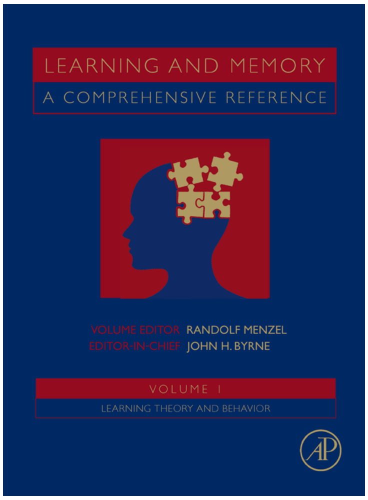

G. Hall. Perceptual Learning. In R. Menzel (Ed.), Learning Theory and Behavior. Vol. [1] of Learning and Memory: A Comprehensive Reference, 4 vols.
(J.Byrne Editor), pp. [103-122] Oxford: Elsevier.

1.07 Perceptual Learning

G. Hall, University of York, York, UK

1.07.1	**Introduction**	103
1.07.2	**Phenomena**	103
1.07.2.1	Simple Sensory Thresholds	104
1.07.2.2	More Complex Stimuli	105
1.07.2.3	Categorization	106
1.07.2.4	Taste and Smell	107
1.07.2.5	Acquired Distinctiveness	108
1.07.3	**Theoretical Issues**	110
1.07.4	**The Role of Associative Processes**	111
1.07.4.1	Acquired Distinctiveness and Acquired Equivalence	111
1.07.4.2	Unitization	113
1.07.4.3	Associative Inhibition	114
1.07.5	**Attentional Learning Processes**	116
1.07.5.1	Latent Inhibition and Associability Modulation	117
1.07.5.2	Habituation and Salience Modulation	118
1.07.6	**Conclusions**	119
References		120

1.07.1 Introduction

It is not the aim of this chapter to provide a review of all the work that has been published in recent years under the heading of perceptual learning. The aim is more specific and is guided in part by theoretical considerations; it is to set perceptual learning in the context of what we know about learning more generally – to assess the extent to which the phenomena of perceptual learning can be explained in terms of known learning processes, to identify aspects of perceptual learning that cannot be explained in this way, and to attempt to outline the nature of any new learning process that may be required.

Nonetheless, the first major section of the chapter will offer a selective review of what may be considered to be the most important recent (and some not-so-recent) experimental findings. This is necessary not only to provide the grist for the explanatory mill of later sections, but also to allow us to define the field. As we shall see, the range of phenomena that have been studied under the heading of perceptual learning is exceedingly broad. However, it is possible to discern a set of important issues that is common to most of them and that constitute the core features of perceptual learning. The attempt to explain – to understand the psychological mechanisms responsible for – these core features is dealt with in the next two major sections of the chapter. The first of these deals with the application of the principles of associative learning theory to perceptual learning effects; the second is concerned with the role of nonassociative (principally attentional) learning processes.

1.07.2 Phenomena

All the studies to be described in this section have been regarded, by their authors or by later commentators, as involving perceptual learning. The variety is impressive. All major sensory systems have been subject to study; vision perhaps predominates, but there are also many studies of hearing, touch, olfaction, and taste. The stimuli used have varied, from the apparently simple (e.g., a touch with a pointer), to the undeniably complex (e.g., pictures of human faces). The stimuli used have often been difficult to discriminate from one another – for simple stimuli because they are chosen to lie close together on the dimension of difference; for complex stimuli, largely because of the presence of a host of irrelevant, nondistinguishing features in each of the displays. But perceptual learning has also been investigated in experiments using stimuli that, on the face of things, seem readily discriminable one from another – one of the earliest and most influential studies in nonhuman animals (Gibson and Walk, 1959) looked at the effects of training with two simple geometrical shapes, a triangle and a circle. This example serves to make the further point that relevant phenomena have been studied in a range of species – our own species is best fitted for psychophysics, but experiments on laboratory animals have some advantages when it comes to investigating basic mechanisms of learning. Finally, although all the work to be considered concerns the effect of experience with stimuli on the subject's subsequent response to them, the exact form of the experience given has been varied – notably some experimenters have been concerned with the effects produced by mere exposure to stimuli, whereas others have given explicit training with feedback (also called knowledge of results, or reinforcement).

In an attempt to impose some order on this apparent chaos, the sample of experimental work reviewed next is organized under a set of convenient headings. It will be evident, however, that these do not form exclusive categories and that many of the experiments described could legitimately be placed under more than one of the headings (*See* Chapters 1.08, 1.13).

1.07.2.1 Simple Sensory Thresholds

The study of perceptual learning appeared early in the history of experimental psychology. No sooner had psychophysics been established as a coherent enterprise than its proponents began to study the effects of experience on the sensory threshold measures that were one of its primary concerns. In 1859 A. W. Volkmann (Fechner's brother-in-law; Woodworth and Schlosberg, 1954) published a study of the effects of practice on the two-point tactile threshold (Volkmann, 1859). He found that with practice, the ability of a subject to make the discrimination – two points or one – was dramatically increased; that after a hundred or so trials, the minimum distance required for a judgment of two points was reduced by about half. The effect was limited to the general area of skin on which the stimuli had been applied, except that positive transfer was observed when the test was carried out on the equivalent region of the other hand or arm.

We can say straightaway, on the basis of these results, that the conventional title for studies of this sort is potentially misleading. Although the stimulus may be simple, the mechanism responsible for the change in threshold is not. The fact that training failed to transfer to an adjacent patch of skin on the same arm indicates that the improvement is not a consequence of some general learning process (such as might result simply from familiarization with the procedure). But the fact that it did transfer to the other limb indicates that it is not a consequence of some change in the particular receptors that were stimulated in training – a more central mechanism must be involved.

Modern studies of difference thresholds have confirmed, for a range of stimuli, the essence of Volkmann's (1858) findings. Practice at the task (usually some version of a two-alternative forced-choice task in which the subject is presented briefly with two events, one after the other, and has to say whether they are the same or different) will produce a reduction in the magnitude of the difference that can be reliably detected. This is true for auditory frequency (e.g., Demany, 1985), for the orientation of visually presented lines (Shiu and Pashler, 1992), for hyperacuity (the ability to judge whether or not two line segments are colinear, e.g., McKee and Westheimer, 1978), for complex sinusoidal gratings (Fiorentini and Berardi, 1980), for the direction of motion of an array of moving dots (Ball and Sekuler, 1982; Liu and Weinshull, 2000), for the discrimination of visual texture (Karni and Sagi, 1991), and for many other tasks (see Fahle and Poggio, 2002). In most of these studies, the subjects were given feedback during training, but its role in producing the effect is unclear. Shiu and Pashler (1992) included subjects given no feedback, and these showed no within-session improvement, but this procedure was still capable of producing learning, as the subjects showed an improvement from one training session to the next.

As for Volkmann, the extent to which the effects of training transfer to other stimuli has been a focus of interest for modern experimenters. The intention has been to identify the stage in the perceptual system at which the training has had its effects – for instance, training that produces effects specific to the area of the visual field to which the stimuli were presented points to processes occurring at an early stage in the visual system, where retinotopic organization is still maintained. And, as was true for Volkmann's study, the pattern of transfer has turned out to be far from simple. For most of the visual tasks described, a degree of retinal specificity has been obtained, with performance falling to the starting level when the stimuli were presented in a different retinal location. But this does not mean that no transfer occurred – both Liu and Weinshull (2000) and Shiu and Pashler (1992) found that the new task was learned more rapidly than the original. And Ahissar and Hochstein (1997), using a version of the texture discrimination task of Karni and Sagi (1991), found transfer to a new retinal location when the initial training had been given on an easy version of the task (Liu and Weinshull, 2000, obtained a similar result for direction of motion). In other cases, the effect of training has proved specific to the particular stimuli used, but not to the receptors stimulated. Demany and Semal (2002) found that the effects of auditory training transferred readily to the other ear but not to a new discrimination involving a different frequency range (see also Amitay et al., 2005). And to return to the topic of tactile discrimination with which we began, Sathian and Zangaladze (1997) have demonstrated that the effects of training on a discrimination of the roughness of a pattern of ridges will transfer to a finger other than that used in training, but not to a new task in which the orientation of the pattern must be judged. The only conclusion justified by these observations at this stage is that although practice can facilitate performance on these simple discriminations, it would be foolish to conclude that these examples of a perceptual learning effect are to be explained solely in terms of processes occurring very early in the sensory/perceptual processing system.

Figure 1 Two of the pictures of objects used by Furmanski and Engel (2000). To render discrimination difficult, the contrast was reduced to 12.5%, and each presentation was followed by the mask shown at the bottom of the figure. Used with permission.

1.07.2.2 More Complex Stimuli

In a study directly inspired by those described in the previous section, Furmanski and Engel (2000) looked for a perceptual learning effect with more complex visual stimuli. These were degraded pictures of everyday objects (e.g., a telephone, a pencil sharpener; see **Figure 1**) presented very briefly. The subjects had to name the object and were told if they were right or wrong. Initially performance was poor, but training (800 trials a day over several days) produced a sizable reduction in threshold (in the exposure duration necessary for correct identification). This effect was specific to the pictures presented in training, but not to retinal location – transfer was good when the size of the image was changed. Enhanced discrimination when the stimulus set is familiar appears to be a quite general phenomenon. An example for a very different procedure (and species) is provided by Todd and Mackintosh (1990). Their subjects, pigeons, were presented with 20 pictorial slides in a training session, each being presented twice, in random order. They were rewarded for pecking at the first but not at the second presentation. The pigeons learned this discrimination, but did so less well when a new set of pictures was used each session than when the same set was used throughout. This result – better performance on a judgment of relative recency than of absolute novelty – may seem surprising, but it makes sense if we accept that prolonged experience with a given set of pictures will enhance the subject's ability to discriminate among them.

