such differential latent inhibition does contribute to the perceptual-learning effects seen in animal studies (McLaren & Mackintosh, 2000).

I said that differential latent inhibition will have the same effect as learning to attend to/ignore unique/common elements. That will often be true. But it is also possible to differentiate the two. One way of teaching an animal a very difficult discrimination, first discovered by Pavlov (1927), is to train them on an easier version of the discrimination—an effect termed *transfer along a continuum* or the *easy-*

*to-hard effect*. Lawrence (1952), for example, trained one group of rats on a difficult brightness discrimination between two mid-grays. Another group was initially trained on an easier brightness discrimination between a dark and a light gray. When shifted to the difficult discrimination, this second group performed substantially more accurately than the first, which had been trained throughout on the difficult problem. One plausible interpretation of this result is that training on the easy discrimination is much more effective at teaching animals to attend to the relevant features of the stimuli—that is, the difference in their brightness—and this efficient attentional learning then transfers to the hard problem. Scahill and Mackintosh (2004) adapted this procedure to our flavor aversion paradigm in an experiment whose design is shown in Table 1, where the saline and sucrose flavors are represented as A and B and the common lemon flavor as X. Group Hard was trained throughout the experiment on a difficult discrimination between two compound solutions, one consisting of weak saline and strong lemon, the other of weak sucrose and strong lemon. Consumption of one was always followed by a very small injection of lithium, whereas consumption of the other had no further consequence. As can be seen from the table, for reasons that will become apparent, we made this discrimination easier in two different ways: one by increasing the concentration of saline and sucrose, the second by eliminating the common lemon flavor. Both manipulations produced an easy-to-hard effect. As can be seen in the left-hand panel of Figure 1, when tested on the hard discrimination, both Groups Easy 1 and 2 discriminated between the two solutions more effectively than did Group Hard.



**Figure 1.** Results of an experiment by Scahill and Mackintosh (2004). The left-hand panel shows Stage 2 performance on the hard discrimination of three groups of rats trained on one of the two easy discriminations or the hard discrimination in Stage 1. The right-hand panel shows the Stage 2 performance on the hard discrimination of groups simply preexposed to the stimuli of the easy and hard discriminations in Stage 1. The control group received no preexposure. The measure of discrimination is the average (over 10 days) amount of the safe fluid consumed on each trial (out of a maximum of 5 ml available) minus the amount of the fluid paired with an LiCl injection consumed.

What would happen if, instead of *training* animals to discriminate between these compound solutions in the first phase of the experiment, we had simply given them unreinforced exposure to these solutions, before finally training all of them on the hard discrimination? The answer is shown in the right-hand panel of Figure 1. There was a perceptual learning effect—at least in two groups. Groups Easy 1 and Hard both learned to discriminate between the two hard solutions significantly more quickly than did the animals that had received no exposure to any compound solutions. But Group Easy 2 did not. Given that Group Hard learned the hard discrimination as fast as Group Easy 1 and significantly faster than Group Easy 2, it is clear that there was no easy-to-hard exposure effect. What is the explanation of these results? They are, in fact, exactly what one would expect on the assumption that perceptual learning here is a consequence of latent inhibition to the common lemon solution. Groups Hard and Easy 1 received the same exposure to lemon and, so, showed the same perceptual learning effect. But Group Easy 2 received no exposure to lemon and so showed no better discrimination than did a group that received no exposure at all. Differential latent inhibition of common and unique elements may sometimes mimic the putative effects of learning to attend to unique elements and ignore common elements. But not necessarily. Put another way, if the perceptual learning produced by mere exposure to two similar stimuli is a consequence of such attentional learning, the results in the right-hand panel of Figure 1 should have looked exactly the same as those in the left-hand panel.

A second way to show that differential latent inhibition of common and unique elements is not the same process as learning to attend to the relevant features and ignore the irrelevant features of two or more stimuli is to put them in opposition and see which wins. McLaren, Leavers, and Mackintosh (1994) required people to categorize two sets of variable exemplars based on two prototypes. If we call the two prototypes A and B, the exemplars of A can be labeled  $a_1, a_2, a_3 \dots a_n$ , and the exemplars of B are  $b_1, b_2, b_3 \dots b_n$ , where each exemplar contained some elements in common with its prototype and some unique elements. The relevant features of each exemplar, therefore, which will enable its correct categorization, will be those it shares in common with its prototype, whereas its unique features will not enable participants to categorize it correctly and are, therefore, irrelevant. Attentional learning therefore implies that it will be peculiarly difficult to learn a new discrimination between two new exemplars of one of the categories: The features that distinguish two exemplars of the same category are their unique features, which have been ignored, whereas categorization training has directed attention to the features that they share in common with their prototype, which are of no help in telling them apart.

