

Learned Changes in Stimulus Representations (A Personal History)

Geoffrey Hall
University of York, UK

Almost 40 years ago I began what turned out to be a programme of research on the way in which experience can change the effectiveness of the events used as stimuli in standard associative learning procedures. In this personal history I will describe my early (failed) attempts to find evidence for the acquired distinctiveness of cues, and my conclusion that experience tends to reduce, not enhance the associability of stimuli. I then go on to describe my attempts to square this conclusion with the stubborn empirical fact that, in some circumstances, pretraining with (or preexposure to) stimuli, can facilitate subsequent discrimination between them. I describe experiments (conducted mostly with rats as the subjects) showing how some of these effects can be explained in associative terms. Others, however, seemed to demand an explanation in terms of a new learning process that modulates the effective salience of stimuli. I go on to describe attempts to specify the nature of this process, and (bringing the story up to date) to describe recent experiments investigating the effects of salience modulation in human perceptual learning. *Keywords: associative learning, stimulus discrimination, pretraining, stimulus salience, perceptual learning*

Hace casi 40 años, empecé lo que con el tiempo se convirtió en un programa de investigación sobre la forma en que la experiencia puede cambiar la efectividad de los eventos empleados como estímulos en procedimientos típicos de aprendizaje asociativo. En esta historia personal, describiré mis primeros (fallidos) intentos de demostrar la distintividad adquirida de las claves, y mi conclusión de que la experiencia tiende a reducir, en vez de a facilitar, la asociabilidad de los estímulos. Después paso a describir mis intentos de hacer compatible esta conclusión con el innegable hecho empírico de que, en algunas circunstancias, el pre-entrenamiento con (o la pre-exposición a) los estímulos puede facilitar la posterior discriminación entre ellos. Describo los experimentos (llevados a cabo con ratas como sujetos) que muestran cómo algunos de estos efectos pueden explicarse en términos asociativos. Sin embargo, otros parecen exigir una explicación en términos de un nuevo proceso de aprendizaje que modula la saliencia efectiva de los estímulos. Paso a describir los intentos de especificar la naturaleza de este proceso y (llegando al momento actual) a describir los experimentos recientes que investigan los efectos de modulación de la saliencia en el aprendizaje perceptual humano. *Palabras clave: aprendizaje asociativo, discriminación estimular, pre-entrenamiento, saliencia estimular, aprendizaje perceptual*

I thank the SEPC (*Sociedad Española de Psicología Comparada -Spanish Society of Comparative Psychology-*) for inviting me to write this paper. Correspondence concerning this article should be addressed to Dr Geoffrey Hall, Department of Psychology, University of York, York YO10 5DD, UK. E-mail: g.hall@psych.york.ac.uk.

As I sat down to prepare this article the (somewhat disconcerting) thought occurred to me that it will appear in print almost exactly 40 years after I first started working on the issues it will discuss. Looking back on what had been done over those years provoked a mixture of feelings. What came to mind first was a consciousness of the mistakes that I had made – hypotheses, carefully nurtured, that turned out to be false; and worse, empirical phenomena, vigorously investigated, that turned out to be blind alleys, leading nowhere. But not all was gloom. I also thought that I could discern a central theme in which (with extensive help from others, both in my own laboratory, and elsewhere) I had been able to contribute to a genuine advance in our understanding. This theme (although this was not immediately apparent to me) concerns the comprehensiveness of associative theory. When I first learned about conditioning and its effects I thought that here was a process that had the potential to provide a firm basis for the analysis of what I took to be the central topic of psychology – the phenomenon of individual adaptation (the way in which individuals interacted with their environment, are changed by the experience and go on to interact again in a different way). I began to hope that the laws of association might provide for the science of psychology what Mendel's laws had given to the science of genetics.