Intriguing as the examples just cited may be, popular interest in the phenomenon of perceptual learning is most readily evoked by description of the special skills shown by experts in dealing with even more complex stimuli – the experienced radiographer who can detect a tumor on an x-ray where the rest of us see only a meaningless blur (Myles-Worsley et al., 1988) or the chicken sexer who can make a determination after inspecting the pinhead-sized genital eminence of the day-old chick for less than 0.5 s (Biederman and Shiffrar, 1987). These are abilities acquired through experience. Knowledge of results may play a role in this (if only because the chicken sexer will soon hear about it if a large proportion of the hens turn out to be cocks), and often there will be explicit training in which an established expert instructs the novice in what to look for (see Biederman and Shiffrar, 1987). And however exotic these specialized skills may seem, there are certain areas in which, even in the absence of explicit instruction, almost all of us have acquired the status of experts – we are all of us experts in dealing with the visual cues responsible for face recognition and the auditory cues underlying our native language.

Despite occasional embarrassing mistakes, our ability to discriminate among people, largely on the basis of their facial characteristics, is impressive. That this ability depends on experience is indicated by the fact that children are less good at it than adults (see Chung and Thomson, 1995). Further evidence comes from the so-called own-race effect – superior discrimination when the faces belong to individuals from our own racial group (Malpass and Kravitz, 1969). This effect depends on the fact that we have (usually) had much more experience of such faces – the magnitude of the effect is much reduced in people who have had extensive interaction with races other than their own (Chiroro and Valentine, 1995). What might be thought as a parallel, own-language effect describes our ability to discriminate the speech sounds of our native language. Native English speakers readily distinguish (on the basis of differences in the third formant) between the phonemes /r/ and /l/, a task that native speakers of Japanese find exceedingly difficult (e.g., Goto, 1971). Japanese speakers, on the other hand, can make distinctions (according to changes in the second formant) within the category of sounds that English speakers regard as all being examples of /r/ (Iverson et al., 2003). Explicit training, in which Japanese subjects were given feedback after being required to distinguish between word pairs such as *lock* and *rock*, has been shown to enhance their discriminative ability (Iverson et al., 2005).

The examples discussed so far in this section have involved training in which feedback has been given; that is, the subjects have been told that their identification of a stimulus has been right or wrong. (This is explicitly arranged in most of the experimental studies, but something equivalent will occur in the natural environment as we learn the discriminations necessary for language or for the recognition of faces.) But perceptual learning effects can be obtained without such feedback, as a result of mere exposure to the stimuli (when this is arranged appropriately). Lavis and Mitchell (2006) required people to discriminate between pairs of checkerboards of the sort shown in **Figure 2**. When they are first presented with the one checkerboard followed, after a short interval, by another, people are poor at answering the question: Same or different? But performance was much enhanced by mere preexposure to the displays when this was organized such that the different stimuli were presented in alternation. Interestingly, preexposure, consisting of a block of trials with one stimulus followed by a block of trials with the other, was much less effective in enhancing subsequent discrimination. This outcome (to be discussed in detail later) is of interest as it accords with the influential analysis of perceptual learning offered by Gibson (1969), who emphasized the role of stimulus comparison in producing the effect (we may assume alternating preexposure is likely to foster the processes involved in comparison).

1.07.2.3 Categorization

The auditory discrimination described earlier is one that involves categorization; that is, faced with a range of different stimuli, the native English speaker learns to put instances of one set (which will differ

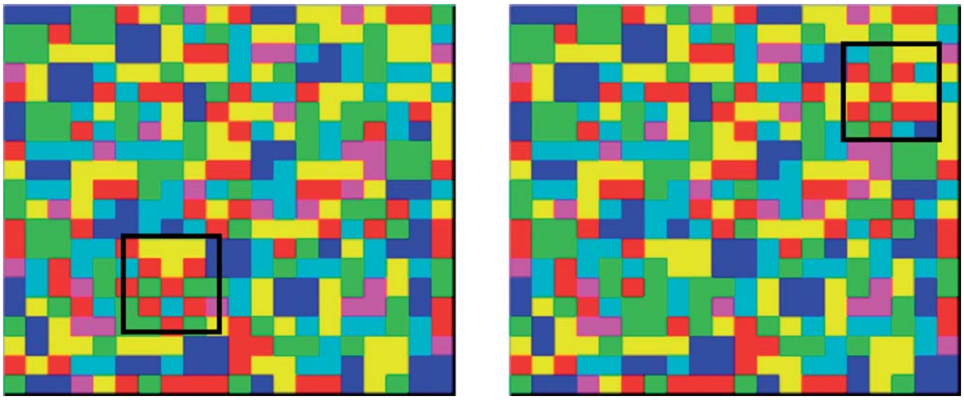

Figure 2 Examples of the colored checkerboard stimuli used by Lavis and Mitchell (2006). The unique elements are indicated (for the purposes of illustration only) by the black squares. Images courtesy of Y. Lavis.

among themselves in some respects) in one category (meriting /r/) and instances of another set in a different category (meriting /l/). The same is true of some of the visual tasks – one male chick will differ from another, but the experienced chicken sexer puts them all in the same category. An experimental demonstration of how such an ability can be acquired through experience comes from a study of face recognition by Quinn et al. (1999). In this experiment, the faces were pictures not of people, but of cats, and the participant's task was to sort them into two categories, male and female. The discrimination was difficult – initial performance was at chance level – but reinforced training with a subset of the pictures, those most easily identified as male or female, produced positive transfer to the ability to categorize other pictures.

Successful performance on a categorization task such as that used by Quinn et al. (1999) may seem to involve a reversal of the sort of perceptual learning effect that we have been concerned with so far – although discrimination between male and female is enhanced, the subjects appear to be less sensitive to differences among individuals that fall into a given category. But whether this is really the case requires explicit investigation. It is quite possible that the within-category discrimination was also enhanced – that, had they been asked, the trained subjects would have been better able to distinguish between Lucky and Widget, while still categorizing both as male. This is certainly true of human faces – we have no trouble in telling Ann from Zoe, and Andrew from Zach, while still distinguishing male from female.

The issue has been investigated experimentally, using complex artificial stimuli, by McLaren et al. (1994); see also McLaren, 1997). The stimuli used were based on the checkerboards shown in **Figure 3**. Changing a proportion of the elements from black to white, or vice versa, produced two sets of exemplars of these prototypes. Training consisted of a categorization task in which subjects learned to assign exemplars to type 1 or to type 2. In the test, two stimuli were presented side by side, and the subjects were reinforced for reliably choosing one rather than the other. They proved to be very good at this discrimination when the two stimuli were the original prototypes and also when the test stimuli were novel exemplars of these prototypes not used in the original training. But although the effect of preexposure was less profound in this case, the participants were also at an advantage when required to discriminate between two exemplars drawn from the same category. These effects are not confined to our own species. Aitken et al. (1996), using similar stimuli, generated essentially the same pattern of result in an experiment that used pigeons as the subjects.

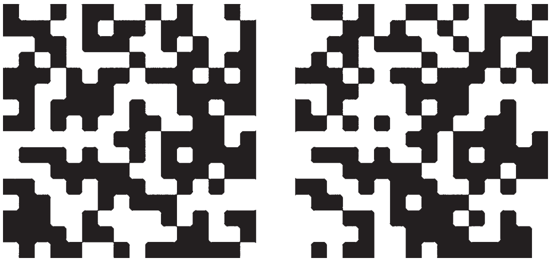

Figure 3 Checkerboard stimuli used by McLaren et al. (1994). Those shown were the prototypes. Discrimination was between exemplars, produced by changing a proportion of the elements of each prototype from black to white, or vice versa. Image courtesy of I. McLaren.

1.07.2.4 Taste and Smell

Unusually among psychologists, students of perceptual learning have paid almost as much attention to the chemical senses as to vision and hearing. Perhaps this derives in part from the fact that some of the most dramatic examples of acquired perceptual skills are found in these modalities. Foremost among these are the well-documented achievements of expert wine tasters (e.g., Soloman, 1990), but who can forget William James's description of "the blind-deaf mute . . . Julia Brace [who] is said to have been employed in the Hartford Asylum to sort the linen of its multitudinous inmates, after it came from the wash, by her wonderfully educated sense of smell" (James, 1890, pp. 509–510).

These are obviously very special cases, but evidence that an approach to skills of this sort can be established in any of us comes from experimental studies. Thus, Peron and Allen (1988) found that novice beer drinkers, who were initially unable to tell one brand from another, became able to do so after training in which they simply sampled a range of beers and reflected on the flavor qualities that came to mind. (Training with the specialist vocabulary of master brewers conveyed no special advantage; see also Melcher and Schooler, 1996.) Rabin (1988), who asked subjects to make same/different judgments after sniffing two unusual odors, found that discrimination was enhanced when the subjects had been given prior exposure to the odors (and in this case, training in which a distinctive label was attached to each odor during preexposure was found to help). Preexposure is not always beneficial, however. Stevenson (2001) gave training in which his subjects sniffed two odor mixtures (call them AX and BY). Subsequent discrimination between components of separate pairs (e.g., A vs. B) was good, but discrimination between components of a compound (e.g., A vs. X) was poor.

Studies of perceptual learning using animal subjects have made extensive use of tastes – adding a flavor to the drinking water of a thirsty rat ensures that the animal receives full exposure to the relevant stimulus, and the flavor-aversion conditioning technique provides an effective way of assessing discrimination. **Figure 4** shows the results of one study (by Symonds and Hall, 1995, Experiment 1) that made use of this procedure. All the rats received a test phase consisting of aversion conditioning with flavor A, followed by a test with flavor B (A and B were solutions of salt and sugar, rendered more similar by the addition of the sour taste of acid to each). In rats given no previous experience of the flavors (group W in the figure), the aversion conditioned to A generalized readily to B; that is, they failed to discriminate between A and B. The same was true of rats given prior exposure either to A or to B. But rats given prior exposure consisting of alternating presentations of A and B (group A/B in the figure) showed poor generalization (i.e., an enhanced ability to discriminate). As in the experiments by Lavis and Mitchell (2006), described earlier, this alternating arrangement turned out to be critical. A subsequent study (Symonds and Hall, 1995, Experiment 3) showed that preexposure consisting of a block of A trials followed by a block of B trials (or vice versa) did not produce the same enhancement. It is interesting to note that closely parallel effects have been obtained with human subjects given initial training in which they tasted the compound flavors saline-lemon and sucrose-lemon either on alternating trials or on separate blocks of trials. Subsequent same/different judgments were found to be more accurate in those given the alternating schedule during preexposure (Dwyer et al., 2004).