The latent inhibition analysis makes a quite different prediction. Exposure to 100 exemplars of each category during the course of categorization training will result in strong latent inhibition to the features that the exemplars share in common and little or no latent inhibition to their

unique features. Differential latent inhibition therefore predicts that learning to discriminate between two new exemplars of the same category will be easy, and this was the result observed. It is worth remarking that this fits well with common experience. Thus, exposure to a wide variety of different red wines (provided it is suitably paced) is what enables the wine expert to discriminate between new instances of red wine, identifying the country or region, perhaps even the vineyard and year, of origin. Expertise, as we all know, is a result of prolonged experience, but here the important aspect of that experience that allows such fine discrimination between different wines is that it reduces the salience of those features that all (well, most) red wines share in common.

**Schedule of preexposure.** From the first, it was always apparent that differential latent inhibition of common and unique features could not be the only mechanism underlying unsupervised perceptual learning (Mackintosh, Kaye, & Bennett, 1991; McLaren, Kaye, & Mackintosh, 1989). One experimental finding that has been taken as reinforcing this conclusion was first reported by Honey, Bateson, and Horn (1994) in an experiment with domestic chicks but has since been replicated in numerous other studies, in experiments on taste aversion in rats (Bennett & Mackintosh, 1999; Symonds & Hall, 1995) and on spatial learning in rats (Prados, Artigas, & Sansa, 2007) and in a variety of different experiments with people (Mitchell et al., 2008; Mundy, Dwyer, & Honey, 2006). The magnitude of any perceptual-learning effect depends not only on the amount of exposure to two or more stimuli, but also on the way that exposure is scheduled. Intermixed or alternating exposure to two similar stimuli, AX and BX, produces substantially better subsequent discrimination between them than does the same total number of exposure trials when all the trials with AX precede all those with BX (or vice versa). On the face of it, the fact that the two schedules can equate the total amount of exposure to both stimuli makes it unlikely that they should result in substantial differences in the amount of latent inhibition accruing to the common element, X. Indeed, at least one study has shown no such difference (Bennett & Mackintosh, 1999), although another has shown some (Mondragón & Hall, 2002). Given the reliability of the substantial difference in the size of the perceptual-learning effect that has repeatedly been observed, it does not seem likely that it can be attributed to a not wholly reliable difference in latent inhibition to X. If that is accepted, how else can the perceptual-learning effect be explained? There has been no shortage of suggestions. Four such explanations will be given below.

1. An associative account will assume that during the course of intermixed exposure to AX and BX, various associations will be established between their elements—for example, between A and X and between B and X. After such excitatory associations have been established, it will follow that on AX trials, X will retrieve a representation of B and, similarly, on BX trials, X will retrieve a representation of A. But on AX trials, A will signal the absence of the otherwise expected B, and on BX trials, B will signal the absence of A. These are precisely the conditions that

will establish mutually inhibitory associations between A and B—thus reducing generalization between AX and BX. If all exposure to AX precedes any trial with BX, there will be no reason to expect A to become a signal for the absence of B, and B will become only a weak inhibitor of A (McLaren & Mackintosh, 2000).

There is good evidence that intermixed exposure to AX and BX does establish such mutually inhibitory associations between A and B and that this inhibition does contribute to the perceptual-learning effect observed in these experiments (Dwyer, Bennett, & Mackintosh, 2001; Dwyer & Mackintosh, 2002; Mundy et al., 2006).

2. Gibson (1969) argued that the opportunity to compare and contrast two similar stimuli will enhance their discriminability, because it is a particularly effective way of drawing attention to their unique differentiating features. Such an opportunity is better provided by an intermixed schedule of exposure than by one in which all exposure to one precedes any exposure to the other. Indeed, as Gibson herself argued, the most effective schedule should be one in which the two stimuli are presented simultaneously. This prediction was confirmed by Mundy et al. (2007) in an experiment on face recognition by people, and also by Mitchell et al. (2008) in an experiment in which people were required to discriminate between two very similar checkerboard patterns. The result is important, since it is clearly inconsistent with the inhibitory account of McLaren and Mackintosh: If AX and BX occur together on the same trial, A and B cannot predict each other's absence. Indeed, Mundy et al. (2007) provided evidence that excitatory associations could be formed between A and B under these circumstances.