But there was a difficulty. It was not so much that many instances of learning, as experienced in everyday life, were not immediately explicable in associative terms – the complexity of these phenomena might make analysis difficult, but this does not, in itself, prove that an associative analysis is impossible. More worrying was the fact that, even with the restricted area of animal learning itself, there appeared to be a range of phenomena that challenged the adequacy of associative theorizing. Foremost among these (to my mind) were observations that seemed to imply a process of perceptual learning. Associative theory concerned itself with the mechanisms by which experience forged links of various sorts between the central representations of events; it found no place for the possibility that the nature of these representations might be modified by experience. Not only would a theory that failed to deal with such things fall short of comprehensiveness; it would also run the risk of being rendered predictively impotent – if the nature of the central representation that is activated by a stimulus changes from one presentation to the next, we cannot be sure that the behaviour established on the first presentation will be evoked by the second. This matter seemed important enough to merit direct investigation. In what follows I will describe my own contributions to this investigation, past and present, ending with some speculations about what the future may hold.

The Past

Analyzer Theory

In the 1960s, Sutherland and his student Mackintosh developed a theory of the role of attention in animal

discrimination learning that became known as analyzer theory (see Sutherland & Mackintosh, 1971). The essential notion was that, in addition to forging associative links, discrimination training also engaged an attentional learning process that focussed attention on the relevant dimension of difference (strengthened the relevant analyzer). Central to this account was the phenomenon of the acquired distinctiveness of cues, first demonstrated for animal discrimination learning by Lawrence (1949). His experiments were taken show that when animals had been trained on one discrimination with a given set of cues, they were better able to learn a new, quite different task, involving these same cues. The change in the nature of the task meant that this transfer could not be response-based and must therefore, it was argued, depend on some change in the way that the stimuli themselves were processed.

I will return to acquired distinctiveness shortly, but when I began my doctoral research in 1968 I focussed on another phenomenon thought to support analyzer theory. This was the overtraining reversal effect, the paradoxical observation that rats trained on a simultaneous discrimination learned the reversal of this task more readily when they had been given overtraining on the original (e.g., Mackintosh, 1969). The interpretation offered by analyzer theory was that overtraining was so effective in strengthening the relevant analyzer that the beneficial effects of sustained attention to the relevant dimension were able to outweigh any potential disadvantages of such training. My own experiments were successful in replicating the basic effect, but revealed a number of features that made me doubt the attentional explanation (e.g., Hall, 1973). Perhaps the most critical was the finding that overtraining was capable of fostering later learning even when the stimuli were quite different from those used in initial training (Hall, 1974). The source of this positive transfer appeared to lie in the development of particular response-based strategies, and did not require the assumption that overtraining enhanced the degree of attention paid to the cues. This did not, of course, constitute a proof that attention is insensitive to modulation by experience; indeed the work I did next confirmed that it was (although not in the way envisaged by the Sutherland and Mackintosh theory).

Stimulus Associability and Latent Inhibition

In the mid 1970s, research on these matters turned to the use of classical conditioning procedures, prompted principally by Kamin's (1968) discovery of blocking—an effect that, at first sight, seemed to indicate a role for learned changes in attention. John Pearce and I decided to use the conditioned suppression procedure to explore the phenomenon of acquired distinctiveness. In the first of our attempts (Pearce & Hall, 1978) we trained rats on a discrimination between a reinforced stimulus (A+) and a nonreinforced compound of this stimulus and another (AB-). We reasoned that, at asymptote, the compound should have a net associative

strength of zero; transfer to further training in which the compound was used as a conditioned stimulus (CS) would therefore be capable of revealing the distinctiveness acquired by the cues, uncontaminated by direct associative transfer. Transfer was indeed obtained, but to our surprise it was the exact opposite of what we had expected. The AB compound was learned about more slowly as a result of this pretraining (see Figure 1). Far from acquiring distinctiveness (or the ability to command attention) during the first phase of training, the cues appeared to have lost it.

Further experiments, using modifications of this general design (e.g., Hall & Pearce, 1979, 1982) confirmed the generality of the finding and led us to adopt the general principle that training in which a cue was reliably followed by a given consequence would cause that cue to lose associability. The phenomenon of latent inhibition (the retardation of conditioning produced by prior nonreinforced exposure to the CS) was already well known. We simply extended the principle, making latent inhibition a special case, in which the stimulus consequence was no event at all. This principle formed the basis of our formal model of conditioning (Pearce & Hall, 1980), and we developed it (drawing liberally on the ideas of Mackintosh, 1975) to accommodate a range of other phenomena, including overshadowing and blocking.