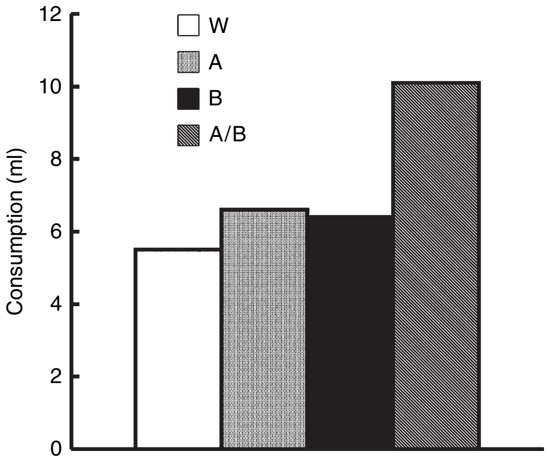

Figure 4 Group mean scores for consumption of flavor B, after aversion conditioning with flavor in the experiment by Symonds and Hall (1995). Before conditioning, different groups had received exposure to A, to B, to both (A/B), or just to plain water (W).

1.07.2.5 Acquired Distinctiveness

The experiments discussed so far have commonly used a procedure in which the subjects received explicit discrimination training, with feedback or knowledge of results being given. But (as the authors of several of these studies have noted), it can often be difficult to be sure of the source of the improved discriminative performance that is obtained in these conditions. Is there a change in the way in which the stimuli are being perceived, or is the learning occurring at a later stage in the sequence of processes that connect input to response? Transfer studies in which the stimuli are presented to a different set of receptors (e.g., at a different retinal location) constitute one way of addressing this issue. It is usually assumed that effects that fail to transfer cannot be a consequence of some general learning process, but must be specific to the stimuli used in training. An alternative strategy is to retain the original stimuli, but to require the subjects to learn some new discrimination on the basis of them. In this case, we are looking for positive transfer. Such transfer could not be based on knowledge of general task requirements (these having been changed) but must be a consequence of some learned change in the properties of the stimuli.

This experimental design was first introduced in the classic study of animal discrimination learning reported by Lawrence (1949). In outline, rats were trained initially on a food-rewarded choice discrimination between black and white. They were then shifted to a new task involving the same cues but with a different response requirement – when given two black cues, they were required to choose the left (for example), and when given two white cues to choose the right. The responses acquired in the first stage (e.g., to approach black and avoid white) will be irrelevant in this new task; nonetheless, Lawrence found positive transfer from stage 1 to stage 2, transfer, he said, that must depend on something that has been learned about the stimuli. The effect has been called the acquired distinctiveness of cues, the implication being that the initial training in which black and white were associated with different responses and different outcomes had rendered those stimuli more distinguishable. The idea that training in which cues are associated with differing outcomes will enhance their subsequent discriminability goes back at least as far as James (1890).

Lawrence's (1949) study was followed by a rush of similar transfer-of-training experiments, some with animal subjects (reviewed by Sutherland and Mackintosh, 1971) and many with human participants. The latter, for the most part, concentrated on a procedure in which participants learned to apply verbal labels in the first stage of training followed by a discrimination involving overt motor responses to the same stimuli in the second stage. It was usually found (see Hall, 1991, for a review) that stage-1 training facilitated learning of the second task, although whether or not this effect was a consequence of changes in the distinctiveness of the cues is open to debate. In many of the classic experiments (e.g., Battig, 1956; Gagné and Baker, 1950; Holton and Goss, 1956), comparison was made with a control condition given no stage-1 training, raising the possibility that the advantage shown by the experimental condition simply reflected some general facilitatory effect produced by the first stage of training. What is needed is to compare the effects of initial discrimination training with those of some control stage-1 procedure that will be equally effective in producing general transfer effects but that does not involve the consistent stimulus-outcome associations characteristic of the experimental condition.

One strategy, illustrated in a study by Goldstone (1994), is to use compound stimuli. In Goldstone's experiment, the stimuli were squares, differing in size and brightness. One aspect of the stimuli (e.g., their size) was irrelevant to the discrimination (based, e.g., on brightness) trained in stage 1. Thus, when it came to the test phase (involving either a further brightness discrimination or one based on shape), all subjects were familiar with the stimuli and had received discriminative pretraining. It was found that the test task was performed more readily by those subjects for whom the same dimension (brightness, in this example) was relevant in both stages. Analogous effects, which have been interpreted as reflecting an acquired enhancement of the distinctiveness of an entire dimension of stimulus variation, have been obtained in studies of animal discrimination learning (see, e.g., Mackintosh and Little, 1969).

An alternative strategy (also based on a design successfully used with animal subjects; see Bonardi et al., 1993) is presented schematically in **Table 1**. (At this point, we are concerned only with the first two stages shown in the table; the implications of Stage 3 will be taken up later). In this example, which comes from an experiment by Hall et al. (2003), people received stage-1 training with four different stimuli, four different geometrical shapes (A–D in the table). Two, A and B, were followed by one outcome (presentation of a red rectangle); two (C and D) by another outcome (a green rectangle). No overt response was required at this stage. Stage 2 consisted of a discrimination learning task in which the subjects had to learn to make one motor response rather than another to each of the shapes. Performance was good when the subjects were required to make different responses to stimuli that had been associated with different outcomes in Stage 1 (the consistent condition of the table) but was relatively poor in the inconsistent condition, when they had to make the same responses to shapes previously associated with different outcomes. Positive transfer in the consistent condition is what would be expected if cues associated with different outcomes had acquired distinctiveness.

It should be noted that with this experimental design (as with most others in this area; see Hall,

Table 1 Experimental design used by Hall et al. (2003)

Stage 1	*Stage 2*	*Stage 3*
Group Consistent		
A → red	A → left	
B → red	B → left	red → left/right?
C → green	C → right	green → left/right?
D → green	D → right	
Group Inconsistent		
A → red	A → left	
B → red	B → right	red → left/right?
C → green	C → left	green → left/right?
D → green	D → right	

Note: A, B, C, and D represent visual stimuli presented on a computer monitor; red and green refer to colored rectangles. Left and right refer to keyboard response required (left = backslash; right = forward slash). Feedback was given after responses in Stage 2. All subjects in a given group received all types of trial listed under a given stage of training. Source of data: Hall G, Mitchell C, Graham S, and Lavis Y (2003) Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *J. Exp. Psychol. Gen*. 132: 266–276.

1991), it is possible that the effect derives, in whole or in part, from negative transfer in the inconsistent condition – that associating cues with the same outcome renders them less distinctive (an effect referred to as acquired equivalence; Miller and Dollard, 1941). But it should also be noted that acquired equivalence in itself constitutes an example of perceptual learning, one that is worth our attention. Our topic is how experience can change the way in which things are perceived. Although in almost all the examples given so far, the change has been for the better (that is, discriminability has been enhanced by experience) there is no reason why this should always be the case, and our analysis of the phenomenon would be incomplete if it failed to encompass acquired reductions in discriminability.

1.07.3 Theoretical Issues

In all the experiments described earlier, the test task requires the subject to discriminate between two similar stimuli. The situation is presented schematically in **Figure 5**. Each circle represents the set of features or elements that define or constitute a particular stimulus, A or B. Each stimulus will have a set of unique features that are not found in the other (and these are represented by the areas containing the elements labeled a and b). The fact that A and B are similar is represented by the overlap – the area marked containing the c elements designates a set of features that they hold in common. Successful discrimination is evident when the subject shows the ability to make one response to A (i.e., ac) and a different response to B (bc). It follows that the job of a theory of perceptual learning is that of explaining how experience or training allows behavior to come to be controlled by the unique features (a and b), rather than by the common features (c). The behavior in question may be a gross overt movement, as when a rat approaches one stimulus object and avoids another; or it might be as minor as the verbal response of "higher" from a human participant presented with one of a pair of tones.

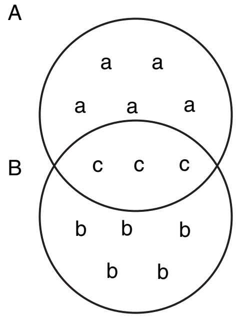

Figure 5 Each circle represents a stimulus (A or B) that is made up of a set of features (or elements). Some features are unique to a given stimulus (the a elements for A; the b elements for B); other features (c) are held in common and thus fall into the area of overlap of A and B.

The scheme shown in **Figure 5** can be applied quite generally. It may need slight modification when the stimuli are drawn from a simple continuum, such as tonal frequency – here a given tone might be regarded as consisting of the elements a, b, c, d (say); its neighbor by the elements b, c, d, e; its neighbor by the elements c, d, e, f; and so on – but the principle remains the same; discrimination between adjacent (similar) tones requires control by the unique elements that distinguish between stimuli.

This characterization prompts an attempt at a definition of perceptual learning. It is the learning process (or processes) that increases the effectiveness of unique stimulus elements and/or reduces that of common stimulus elements, thus facilitating discrimination between similar stimuli. (Although this will serve for almost all the cases discussed in the previous section, we should note the possibility of instances in which training reduces discriminability. For these we must assume that the effectiveness of common elements is increased, that of the unique elements is reduced, or both; *See* Chapter 1.08 for configural processing).