Gibson's (1969) analysis seems intuitively plausible—certainly, for experiments with human participants. I do not think it a coincidence that I was able to argue above that the instructions given to the participants in these two particular studies explicitly encouraged them to look for differences between the stimuli they were shown in the exposure phase. If one is trying to look for difference between two stimuli, it does not seem surprising that simultaneous presentation of two stimuli, or very rapid alternation between the two, should make one's task easier. Would the same effect happen in animal experiments? In fact, there is evidence that it does not. Both Honey and Bateson (1996) and Bennett and Mackintosh (1999) found *less* perceptual learning after very rapid alternation between two similar stimuli than after somewhat more spaced alternation, and Alonso and Hall (1999) found that rats offered two different flavors to drink (from different bottles) on each exposure trial subsequently discriminated between them no better than did a group given blocked exposure to the two.

3. Gibson's (1969) analysis may make intuitive sense, but one could still ask *why* simultaneous presentation of similar stimuli should make it easier to detect any difference between them. There is, in fact, a lower level explanation that appeals to the phenomenon of habituation. As Honey et al. (1994) and Mundy et al. (2007) noted (see also Mitchell et al., 2008), intermixed and blocked presentations of AX and BX can equate the total number of times

each stimulus is shown, but they cannot easily equate the intervals separating the presentations of the two. More important, unlike blocked exposure, in intermixed exposure (AX, BX, AX, BX, etc.), the interval between successive presentations of the common element X will necessarily be shorter than the interval between successive presentations of either A or B. Rapid repetition of the same stimulus results in habituation of responding to that stimulus—an effect that is greater the shorter the interval between each presentation. Intermixed exposure to AX and BX, especially with relatively short intertrial intervals, will thus ensure greater habituation of X than of A or B on each trial. The relatively common assumption (e.g., Hall, 2003) that habituation produces a loss of salience will mean that intermixed exposure to AX and BX ensures that X is less salient than A or B on each exposure trial. As Bennett and Mackintosh (1999) argued, this difference in salience will not be sufficient in and of itself to ensure better discrimination of AX from BX after intermixed exposure to the two, since the relatively long interval between exposure and subsequent testing will mean that any difference in salience between A or B and X will have worn off by the time of testing. But it would still be enough to underpin the comparison-and-contrast process envisaged by Gibson. If X is less salient than A or B, that amounts to saying that attention will be drawn to the features that differentiate AX from BX.

4. A final explanation has been advanced by Hall (2003). Hall also assumes that the habituation produced by repeated presentations of AX and BX will reduce the salience of the elements making up these stimuli. But he suggests that an additional factor may come into play: A reverse process of dishabituation partially restores the salience of a stimulus whenever its representation is retrieved but the stimulus itself does not occur. The excitatory associations formed between A and X and between B and X will mean that, during intermixed exposure, a representation of B is retrieved on each AX trial and one of A on each BX trial. This will at least partially restore the salience of A and B—but not, of course, of X, which is actually present on each AX and BX trial. As in the associative account offered by McLaren and Mackintosh (2000), such an effect will be much less likely to occur during blocked exposure to the two stimuli.

The results of a number of experiments have provided evidence for this account (Blair & Hall, 2003; Blair, Wilkinson, & Hall, 2004). They have shown that if rats are exposed to AX and X in alternation and, in a separate block of trials, to CX, subsequent conditioning proceeds more rapidly to A than to C. Intermixed exposure to AX and X preserves the salience of A, whose representation is retrieved on the intervening X-alone trials, whereas no such effect is observed to C after a block of CX exposure trials. Other studies have revealed some problems for this interpretation (Dwyer & Honey, 2007; Hall & Rodriguez, in press), and I think I can safely leave it to Hall's own contribution to this symposium (Hall, 2009) to make the case for his account. But, as was noted by Dwyer and Honey, most of these findings can probably also be derived from the McLaren and Mackintosh (2000) model.