	Train	Test
Group E	A+/AB-	AB+
Group C	X+/Y-	AB+

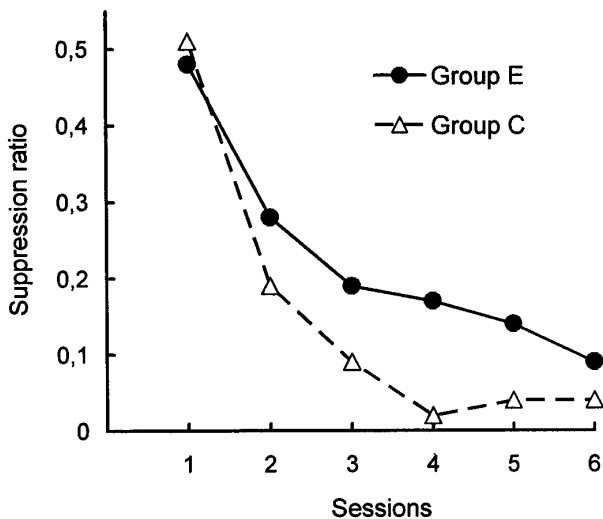


Figure 1. Design and results of an experiment by Pearce and Hall (1978), using rats and the conditioned suppression procedure. Rats in the experimental condition (Group E) were trained initially on a discrimination in which a compound was nonreinforced (AB-), but one of its elements was reinforced (A+). Rats in the control condition (Group C) were trained on a discrimination between two different cues. The graph shows the results of the subsequent test phase in which the AB compound was reinforced for all subjects.

Enhanced Discrimination

We were now confronted with a problem. The experiments (and our theory) said that exposure to a stimulus (at least when it signalled a consistent consequence) led to a loss of associability and poorer subsequent learning. What then, were we to make of the large body of evidence that seemed to suggest that prior experience could facilitate learning? This evidence came in two forms: One set of studies dealt with the effects of prior discrimination training on subsequent discrimination; the second set dealt with the effects of mere exposure to the cues.

Lawrence's (1949) experiment was the forerunner of a large number of studies, conducted during the 1950s and 1960s, mostly with human participants (see Hall, 1991, for a review), to demonstrate the reality of acquired distinctiveness effects – transfer from one discrimination task to another involving the same cues. My initial reaction to these studies was to try to find some sort of artifact in

	Discrimination	Conditioning	Test
Group 1	A → 0	C → shock	A and B
	B → food		
	C → 0		
Group 2	A → 0	C → shock	A and B
	B → food		
	C → food		

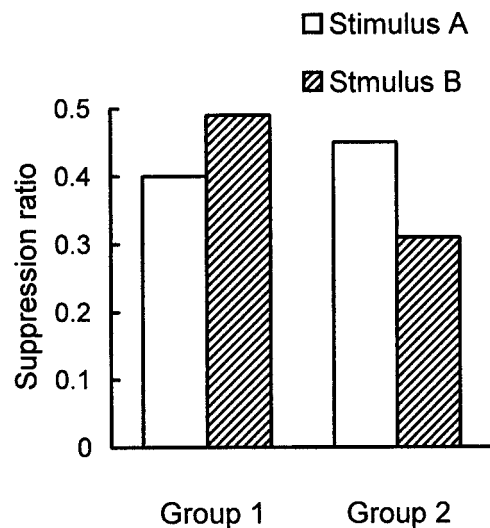


Figure 2. Design and results of an experiment by Honey and Hall (1989). Two groups of rats received initial discrimination training with three cues (A, B, and C). For Group 1, A and C had a common consequence; for Group 2, B and C had a common consequence. After C had been associated with shock (the Conditioning phase), generalization to A and B was tested using the conditioned suppression procedure. The lower part of the figure shows the result of this test.

each of them, and thus deny the reality of the effect. But I had to abandon this strategy when we succeeded in generating the effects in question, in our own laboratory and in the simplest of conditioning procedures. Figure 2 shows the design and result of one of these experiments (Honey & Hall, 1989). This experiment showed that rats will generalize more readily between two cues that have previously shared the same consequence than between two that have been pretrained with different consequences. The latter effect is the essence of acquired distinctiveness; the former has been termed acquired equivalence.

At the cost of disrupting the historical sequence, I will describe a recent experiment in which the same basic effect has been demonstrated for human subjects. The design of the experiment (Hall, Mitchell, Graham, & Lavis, 2003) 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). 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.