This definition may allow us to rule out, as instances of true perceptual learning, some of the processes that result in improved performance during practice on a discrimination task. It is often found, for instance, that performance on even the simplest difference threshold task can show a dramatic improvement on the early trials, and this improvement could well be a consequence of the participant learning to deal with the requirements of the procedure. If, for example, the subject is initially a little unclear as to which button to press for the high tone and which for the low tone, practice will establish the relevant associations and remove one obstacle to accurate performance – but this improvement would not be a consequence of a change in the effectiveness of the unique or common features of the stimuli and thus would not count as perceptual learning.

This is not to say, however, that we would want to dismiss what some (e.g., Liu and Weinshull, 2000) have called 'cognitive' learning, as a possible mechanism of perceptual learning. For example, McLaren and Mackintosh (2000) suggest that we should exclude from consideration cases in which participants, instructed to look for differences between stimuli, learn to focus on aspects that enable them to solve the task required of them. But surely a strategy as simple as learning to attend to or fixate on a particular part of a visual display deserves to be regarded as a mechanism of perceptual learning if it serves to increase the effectiveness of unique features (which happen to be located at a particular point in space). Again, Bruce and Burton (2002), discussing the discrimination of faces, question the extent to which enhanced discrimination depends on the acquisition of verbal labels for the different faces as opposed to reflecting (true) perceptual learning. But the distinction may not be a useful one. It has long been thought (e.g., James, 1890) that the acquisition of associates, such as verbal labels, might be a way of increasing the range and number of distinctive features activated by the presentation of a given stimulus; if this is so, the consequent improvement in discrimination becomes an example of perceptual learning under our definition.

The general point is that our proposed definition is silent as to the nature of the learning processes involved. The approach taken here contrasts with that sometimes taken by other students of perceptual learning. For example, Fahle and Poggio (2002) began their survey of the topic by ruling out, by means of their definition of perceptual learning, a number of possibilities that we would want to consider. Perceptual learning, they say, is independent from conscious experience and leads to implicit memory; it is not declarative, as it does not consist of consciously memorized facts or events; it is not associative, as it does not bind things together, and does not rely on the mechanisms of classical and operant conditioning; it differs from other forms of learning in that it principally involves functional and anatomical changes in primary sensory cortex. It may well be that some of the examples of enhanced discrimination that were described earlier in this chapter are a consequence of the type of learning envisaged by Fahle and Poggio (and we will try to identify them in subsequent sections of the chapter). But declining to subscribe to such a restrictive definition leaves us free to consider a range of other possible mechanisms, including several (such as various cognitive, associative, and attentional learning processes) that have been well studied in other contexts. We begin by considering the extent to which the associative analysis of learning can supply an explanation for perceptual learning phenomena.

1.07.4 The Role of Associative Processes

It may seem odd to give pride of place to a learning process that some have emphatically asserted is not responsible for perceptual learning effects; thus Gibson and Levin (1975, p. 23) wrote: “this simple and ancient notion does not work for perceptual learning, because what is learned [in perceptual learning] is not addition of something but rather extraction of something.” There are, however, at least two good reasons for doing so. First, associative learning theory, in its modern form (see, e.g., Wagner, 1981; Mackintosh, 1983), provides by far the best worked out and most comprehensive account of basic learning mechanisms, and it seems a sensible first step to attempt to explain some (supposedly) new form of learning in terms of what we already know about learning more generally. Second, one of the earliest attempts to explain an instance of perceptual learning was, in fact, precisely in terms of the notion that it depended on the associative ‘addition of something.’

1.07.4.1 Acquired Distinctiveness and Acquired Equivalence

Figure 6 (top part) presents a schematic version of the associative account offered by James (1890) for the acquired distinctiveness of cues. Recall that in this procedure, discrimination training, in which the cues are followed by different outcomes, enhances their subsequent discriminability. **Figure 6** shows the associations assumed to be formed when two similar cues (A and B) have been given training in which each has become linked to a different associate (X and Y); the associates are less similar (they share few common features) than are A and B. Discrimination between A and B prior to training will be difficult as they share many common elements. But the formation of the associative links means that presentation of A will produce associative activation of the representation of X, and presentation of B will associatively activate Y. As a result, discrimination between A and B will be enhanced because the proportion of common features present in the overall patterns of activation produced by these stimuli will

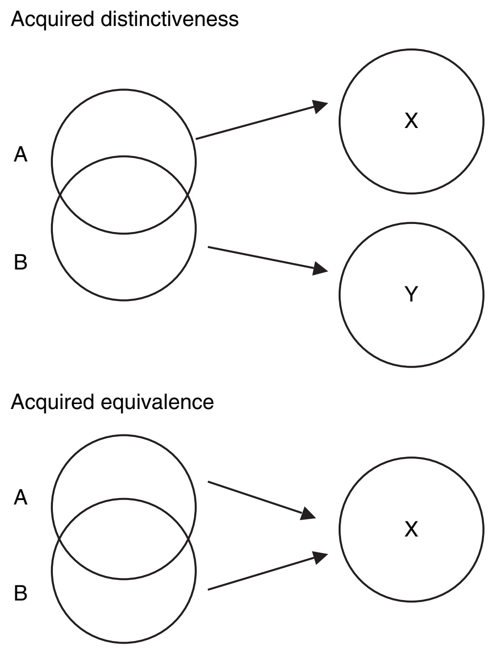

Figure 6 Associative structures in acquired distinctiveness and acquired equivalence. The overlapping circles represent two similar stimuli, A and B (see Figure 5); arrows represent associative links. In the acquired distinctiveness case, A and B have formed associations with quite different stimuli (X and Y). In the acquired equivalence case, both have become associated with the same stimulus.

be low, given the distinctiveness of their associates. The lower part of **Figure 6** shows the situation for the acquired equivalence procedure in which the two stimuli are given the same associate. Here the proportion of common elements is increased, and discriminability should go down. James focused on the case in which the associates were distinctive names, but the analysis applies (and the effect is found) when other events are used (as in experiments with animal subjects, or that presented in **Table 1**).

That stimuli can acquire new, and potentially distinctive, properties by associative means is nicely demonstrated by studies of odor perception. It has been noted that an odor such as vanilla is often described as smelling sweet, even though it is in itself tasteless. The suggestion that this quality is acquired as a result of associative learning (vanilla is often present in sweet foods) is supported by the results of a study by Stevenson et al. (1998), who gave subjects presentations of a novel odor (e.g., lychee) along with a sweet taste. Their subjects started to describe lychee as smelling sweet, a phenomenon that Stevenson et al. referred to as learned synesthesia. (Parallel effects have also been found with sour and bitter tastes.)

Whether associations of this sort are actually responsible for acquired distinctiveness effects is another matter. Gibson (1969), in her influential discussion of the topic, argued that although associations might indeed be formed during training designed to establish acquired distinctiveness, they did not provide the mechanism for the effect. Rather, the role of discrimination training was simply to ensure that the subjects concentrated on the stimuli, noting their similarities and differences, with the result that there was an increase in the perceptual effectiveness of (and attention paid to) their intrinsic distinctive features. This is what was meant by the phrase ‘extraction of something,’ rather than ‘addition of something.’ The evidence that is currently available to us does not allow a clear choice between the alternatives, but suggests, rather, that both processes may play a role.

Evidence that the associations formed during initial training can influence subsequent discrimination performance comes from the experiment by Hall et al. (2003), outlined in **Table 1**. Why, in Stage 2, did subjects in the inconsistent condition find it difficult to assign different responses to cues (such as A and B) that had shared a common associate (red) in Stage 1? According to the associative theory, this is because the representation of red was activated by both A and B during Stage 2. When subjects learned to respond left to A, this response would come under the control both of the cue actually presented (A) and its associate (red). Small wonder, then, that they found it difficult to respond right to B, given that its associate already controlled a tendency to make a different response. For subjects in the consistent condition, on the other hand, the response tendency acquired by red on a trial with A would allow the correct response to emerge immediately when B was presented. In a final test (Stage 3 of the table), subjects were asked to choose left or right when presented with the colors used as associates in Stage 1. Those in the consistent condition made appropriate choices (choosing left for red and right for green, in our example), as would be expected if the associatively activated representations of these colors had been involved in the discrimination test of Stage 2.

A reason to think that associative processes are not wholly responsible for acquired distinctiveness effects comes from the observation that mere exposure to a pair of similar stimuli (if this is appropriately arranged – as we have seen, alternating presentations that give the opportunity for comparison to occur are particularly effective) can enhance subsequent discriminability. That is, distinctiveness can be acquired when the preexposure procedure seems to preclude the formation of links with distinctive associates. For humans, this argument may not be particularly convincing – members of this species may well provide their own associates (names or labels), no matter what the experimenter tries to arrange. But the effect is also seen with animal subjects (rats given alternating presentations of two flavors are better able subsequently to discriminate between them). If another, nonassociative, perceptual learning process is at work during mere exposure, might it not also be active when the exposure phase involves explicit discrimination training? Evidence in favor of this view comes from a study by Bonardi et al. (2005), who modified the procedure used by Hall et al. (2003) to rule out associative mediation and yet still found an acquired distinctiveness effect. Bonardi et al. concluded that the preexposure procedure had, in addition to establishing associations, modified the perceptual effectiveness of the various elements of the stimuli (specifically had enhanced the attention controlled by their distinctive features). The learning mechanisms that might be responsible for such an effect will be taken up in the final major section of this chapter, after we have considered the contribution of other associative mechanisms to perceptual learning effects.

1.07.4.2 Unitization

It has frequently been suggested (e.g., Goldstone, 2000; McLaren et al., 1989) that a process of unitization is, in part, responsible for perceptual learning effects. The idea has been expressed in a variety of different ways, but the central notion is that exposure to a complex and multifaceted stimulus will result in the formation of a unitary representation of that stimulus, in which the various features are somehow bound together. The concept has been much used in the study of complex visual perception (with its talk of object representations, face representations, and so on) but can readily be applied to other modalities (it forms the basis of the account of odor perception proposed by Stevenson and Boakes, 2003), and to seemingly simple stimuli (even a pure tone, for example, has many features – frequency, intensity, location, duration, etc).