## Conclusion

The term *perceptual learning* refers to a number of different phenomena, which require a variety of different explanations. One important distinction is surely that between supervised and unsupervised perceptual learning. That differential reinforcement or provision of feedback should improve discriminative performance is hardly surprising. Theories of discrimination learning that make no mention of the concept of attention (such as Rescorla & Wagner, 1972) have no difficulty in explaining successful discrimination. William James's (1890) venerable theory of perceptual learning, now known as the *acquired distinctiveness of cues* (Hall, 1991), can explain successful discrimination when response to one stimulus is reinforced and to the other is not—as, of course, can theories of attention. The question is whether such theories are sufficient to explain instances of unsupervised perceptual learning.

One issue that needs to be addressed here is whether many experiments with humans that explicitly provide *no* feedback to their participants really constitute evidence of genuinely unsupervised learning. I argued above that many do not. But I should certainly acknowledge that this is not true of all such experiments. The studies reviewed by Watanabe (Tsushima & Watanabe, 2009) in this symposium have convincingly shown that exposure to stimuli that are presented as distractors from the participants' main task can produce a striking perceptual-learning effect. It is noteworthy that the critical feature of these stimuli (the direction of optical flow) was below threshold—and remained so throughout the experiment. Perceptual learning was shown for an easier-to-detect set of stimuli. Whether this is a critical feature of these successful demonstrations remains to be seen. For what it is worth, I would be the first to admit that none of the ideas I have advanced here shed much light on these demonstrations.

The experiments I have described seem to me to show that many instances of unsupervised perceptual learning in humans and other animals can be understood in terms of some rather simple principles of traditional associative learning. Latent inhibition and habituation, although seemingly the antithesis of perceptual learning, actually play an important role in many instances of the phenomenon. That there are other processes going on as well it would be foolish to deny. Perceptual learning, like virtually every other interesting example of a psychological phenomenon, is surely multiply determined.

## AUTHOR NOTE

Most of the ideas expressed here are really Ian McLaren's, and the original experimental work was undertaken by Claire Bennett, Dominic Dwyer, and Vicky Scahill. I am indebted to them all. Correspondence concerning this article should be addressed to N. J. Mackintosh, Department of Experimental Psychology, Cambridge University, Downing Street, Cambridge CB2 3EB, England (e-mail: n.mackintosh@psychol.cam.ac.uk).

## REFERENCES

ALONSO, G., & HALL, G. (1999). Stimulus comparison and stimulus association processes in the perceptual learning effect. *Behavioural Processes*, **48**, 11-23.

- BENNETT, C. H., & MACKINTOSH, N. J. (1999). Comparison and contrast as a mechanism of perceptual learning? *Quarterly Journal of Experimental Psychology*, **52B**, 253-272.
- BLAIR, C. A. J., & HALL, G. (2003). Perceptual learning in flavor aversion: Evidence for learned changes in stimulus effectiveness. *Journal of Experimental Psychology: Animal Behavior Processes*, **29**, 39-48.
- BLAIR, C. A. J., WILKINSON, A., & HALL, G. (2004). Assessments of changes in the effective salience of stimulus elements as a result of stimulus preexposure. *Journal of Experimental Psychology: Animal Behavior Processes*, **30**, 317-324.
- DWYER, D. M., BENNETT, C. H., & MACKINTOSH, N. J. (2001). Evidence for inhibitory associations between the unique elements of two compound flavors. *Quarterly Journal of Experimental Psychology*, **54B**, 97-108.
- DWYER, D. M., & HONEY, R. C. (2007). The effects of habituation training on compound conditioning are not reversed by an associative activation treatment. *Journal of Experimental Psychology: Animal Behavior Processes*, **33**, 185-190.
- DWYER, D. M., & MACKINTOSH, N. J. (2002). Alternating exposure to two compound flavors creates inhibitory associations between their unique features. *Animal Learning & Behavior*, **30**, 201-207.
- FAHLE, M. (2002a). Introduction. In M. Fahle & T. Poggio (Eds.), *Perceptual learning*. Cambridge, MA: MIT Press, Bradford Books.
- FAHLE, M. (2002b). Learning to perceive features below the foveal photoreceptor spacing. In M. Fahle & T. Poggio (Eds.), *Perceptual learning* (pp. 197-218). Cambridge, MA: MIT Press, Bradford Books.
- GIBSON, E. J. (1969). *Principles of perceptual learning and development*. New York: Appleton-Century-Crofts.
- GIBSON, E. J., & WALK, R. D. (1956). The effect of prolonged exposure to visually presented patterns on learning to discriminate them. *Journal of Comparative & Physiological Psychology*, **49**, 239-242.
- HALL, G. (1980). Exposure learning in animals. *Psychological Bulletin*, **88**, 535-550.
- HALL, G. (1991). *Perceptual and associative learning*. Oxford: Oxford University Press, Clarendon Press.
- HALL, G. (2003). Learned changes in the sensitivity of stimulus representations: Associative and nonassociative mechanisms. *Quarterly Journal of Experimental Psychology*, **56B**, 43-55.
- HALL, G. (2009). Perceptual learning in human and nonhuman animals: A search for common ground. *Learning & Behavior*, **37**, 133-140.
- HALL, G., & RODRIGUEZ, G. (in press). Factors determining the effects of associative activation on habituation. *Journal of Experimental Psychology: Animal Behavior Processes*.
- HONEY, R. C., & BATESON, P. (1996). Stimulus comparison and perceptual learning: Further evidence and evaluation from an imprinting procedure. *Quarterly Journal of Experimental Psychology*, **49B**, 259-269.
- HONEY, R. C., BATESON, P., & HORN, G. (1994). The role of stimulus comparison in perceptual learning: An investigation with the domestic chick. *Quarterly Journal of Experimental Psychology*, **47B**, 83-103.
- JAMES, W. (1890). *Principles of psychology*. New York: Holt.
- KRUSCHKE, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, **45**, 812-863.
- LAWRENCE, D. H. (1952). The transfer of a discrimination along a continuum. *Journal of Comparative & Physiological Psychology*, **45**, 511-516.
- LE PELLEY, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology*, **57B**, 193-243.
- LOVEJOY, E. (1968). *Attention in discrimination learning*. San Francisco: Holden-Day.
- MACKINTOSH, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, **82**, 276-298.
- MACKINTOSH, N. J., KAYE, H., & BENNETT, C. H. (1991). Perceptual learning in flavour aversion conditioning. *Quarterly Journal of Experimental Psychology*, **43B**, 297-322.
- MCLAREN, I. P. L., KAYE, H., & MACKINTOSH, N. J. (1989). An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 102-130). Oxford: Oxford University Press, Clarendon Press.
- MCLAREN, I. P. L., LEEVERS, H. L., & MACKINTOSH, N. J. (1994).