Turning now to the effects of mere exposure; one of the best demonstrations that facilitation effects can be obtained even when the preexposure phase does not involve explicit discrimination training comes from the classic study by Gibson

and Walk (1956). They gave infant rats prolonged exposure to geometrical shapes (presented in their home cages) and showed that this treatment improved subsequent performance on a food-rewarded simultaneous discrimination between these shapes. Initially I sought refuge in the possibility that there might something special about the developing nervous system of infant animals, and that adults given the same treatment might show a quite different effect (i.e., retarded learning, latent inhibition). But I quickly had to abandon this notion. Figure 3 shows the acquisition of a simultaneous discrimination by rats given preexposure to the cues or no

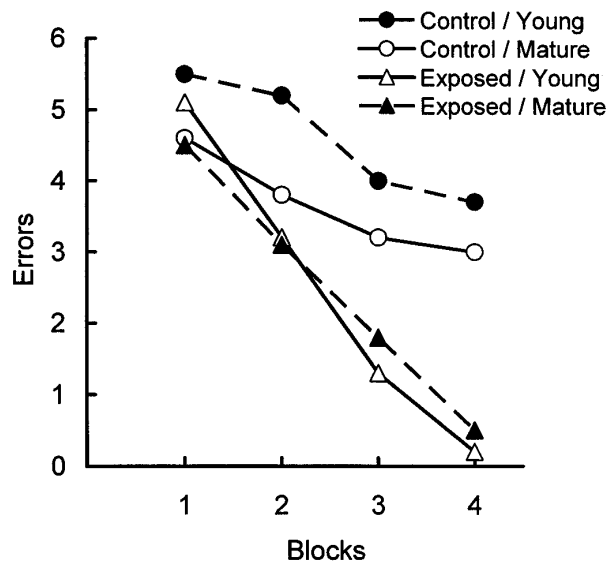


Figure 3. Results of an experiment by Hall (1979) showing the acquisition of a simultaneous discrimination, over four 5-day training blocks, by four groups of rats. Rats in the Exposed condition had been given prolonged exposure to the stimuli. For Young animals, this occurred between 40 and 80 days of age; for Mature animals, this occurred between 120 and 160 days of age.

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 = back slash; 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.

such preexposure. It is evident that preexposure helps; and that it does so whether the rats were preexposed as infants or when mature (Hall, 1979). Subsequent experiments confirmed the reliability of this effect and demonstrated that it was not necessary for the exposure to be given in the home cage (Channell & Hall, 1981). Here then was a procedure that clearly ought to produce latent inhibition, and yet learning was enhanced. How could this be?

The answer (in retrospect) seems obvious. Latent inhibition refers to retarded acquisition of the CR to a single cue; our examples of enhanced learning, on the other hand, all came from procedures in which the subject was required to discriminate between two similar cues. It is quite possible, then, that prior exposure to a pair of cues might reduce the associability of each of them, but that performance on a discrimination task could be enhanced if the preexposure also engaged some other learning process that enabled the animal to tell them apart better. A series of experiments using flavour-aversion learning procedures (e.g., Honey & Hall, 1989) found just this – acquisition of a conditioned aversion was retarded by prior exposure to the flavour used as the CS but discrimination between this and another similar flavour (as assessed by a generalization test) was found to be enhanced. The challenge, then, became that of identifying the nature of the learning process (or processes) that promotes the ability to distinguish between similar stimuli.


Associative Mediation in Acquired Distinctiveness (and Equivalence)

In fact an explanation for acquired distinctiveness was already available (and had been so since 1890 when it was first proposed by James). Figure 4 (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 4 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 be low, given the distinctiveness of their associates. The lower part of Figure 4 shows the situation for the procedure (known as acquired equivalence) in which the two stimuli are given the same associate. Here the proportion of common elements is increased, and discriminability should

go down. James focussed on the case in which the associates were distinctive names, but the analysis applies when other events are used (as in experiments with animal subjects, or that presented in Table 1).