When it comes to specifying the learning mechanism responsible for unitization, the only developed proposal has been in terms of associative processes (see Goldstone, 1998). Presentation of a complex stimulus will, it is assumed, activate a set of units that correspond to its various constituent features. Concurrent activation of these units will lead to the formation of a network of excitatory links among them, and it is this network, in the simplest interpretation of the idea, that is taken to constitute the unitized representation. In terms of the diagram of **Figure 5**, experience of stimulus A will establish connections among all the a elements. (It will also, initially, allow the formation of a–c connections too; but because the c elements may also be activated in the absence of the a elements, when a stimulus such as B is presented, these will be weakened and eventually drop out of the picture.) A slightly more elaborate version of this analysis supposes that experience of the stimulus results in the formation of a separate configural unit (see, e.g., Pearce, 1994) that is not directly activated by any feature of the stimulus itself but that comes to receive its input from the units that correspond to those features.

The formation of a simple network of excitatory associations is enough in itself to explain a range of perceptual learning phenomena. As we have already seen, the learned synesthesia of Stevenson et al. (1998) depends on associations formed between two aspects of a compound stimulus (its taste and smell). And our difficulty in discriminating among the aspects of a previously experienced compound odor (Stevenson, 2001) can be explained in similar terms. When people have sniffed the AX compound, excitatory associations will form between A and X. Discrimination between A and X will be poor because the presentation of A will produce associative activation of the representation of X, and presentation of X will associatively activate A. It is usually assumed that the state produced by associative activation of a representation will be distinguishable from that produced by direct activation (if nothing else, the intensity of activation is likely to be less; see Hall, 1996), and accordingly, discrimination between A and X should still be possible. But the existence of association between A and X will increase the similarity of the overall patterns of activation elicited by A and X and render the task more difficult.

Further support for this interpretation of unitization comes from its ability to explain the finding that preexposure to a complex stimulus or event can enhance the ease with which that complex is subsequently learned about. An example from animal conditioning is supplied by Kiernan and Westbrook (1993), who gave rats a few minutes of exposure to a distinctive context prior to a session in which the context was paired with shock. On a subsequent test, these rats showed more evidence of learned fear in that context than did rats not given preexposure. (See also Fanselow, 1990; Bennett et al., 1996.) The only additional assumption needed to deal with this finding is that the animal's capacity to process all aspects of a complex stimulus will be limited, so that only a subset will be sampled at any one time. As the rat explores the context during preexposure, it will sample a range of contextual features, and connections will form among them. On the subsequent conditioning trial, only some of these features will be sampled, and only these will become associated with shock. A different set may be sampled on the test trial, but for animals given preexposure, the conditioned response should still be evoked, as those sampled on the test will be able to activate those that formed links with the shock during conditioning. Although presented in modern associative terminology, the central idea is essentially that popularized long ago as redintegration (by Hollingworth, 1928).

It will be apparent that the principle illustrated by this example from fear conditioning in the rat will be applicable to any case in which animals (including people) are given exposure to a stimulus containing more features than can be processed all at once. Appropriate response to a complex visual event (such as putting a name to a face) will be able to proceed more rapidly if inspection of one part of the display is able to activate representations of other features, or to activate a configural unit that is connected to the response-output mechanism. In the case of faces (and indeed other complex visual stimuli), it must be assumed that the configural unit is sensitive not just to the co-occurrence of various features but also to their spatial relationships. One of the classic findings from studies of face recognition is the inversion effect – the finding that faces are so much less well recognized when upside down; or, put another way, are especially well dealt with when they are the usual way up (Yin, 1969). This phenomenon is explained by assuming that experience of a face in the normal orientation establishes a configural unit that encodes the spatial relationships of the various features; when the face is inverted, these relationships are disrupted, and the unit will not be activated. Evidence for this interpretation comes from a study by Gauthier and Tarr (1997), who required people to learn the names of nonsense objects of the sort shown in **Figure 7**. An inversion effect was found for these stimuli too – performance was poor when the stimuli were presented upside down – but only for people who had received extensive prior experience with the cues. The effect is specific not to faces but to complex familiar visual cues.

Figure 7 Examples of the figures (referred to as Greebles) used by Gauthier and Tarr (1997) and of the names that subjects had to learn to apply to them. Used with permission.

1.07.4.3 Associative Inhibition

McLaren et al. (1989; see also McLaren and Mackintosh, 2000) have pointed out that standard associative theory predicts that certain schedules of exposure to a pair of similar stimuli (such as A and B of **Figure 5**) will allow the formation of other associations, in addition to those considered so far. They were concerned in particular with a preexposure schedule in which A and B were presented in alternation. Gibson (1969) has argued that a process of comparison plays an important part in the enhancement of discrimination that follows exposure to similar stimuli. The exact nature of the comparison process was not specified, but it can be agreed that alternating presentations are likely to enhance its operation, making this schedule of special interest.

We have already discussed how exposure to the stimuli will allow unitization to occur, with the various elements of each stimulus becoming linked together. McLaren et al. (1989) noted that the within-stimulus links formed would include excitatory connections between common and unique elements (which we may summarize as c–a, for stimulus A, and c–b, for stimulus B). As a consequence, presentation of A would, by way of the c–a association, be able to activate the representation of the unique features of B (the b elements); similarly, presentation of B by way of the c–a link would activate the a elements. On the face of things, therefore, preexposure might be expected to hinder subsequent discrimination, as these excitatory links would render the patterns of activation produced by A and B more similar to one another. Such indeed would be the outcome but for another factor that comes into play when the alternating schedule is used – this schedule will allow the formation of inhibitory associative links. According to standard associative theory (e.g., Wagner, 1981; espoused by McLaren and Mackintosh, 2000) an inhibitory link will form between a directly activated stimulus representation and one that is activated associatively. The alternating schedule ensures that the subject experiences a sequence of trials in which a is associatively activated in the presence of b, and b is associatively activated in the presence of a. In these circumstances, inhibitory associations will form between a and b. Activation of an inhibitory link is assumed to oppose the effects of excitatory influences acting on a given representation, restricting the ability of that representation to be activated.

This analysis can supply an explanation for the results of the experiments described previously, in which discrimination was enhanced by intermixed preexposure (compared to a control condition in which the stimuli were presented in separate blocks of trials). Consider the experiment by Lavis and Mitchell (2006). When asked to make same–different judgments about stimuli like those shown in **Figure 2**, the excitatory links formed by subjects in the control condition mean that each stimulus would tend to activate, to some degree, the unique features of the other. Presentation of A would activate a and c directly and b associatively; presentation of B would activate b and c directly and a associatively. That is, all the same units would be activated in each case, and the judgment would therefore have to be made in terms of degree (or type) of activation. For subjects given intermixed exposure, on the other hand, inhibition between a and b would mean that A would activate only a and c, and B would activate only b and c, qualitatively different patterns of activation.

The difference between intermixed and blocked preexposure schedules has also been obtained in studies of animal conditioning (e.g., Symonds and Hall, 1995; see also Bennett and Mackintosh, 1999; Mondragón and Hall, 2002; Dwyer et al., 2004), and the associative inhibition analysis applies readily to this case too. In these experiments, subjects received conditioning with stimulus A followed by a generalization test with stimulus B. The poor discrimination (good generalization) produced by blocked exposure is explained in terms of the within-stimulus excitatory links formed during preexposure; specifically when tested with B, the c–a link will allow this cue to activate a stimulus representation (a) that had acquired associative strength during conditioning with A. For the intermixed condition, on the other hand, the inhibitory association between b and a established during preexposure would prevent activation of a on the test trial, eliminating its contribution to responding and thus reducing the degree of generalization.

The role of associative inhibition as a mechanism of perceptual learning deserves our serious attention because it provides one of the few fully worked out accounts of the processes by which stimulus comparison might have its effects. It remains to be determined, however, whether or not this theoretical possibility is, in fact, responsible for the effects observed. There is some supportive evidence from a study by Dwyer et al. (2001; see also Dwyer and Mackintosh, 2002). Rats were exposed to intermixed presentations of the compound flavors, AX and BX, where A and B represent unique flavors, and X an explicitly added common element. According to the associative principles outlined earlier, alternating presentations of AX and BX should establish inhibition between A and B (just as alternating presentations of ac and bc establish inhibition between a and b). Dwyer et al. found that after extensive intermixed exposure, animals were retarded in learning an excitatory association between A and B – what would be expected if the prior training had established inhibitory associations between these cues. It seems likely that associative inhibition plays a role in the perceptual learning effect under these training conditions.