- Recognition, categorization and perceptual learning. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance XV: Conscious and nonconscious information processing* (pp. 889-909). Cambridge, MA: MIT Press.
- MCLAREN, I. P. L., & MACKINTOSH, N. J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior*, **28**, 211-246.
- MITCHELL, C., KADIB, R., NASH, S., LAVIS, Y., & HALL, G. (2008). Analysis of the role of associative inhibition in perceptual learning by means of the same-different task. *Journal of Experimental Psychology: Animal Behavior Processes*, **34**, 475-485.
- MOLLON, J. D., & DANILOVA, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision*, **10**, 51-58.
- MONDRAGÓN, E., & HALL, G. (2002). Analysis of the perceptual learning effect in flavour aversion learning: Evidence for stimulus differentiation. *Quarterly Journal of Experimental Psychology*, **55B**, 153-169.
- MUNDY, M. E., DWYER, D. M., & HONEY, R. C. (2006). Inhibitory associations contribute to perceptual learning in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, **32**, 178-184.
- MUNDY, M. E., HONEY, R. C., & DWYER, D. M. (2007). Simultaneous presentation of similar stimuli produces perceptual learning in human picture processing. *Journal of Experimental Psychology: Animal Behavior Processes*, **33**, 124-138.
- PAVLOV, I. P. (1927). *Conditioned reflexes* (G. V. Anrep, Trans.). London: Oxford University Press.
- PRADOS, J., ARTIGAS, A. A., & SANSÀ, J. (2007). Preexposure effects in the spatial domain: Dissociation between latent inhibition and perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **33**, 115-123.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- SCAHILL, V. L., & MACKINTOSH, N. J. (2004). The easy to hard effect and perceptual learning in flavor aversion conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, **30**, 93-103.
- SUTHERLAND, N. S., & MACKINTOSH, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- SYMONDS, M., & HALL, G. (1995). Perceptual learning in flavor aversion conditioning: Roles of stimulus comparison and latent inhibition of common elements. *Learning & Motivation*, **26**, 203-219.
- TSUSHIMA, Y., & WATANABE, T. (2009). Roles of attention in perceptual learning from perspectives of psychophysics and animal learning. *Learning & Behavior*, **37**, 126-132.

(Manuscript received August 5, 2008;  
revision accepted for publication December 23, 2008.)