It should be noted, however, that this example of perceptual learning (better discrimination after intermixed than after blocked preexposure) can be obtained in the absence of associative inhibition. Inhibitory learning can take many trials to develop (Dwyer et al., 2001, gave extensive initial training), but the perceptual learning effect can be obtained after just a few preexposure trials, well before there is any evidence of inhibition between A and B (Artigas et al., 2006). What is more, the effect can be observed when the training procedure is modified so as to preclude the formation of inhibitory links. **Table 2** (Experiment 1) shows the design of an experiment (Blair and Hall, 2003, Experiment 1a) devised to demonstrate the basic effect using a within-subject design and a conditioning procedure. The subjects (rats) received intermixed preexposure to the compound stimuli AX and BX and a separate block of trials with CX. A response subsequently conditioned to AX was found to generalize less readily to BX (the stimulus presented intermixed with AX during preexposure) than to the control stimulus CX. This result in itself is compatible with the suggestion that the

Table 2 Experimental designs used to investigate the effects of preexposure on flavor discrimination in rats

	Preexposure	*Conditioning*	*Test*
Experiment 1	AX/BX _ CX	AX +	BX and CX
Experiment 2	AX/BX _ CX	X +	BX and CX
Experiment 3	X/BX _ CX	AX +	BX and CX
Experiment 4	AX/BX _ CX	—	B+ or C+

Note: A, B, C represent flavors that could be presented in compound with flavor X. In preexposure, AX (or X, in Experiment 3) was presented in alternation with BX. The CX compound was presented in a separate block of trials. In the conditioning phase of experiments 1–3 and the test phase of Experiment 4, flavors were presented along with (+) an aversive reinforcer. The test phase measured the extent of the aversion shown to the test stimuli (a generalized aversion in the case of experiments 1–3). Experiments 1 and 2 were fully reported by Blair CAJ and Hall G (2003) Perceptual learning in flavor aversion: Evidence for learned changes in stimulus effectiveness. *J. Exp. Psychol. Anim. B*. 29: 39–48; Experiment 3 by Hall G, Blair CAJ, and Artigas AA (2006) Associative activation of stimulus representations restores lost salience: Implications for perceptual learning. *J. Exp. Psychol. Anim. B*. 32: 145–155; and Experiment 4 by Blair CAJ, Wilkinson A, and Hall G (2004) Assessments of changes in the effective salience of stimulus elements as a result of stimulus preexposure. *J. Exp. Psychol. Anim. B*. 30: 317–324.

presence of B in the test inhibits activation of the representation of the (conditioned) A element. But this argument cannot apply to the modified design shown as Experiment 2 in the table (Blair and Hall, Experiment 5a). Here the same result was found (less generalization to BX than to CX) despite the fact that conditioning was given to the X element alone. If A has not been conditioned, then any ability that B might have to inhibit activation of the A representation would be irrelevant to the outcome of the procedure. The experiment presented as Experiment 3 in the table (Hall et al., 2006) makes the same point in a different way. Here the intermixed preexposure procedure involved alternation, not of AX and BX, but of BX and X alone. Obviously, inhibition between A and B cannot be established with this procedure. Nonetheless, conditioning to AX was still found to generalize less well to BX than to CX on the test (see also Rodriguez and Alonso, 2004).

The new results just described serve to support Gibson's (1969) suggestion that a preexposure procedure that allows comparison between two similar stimuli (such as alternating presentations of AX and BX) is particularly effective in enhancing discrimination between them. They also show that associative inhibition mechanism can supply only a partial explanation for these effects. Gibson's own interpretation was that comparison served to enhance the perceptual effectiveness of the distinguishing features of the stimuli (A and B in this case). This notion can help explain the results produced by the experiments summarized in **Table 2**. In these experiments, performance on the generalization test will be largely determined by the response controlled by the X element, the response that was established during the conditioning phase. To the extent that the presence of another element (such as B or C) detracts from the ability of the animals to perceive stimulus X, the magnitude of the response will be reduced. It follows that if alternating preexposure enhances the perceptual effectiveness of the B element, this element will be better able to interfere with the ability of X to evoke its response on test and generalization will be restricted – the result obtained. What we need to consider now, therefore, is the mechanisms by which the perceptual effectiveness of stimuli might be modified; this issue is taken up in the next section of the chapter.

1.07.5 Attentional Learning Processes

The only learning process utilized so far has been one that results in the formation of links (excitatory or inhibitory) between the central representations of stimulus elements. Perhaps surprisingly, this notion has proved helpful in explaining some perceptual effects. But however powerful an explanatory tool this "simple and ancient notion" may be, it is not, on its own, enough to explain even simple associative learning. A number of learning theorists (e.g., Mackintosh, 1975; Pearce and Hall, 1980; McLaren and Mackintosh, 2000) have argued that the associative principle needs to be supplemented by another learning process, one that is capable of changing the properties of the stimulus representation, modulating its sensitivity to activation, for example, or modulating the readiness with which it will enter into association.

Because these theories are usually described as involving a process of attentional learning, it would be useful to clarify what is meant by attentional in his context, as the use of the term (which is really more of a chapter heading than a well-defined psychological construct) can vary widely. It is not meant to indicate a form of learning that occurs only when learners focus their attention on the task at hand (a characteristic that, in any case, would be difficult to identify in the experiments using animal subjects that will be considered shortly). Rather, it indicates a form of learning that modifies the processing that a stimulus will receive. As for the nature of the modification, two principal suggestions have been advanced. One is that experience might change the effective salience of a stimulus (making a dim light function as if it were bright or a loud noise as if it were soft) and thus modify the ability of the stimulus to command attention. The other, not necessarily alternative, suggestion is that experience might change the associability of a stimulus, the readiness with which it will be learned about. (A significant event might deserve attention, even if its salience is low.) Both of these possibilities merit the description attentional, but it will be noted that they have somewhat different implications, and where necessary, they will be distinguished in what follows.

1.07.5.1 Latent Inhibition and Associability Modulation

Prior exposure to an event that is to be used as the conditioned stimulus (CS) in a classical conditioning procedure produces a marked retardation in the subsequent rate of learning. The source of this phenomenon (known as latent inhibition; Lubow, 1989) is still a matter for debate (some possibilities will be considered shortly). For our present purposes, however, we may simply note that mere exposure to a stimulus can produce a reduction in the readiness with which it can be learned about, and then go on to explore the implications of this fact of perceptual learning.

McLaren et al. (1989) have proposed an interpretation of perceptual learning effects in which latent inhibition plays an important part. They point out that it is important to distinguish between the ease with which a stimulus can be learned about and the ease with which it can be discriminated from other similar stimuli (which is our major concern). A process that reduced the former might enhance the latter. Consider the stimuli of **Figure 5**. A response conditioned to A will generalize to B (i.e., a failure of discrimination will occur) because the common (c) elements acquire strength during conditioning with A (ac) and are present in the test stimulus B (bc). Prior exposure to A will reduce generalization (enhance discrimination) because the c elements will suffer latent inhibition and thus acquire little strength during conditioning. This effect will be most marked if the subjects are given preexposure to both A and B, as the c elements will be present in both types of exposure trial, thus having twice the opportunity to acquire latent inhibition. Generalization should, therefore, be particularly weak after preexposure to both A and B – just the result obtained by Symonds and Hall (1995) and shown in **Figure 4**.

This simple notion generates an interesting prediction that has received experimental support. A perceptual learning effect (enhanced discrimination after preexposure) should only be found when the stimuli share a substantial number of common elements. When the stimuli are very different (as, in the limiting case, when A consists only of a elements and B only of b elements), latent inhibition of the c elements can play no part, and discrimination between A and B will be poor, as the latent inhibition suffered by the a and b elements will retard the acquisition of the (different) responses required to these stimuli. Trobalon et al. (1991; see also Prados et al., 1999) have demonstrated this result in a study of maze learning in rats. When the two maze arms that the rats had to choose from were very similar, preexposure to these arms facilitated discrimination learning; when the arms were made distinctively different, preexposure hindered learning.

Latent inhibition can provide an explanation for the result reported by McLaren et al. (1994) that people who had learned to assign checkerboard patterns (**Figure 2**) to different categories showed an enhanced ability to discriminate between new examples drawn from the same category. This result is unexpected, given that examples from the same category will have a common associate, in that both will elicit the same category label. Associative processes might be expected act to hinder discrimination between stimuli that have a common associate (the acquired equivalence effect). But this is to reckon without latent inhibition. McLaren et al. (1994) point out that during initial categorization training, the features common to all exemplars of that category occur on every trial. These features will therefore suffer extensive latent inhibition. Performance on the within-category discrimination will be facilitated, as this task requires precisely that the choice response should come under the control of features that distinguish the displays rather than features they hold in common.

The arguments just advanced hold whatever the mechanism of latent inhibition. Of the various possibilities (see Hall, 1991, for a review), perhaps the most widely accepted is that it reflects a loss of associability, this being expressed in formal terms as a reduction in the value of a stimulus-specific learning rate parameter (symbolized alpha in the influential learning model of Rescorla and Wagner, 1972; see Pearce and Hall, 1980). To adopt this interpretation raises a further interesting possibility. Latent inhibition itself involves only reduction in associability, but might it not be possible for the alpha-value of a stimulus to be increased under appropriate conditions? Mackintosh (1975) has adopted this proposal and devised a theory in which the associability of a stimulus is held to increase as a result of training in which it is a reliable predictor of its consequences (see also Kruschke, 2001). Direct tests of the validity of this proposal have generated mixed results (e.g., Hall and Pearce, 1979; Le Pelley, 2004). But if it could be confirmed, it would usefully extend the explanatory reach of the associability concept. In particular, it could supply an explanation for acquired distinctiveness effects. The acquired distinctiveness training procedure is one in which the subject experiences each of the critical stimuli in reliable association with another event (in **Table 1**, for example, A reliably predicts red, and B reliably predicts green). In addition to any associations that may be formed, Mackintosh's theory says that the associability of A and B will go up under these conditions. Subsequent discrimination between these cues would be enhanced, even in circumstances in which associative mechanisms do not seem to operate (Bonardi et al., 2005).

Whatever the fate of Mackintosh's (1975) theory of associability change, it will be evident that the basic latent inhibition process plays an important part in many perceptual learning effects. There is, however, one critical version of perceptual learning that defies explanation in terms of latent inhibition. This is the well-established, and already much-discussed, finding that discriminability is especially enhanced when the subject is able to compare the stimuli during preexposure. **Table 2** (Experiment 1) presents a simple experimental demonstration of the effect. In this experiment, the subjects can compare A and B during preexposure (they are presented on alternate trials), but will be less able to compare C with the others, as this stimulus is presented on a separate block of trials. But because the subjects experience the critical cues, A, B, and C, the same number of times, all three cues should acquire latent inhibition to the same extent. There are no grounds, therefore, for the latent inhibition account to predict the result obtained – poorer generalization from AX to BX than to CX. Further analysis of this finding is one of the topics of the final section of the chapter.

1.07.5.2 Habituation and Salience Modulation

Whatever other factors may play a part, there is no doubt that the effectiveness of a stimulus depends on its intensity. A strong stimulus will normally elicit a more vigorous response than a weak one (we show a bigger startle response to a loud noise than a soft one); associative learning occurs more rapidly when the events to be associated are intense. Formal theories of these phenomena (see, e.g., Hall, 1994) incorporate a notion of salience, a parameter associated with each stimulus and set by its intensity.

Stimulus salience will influence performance on the tests used in studies of perceptual learning. Subjects will be best able to discriminate (on a same-different test, say) between A and B when the unique features (a and b) are intense, and the common features (c) are not. And generalization between such stimuli will be poor, as the a element will dominate during conditioning with A, restricting the acquisition of control by the c element, and the b element will dominate on test with stimulus B, restricting the ability of c to influence performance. In most studies of the topic, we use stimuli with nonsalient unique features and salient common features; that is, we study the effects of experience on discriminations that are difficult. Perceptual learning effects would be obtained, then, if experience with stimuli was capable of boosting the effective salience of the unique features of stimuli (or of lowering that of common features, or both). What evidence is there that effective salience can change? We have discussed how simple exposure to repeated presentations of a stimulus can produce a loss of associability (latent inhibition), but there is reason to think that this procedure can also bring about a change in effective salience.

Repeated stimulus presentation results in habituation – the waning of the response unconditionally elicited by that stimulus. Explanation of this simple phenomenon turns out to be surprisingly complex (see Hall, 1991, for a review in the context of perceptual learning). But what we need to note for our present purposes is that the habituation procedure makes a salient stimulus behave like a less-salient one. After extensive habituation training, the startle response evoked by a loud noise will be much the same as the (weaker) response evoked by the first presentation of a softer noise. The habituation effect is most easily observed with motivationally significant events, as these evoke obvious responses; but the learning process responsible for it presumably operates for any stimulus, including those used as cues in experiments on perceptual learning. For a few of these, the effect can in fact be observed directly. Rats show neophobia to foods (that is, they are reluctant to consume a substance with a novel taste), and habituation of neophobia is commonly observed over the preexposure phase of perceptual learning experiments using flavor stimuli (Blair et al., 2004). Observations like these make it a reasonable presumption that the (unobservable) response evoked by presentation of a checkerboard, say, will also undergo habituation, or, in other words, that these stimuli too will lose effective salience with repeated exposure.

Blair et al. (2004) have investigated the role of salience modulation in perceptual learning, focusing on the differing effects of intermixed and blocked preexposure. Recall for the experiment shown as Experiment 1 in **Table 2** that rats consume more of BX than of CX on test, and that this difference is explained by the fact that at the end of preexposure, B has greater effective salience than C (the more salient a cue, the more it will interfere with expression of the response controlled by X). Blair et al. (2004, Experiment 3) tested the salience of B and C using the design shown in as Experiment 4 in **Table 2**. After preexposure, some rats received conditioning trials with B alone as the conditioned stimulus (CS); others received C alone as the CS. Acquisition occurred more rapidly to B than to C, as would be expected if B were higher in salience than C. In a further study, Blair et al. simply monitored the unconditioned response evoked by B and C at the end of preexposure. The neophobic reaction evoked by these flavors was found to have habituated to some extent over the course of preexposure, but it was still observable, particularly for flavor B. Thus the effective salience of both B and C was reduced by preexposure, but critically, the reduction was less for the cue presented in alternation with a similar cue in preexposure.

The conclusion that emerges from these and related studies (Hall, 2003; Hall et al., 2006) is that mere exposure to a stimulus will cause a loss of effective salience, but that with some schedules of preexposure, this loss can be attenuated or reversed. The critical arrangement appears to be one in which the cue in question is presented in alternation with another similar cue. Why this schedule should have the effects it does is not yet clear. There is some evidence from the experiments by Hall et al. (2006) to suggest that an important feature of this schedule is that on each trial the subject is likely to be (slightly) surprised at the omission of one of the unique features and at the occurrence of the other. (With the blocked schedule the same stimulus occurs trial after trial.) It seems plausible that an event that evokes surprise might also maintain its salience, but the precise learning mechanisms that might underlie such an effect remain to be specified.

1.07.6 Conclusions

The material reviewed in this chapter has covered a wide range; this is true both for the empirical phenomena considered in the first part and the theoretical analyses dealt with in the second part. The latter point might seem to be a cause for concern, given our customary aspiration to achieve parsimony in explanation. But the concern would probably be misplaced. There is every reason to think that perceptual learning effects are the product (usually the joint product) of several different processes. On the basis of the evidence reviewed in this chapter, a place should be found for associatively mediated acquired equivalence and distinctiveness, effects based on within-stimulus association (unitization) and between-stimulus associations (associative inhibition), latent inhibition (and possibly other learned changes in associability), habituation, and salience modulation.

It may have been noticed that the operation of these various learning mechanisms in perceptual learning has been demonstrated for only a subset of the phenomena described in the first section of the chapter. Analytic studies have, for the most part, made use of just a few well-established and tractable experimental procedures. This may raise the fear that detailed exploration of other paradigms would uncover a whole new set of explanatory principles, in addition to those already listed. But this fear is not justified. As was suggested earlier, the job of a theory of perceptual learning is to explain how experience of similar stimuli can enhance the perceptual effectiveness of features that distinguish them and reduce the perceptual effectiveness of features that they have in common. This description is valid generally – it applies equally, for example, to rats learning to discriminate between flavors and to people learning to distinguish between speech sounds. We have every reason to hope that explanatory principles established in one of these paradigms will also apply in the other.

References

Ahissar M and Hohstein S (1997) Task difficulty and the specificity of perceptual learning. *Nature* 387: 401–406.

Amitay S, Hawkey DJC, and Moore DR (2005) Auditory frequency discrimination learning is affected by stimulus variability. *Percept. Psychophys.* 67: 691–698.

Aitken MRF, Bennett CH, McLaren IPL, and Mackintosh NJ (1996) Perceptual differentiation during categorization learning by pigeons. *J. Exp. Psychol. Anim. B.* 22: 43–50.

Artigas AA, Sansa J, Blair CAJ, Hall G, and Prados J (2006) Enhanced discrimination between flavor stimuli: Roles of salience modulation and inhibition. *J. Exp. Psychol. Anim. Behav. Process* 32: 173–177.

Ball K and Sekuler R (1982) A specific and enduring improvement in visual motion discrimination. *Science* 218: 697–698.

Battig WF (1956) Transfer from verbal pretraining to motor performance as a function of motor task complexity. *J. Exp. Psychol.* 51: 371–378.

Bennett CH and Mackintosh NJ (1999) Comparison and contrast as a mechanism of perceptual learning? *Q. J. Exp. Psychol.* 52B: 253–272.

Bennett CH, Tremain M, and Mackintosh NJ (1996) Facilitation and retardation of flavour aversion conditioning following prior exposure to the CS. *Q. J. Exp. Psychol.* 49B: 220–230.

Biederman I and Shiffrar MM (1987) Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual-learning task. *J. Exp. Psychol. Learn.* 13: 640–645.

Blair CAJ and Hall G (2003) Perceptual learning in flavor aversion: Evidence for learned changes in stimulus effectiveness. *J. Exp. Psychol. Anim. Behav. Process* 29: 39–48.

Blair CAJ, Wilkinson A, and Hall G (2004) Assessments of changes in the effective salience of stimulus elements as a result of stimulus preexposure. *J. Exp. Psychol. Anim. Behav. Process* 30: 317–324.

Bonardi C, Graham S, Hall G, and Mitchell C (2005) Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychon. Bull. Rev.* 12: 88–92.

Bonardi C, Rey V, Richmond M, and Hall G (1993) Acquired equivalence of cues in pigeon autoshaping: Effects of training with common consequences and common antecedents. *Anim. Learn. Behav.* 21: 369–376.

Bruce V and Burton M (2002) Learning new faces. In: Fahle M and Poggio T (eds.) *Perceptual Learning*, pp. 317–334. Cambridge, MA: MIT Press,.

Chiroro P and Valentine T (1995) An investigation of the contact hypothesis of the own-race bias in face recognition. *Q. J. Exp. Psychol.* 48A: 879–894.

Chung MS and Thomson DM (1995) Development of face recognition. *Br. J. Psychol.* 86: 55–87.

Demany L (1985) Perceptual learning in frequency discrimination. *J. Acoust. Soc. Am.* 78: 1118–1120.

Demany L and Semal C (2002) Learning to perceive pitch differences. *J. Acoust. Soc. Am.* 111: 1377–1388.

Dwyer DM, Bennett CH, and Mackintosh NJ (2001) Evidence for inhibitory associations between the unique elements of two compound flavours. *Q. J. Exp. Psychol.* 54B: 97–107.

Dwyer DM and Mackintosh NJ (2002) Alternating exposure to two compound flavors creates inhibitory associations between their unique features. *Anim. Learn. Behav.* 30: 201–207.

Dwyer DM, Hodder KI, and Honey RC (2004) Perceptual learning in humans: Roles of preexposure schedule, feedback, and discrimination assay. *Q. J. Exp. Psychol.* 57B: 245–259.

Fanselow MS (1990) Factors governing one-trial contextual conditioning. *Anim. Learn. Behav.* 18: 264–270.

Fahle M and Poggio T (2002) *Perceptual Learning.* Cambridge, MA: MIT Press.

Fiorentini A and Berardi N (1980) Perceptual learning specific for orientation and spatial frequency. *Nature* 287: 43–44.

Furmanski CS and Engel SA (2000) Perceptual learning in object recognition: Object specificity and size invariance. *Vision Res.* 40: 473–484.

Gagné RM and Baker KE (1950) Stimulus predifferentiation as a factor in transfer of training. *J. Exp. Psychol.* 40: 439–451.

Gauthier I and Tarr MJ (1997) Becoming a ''greeble'' expert: Exploring mechanisms for face recognition. *Vision Res.* 37: 1673–1682.

Gibson EJ (1969) *Perceptual Learning and Development.* New York: Appleton-Century-Crofts.

Gibson EJ and Levin H (1975) *The Psychology of Reading.* Cambridge, MA: MIT Press.

Gibson EJ and Walk RD (1956) The effect of prolonged exposure to visually presented patterns on learning to discriminate them. *J. Comp. Physiol. Psychol.* 49: 239–242.

Goldstone RL (1994) Influences of categorization on perceptual discrimination. *J. Exp. Psychol. Gen.* 123: 178–200.

Goldstone RL (1998) Perceptual learning. *Annu. Rev. Psychol.* 49: 585–610.

Goldstone RL (2000) Unitization during category learning. *J. Exp. Psychol. Gen.* 26: 86–112.

Goto H (1971) Auditory perception by normal Japanese adults of the sounds ''L'' and ''R''. *Neuropsychologia* 9: 317–323.

Hall G (1991) *Perceptual and Associative Learning.* Oxford, UK: Clarendon Press.

Hall G (1994) Pavlovian conditioning: Laws of association. In: Mackintosh NJ (ed.) *Handbook of Perception and Cognition, Vol. 9: Animal Learning and Cognition*, pp. 15–43. San Diego, CA: Academic Press.

Hall G (1996) Learning about associatively activated representations: Implications for acquired equivalence and perceptual learning. *Anim. Learn. Behav.* 24: 233–255.

Hall G (2003) Learned changes in the sensitivity of stimulus representations: Associative and nonassociative mechanisms. *Q. J. Exp. Psychol.* 56B: 43–55.

Hall G, Blair CAJ, and Artigas AA (2006) Associative activation of stimulus representations restores lost salience: Implications for perceptual learning. *J. Exp. Psychol. Anim. Behav. Process* 32: 145–155.

Hall G, Mitchell C, Graham S, and Lavis Y (2003) Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *J. Exp. Psychol. Gen.* 132: 266–276.

Hall G and Pearce JM (1979) Latent inhibition of a CS during CS-US pairings. *J. Exp. Psychol. Anim. Behav. Process* 5: 31–42.

Hollingworth HL (1928) *Psychology: Its Facts and Principles.* New York: Appleton-Century.

Holton RB and Goss AE (1956) Transfer to a discriminative motor task as a function of amount and type of preliminary verbalization. *J. Gen. Psychol.* 55: 117–126.

Iverson P, Kuhl PK, Akahane-Yamada R, et al. (2003) A perceptual interference account of acquisition difficulties for non-native phonemes. *Cognition* 87: B47–B57.

Iverson P, Hazan V, and Bannister K (2005) Phonetic training with acoustic cue manipulations: A comparison of methods for teaching English /r/-/l/ to Japanese adults. *J. Acoust. Soc. Am.* 118: 3267–3278.

James W (1890) *The Principles of Psychology.* New York: Holt.

Karni A and Sagi D (1991) Where practice makes perfect in texture discrimination: Evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Sci. USA* 88: 4966–4970.

Ketermann A and Siebert C (2003) A perceptual interference account of acquisition difficulties for non-native phonemes. *Cognition* 87: B47–B57.

Kiernan MJ and Westbrook RF (1993) Effects of exposure to a to-be-shocked environment upon the rat's freezing response: Evidence for facilitation, latent inhibition, and perceptual learning. *Q. J. Exp. Psychol.* 46B: 271–288.

Kruschke JK (2001) Toward a unified model of attention in associative learning. *J. Math. Psychol.* 45: 812–863.

Lavis Y and Mitchell C (2006) Effects of preexposure on stimulus discrimination: An investigation of the mechanisms responsible for human perceptual learning. *Q. J. Exper. Psychol.* 59: 2083–2101.

Lawrence DH (1949) Acquired distinctiveness of cues: I. Transfer between discriminations on the basis of familiarity with the stimulus. *J. Exp. Psychol.* 39: 770–784.

Le Pelley ME (2004) The role of associative history in associative learning: A selective review and a hybrid model. *Q. J. Exp. Psychol.* 57B: 193–243.

Liu Z and Weinshall D (2000) Mechanisms of generalization in perceptual learning. *Vision Res.* 40: 97–109.

Lubow RE (1989) *Latent Inhibition and Conditioned Attention Theory.* New York: Cambridge University Press.

Mackintosh NJ (1975) A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychol. Rev.* 82: 276–298.

Mackintosh NJ (1983) *Conditioning and Associative Learning.* Oxford, UK: Clarendon Press.

Mackintosh NJ and Little L (1969) Intradimensional and extradimensional shift learning by pigeons. *Psychon. Sci.* 14: 5–6.

Malpass RS and Kravitz J (1969) Recognition of faces of own or other race. *J. Pers. Soc. Psychol.* 13: 330–334.

McKee SP and Westheimer G (1978) Improvement in vernier acuity with practice. *Percept. Psychophys.* 24: 258–262.

McLaren IPL (1997) Categorization and perceptual learning: An analogue of the face inversion effect. *Q. J. Exp. Psychol.* 50A: 257–273.

McLaren IPL, Kaye H, and Mackintosh NJ (1989) An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition. In: Morris RGM (ed.) *Parallel Distributed Processing: Implications for Psychology and Neurobiology*, pp. 102–130. Oxford, UK: Clarendon Press.

McLaren IPL, Leevers HJ, and Mackintosh NJ (1994) Recognition, categorization, and perceptual learning (or, how learning to classify things together helps one to tell them apart). In: Umilta C and Moscovitch M (eds.) *Attention and Performance XV: Conscious and Nonconscious Information Processing*, pp. 889–909. Cambridge, MA: MIT Press,.

McLaren IPL and Mackintosh NJ (2000) An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Anim. Learn. Behav.* 28: 211–246.

Melcher JM and Schooler JW (1996) The misremembrance of wines past: Verbal and perceptual expertise differentially mediate verbal overshadowing of taste memory. *J. Mem. Lang.* 35: 231–245.

Miller NE and Dollard J (1941) *Social Learning and Imitation.* New Haven, CT: Yale University Press.

Mondragón E and Hall G (2002) Analysis of the perceptual learning effect in flavour aversion learning: Evidence for stimulus differentiation. *Q. J. Exp. Psychol.* 55B: 153–169.

Myles-Worsley M, Johnston WA, and Simons MA (1988) The influence of expertise on X-ray image processing. *J. Exp. Psychol. Learn.* 14: 553–557.

Pearce JM (1994) Similarity and discrimination: A selective review and a connectionist model. *Psychol. Rev.* 101: 587–607.

Pearce JM and Hall G (1980) A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* 87: 532–552.

Peron RM and Allen GL (1988) Attempts to train novices for beer flavor discrimination: A matter of taste. *J. Gen. Psychol.* 115: 403–418.

Prados J, Chamizo V, and Mackintosh NJ (1999) Latent inhibition and perceptual learning in a swimming pool task. *J. Exp. Psychol. Anim. B.* 25: 37–44.

Quinn PC, Palmer V, and Slater AM (1999) Identification of gender in domestic-cat faces with and without training: Perceptual learning of a natural categorization task. *Perception* 28: 749–763.

Rabin MD (1988) Experience facilitates olfactory quality discrimination. *Percept. Psychophys.* 44: 532–540.

Rescorla RA and Wagner AR (1972) A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black AH and Prokasy WF (eds.) *Classical Conditioning II: Current Research and Theory*, pp. 64–99. New York: Appleton-Century-Crofts.

Rodriguez G and Alonso G (2004) Perceptual learning in flavor aversion learning: Alternating and blocked exposure to a compound of flavors and to an element of that compound. *Learn. Motiv.* 35: 208–220.

Sathian K and Zangaladze A (1997) Tactile learning is task specific but transfers between fingers. *Percept. Psychophys.* 59: 119–128.

Shiu L-P and Pashler H (1992) Improvement in line orientation discrimination is retinally local but dependent on cognitive set. *Percept. Psychophys.* 52: 582–588.

Soloman GEA (1990) Psychology of novice and expert wine talk. *Am. J. Psychol.* 103: 495–517.

Stevenson RJ (2001) Perceptual learning with odors: Implications for psychological accounts of odor quality perception. *Psychon. Bull. Rev.* 8: 708–712.

Stevenson RJ and Boakes RA (2003) A mnemonic theory of odor perception. *Psychol. Rev.* 110: 340–364.

Stevenson RJ, Boakes RA, and Prescott J (1998) Changes in odor sweetness resulting from implicit learning of a simultaneous odor-sweetness association: An example of learned synesthesia. *Learn. Motiv.* 29: 113–132.

Sutherland NS and Mackintosh NJ (1971) *Mechanisms of Animal Discrimination Learning.* New York: Academic Press.

Symonds M and Hall G (1995) Perceptual learning in flavor aversion learning: Roles of stimulus comparison and latent inhibition of common elements. *Learn. Motiv.* 26: 203–219.

Todd IA and Mackintosh NJ (1990) Evidence for perceptual learning in pigeon's recognition memory for pictures. *Q. J. Exp. Psychol.* 42B: 385–400.

Trobalon JB, Sansa J, Chamizo V, and Mackintosh NJ (1991) Perceptual learning in maze discriminations. *Q. J. Exp. Psychol.* 43B: 389–402.

Volkmann AW (1858) Über den Einfluss der Uebung auf das Erkennen raumlicher Distanzen. *Ber. Verh. Sachs. Ges. Wiss. Leipzig Math. Phys. Kl.* 10: 38–69.

Woodworth RS and Schlosberg H (1954) *Experimental Psychology.* London: Methuen.

Wagner AR (1981) SOP: A model of automatic memory processing in animal behavior. In: Spear NE and Miller RR (eds.) *Information Processing in Animals: Memory Mechanisms*, pp. 5–47. Hillsdale, NJ: Erlbaum.

Yin RK (1969) Looking at upside-down faces. *J. Exp. Psychol. Gen.* 81: 141–145.