Whether associations of this sort are actually responsible for acquired distinctiveness effects needs to be demonstrated. 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. As we shall shortly see, there are demonstrations of perceptual learning that seem to demand explanation in Gibsonian terms, but this

Acquired distinctiveness



Acquired equivalence




Figure 4. Associative structures in acquired distinctiveness and acquired equivalence. 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 overlapping circles represent two similar stimuli, A and B (the area of overlap representing features held in common); 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.

does not preclude the possibility that in some circumstances, associative processes can also 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 colours 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 colours had been involved in the discrimination test of Stage 2.

Effects of Mere Exposure

My story has now run ahead of itself, and I must go back to describe the work that we were doing in parallel with that on acquired equivalence/distinctiveness in order to investigate the enhanced discrimination produced by mere preexposure to stimuli. For, however striking the results supporting James' associative theory, it was, of course, obvious that this could not be the whole story – preexposure is capable of enhancing discrimination even when the stimuli are not followed by different consequences during the preexposure phase. The effect was readily demonstrated in experiments using the flavour-aversion learning procedure in which rats were conditioned with one compound flavour (AX) and generalization to another (BX) was tested. (Here A and B represent distinctive flavours, and X a third, added to each to render them more similar.) It was easy to show that preexposure, consisting of alternating trials with AX and BX, reduced generalization from AX to BX – that is, increased the ability of the rat to discriminate between them. This was true even when comparison was made with a control procedure in which the rats were given equivalent exposure to AX and BX, but with the stimuli presented in separate blocks of trial rather than intermixed. Figure 5 shows the design and results of an experiment (Symonds & Hall, 1995) demonstrating just such an effect.

This result turned out to be of particular significance. For one thing, it appeared to rule out the possibility that preexposure effects might be explained solely in terms of latent inhibition. The opportunity for latent inhibition to occur was the same in the two preexposure conditions—the groups differed only in the schedule of stimulus presentation, not in the amount of exposure—and yet the intermixed group showed better discrimination than the blocked group. Secondly, the result is just what would be expected on the basis of Gibson's (1969) understanding of the nature of perceptual learning. Gibson took perceptual learning to be a process that heightens the perceptual effectiveness of features that distinguish between similar stimuli, a process that is promoted by exposure in which the subject can compare the stimuli. Reduced generalization after intermixed preexposure matches this analysis perfectly. The intermixed schedule is just the sort of arrangement that might be expected to foster a comparison process; and one might expect that generalization between AX and BX would be reduced if the effective salience of the distinctive features, A and B, was enhanced by preexposure. In this training procedure, generalization will depend on the CR controlled by the X component of the stimuli, and the presence of a salient cue can be expected to interfere both with the acquisition and the expression of this CR. That is, the salient A will overshadow acquisition by X during conditioning; the salient B will interfere with the ability of X to evoke its CR in the generalization test.

	Preexposure	Conditioning	Test
Group C	–	AX+	BX
Group I	AX/BX/AX/BX...	AX+	BX
Group B	AX, AX...BX, BX	AX+	BX

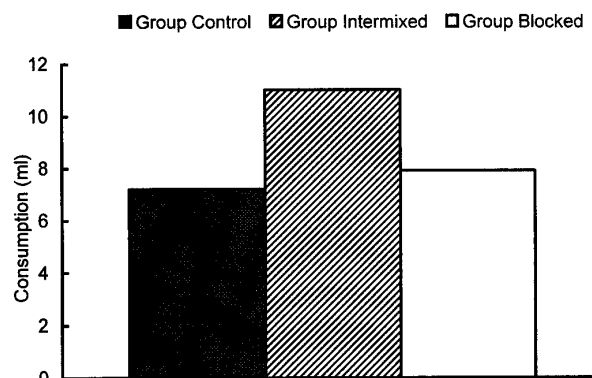


Figure 5. Design and results of an experiment by Symonds and Hall (1995). Rats were given access to compound flavours, AX and BX. Group I (intermixed) received preexposure in which the compounds were presented on alternate trials; Group B (blocked) experienced the compounds on separate blocks of trials. After aversion conditioning with AX (+ represents an injection of lithium chloride) generalization of the aversion to BX was tested.

But to say that this result is consistent with Gibson's (1969) account is not to show that that account is right. Indeed, at the time we published our results we were aware that the effect of schedule of exposure on generalization could be explained in terms of an associative theory proposed some years earlier, by McLaren, Kaye, & Mackintosh (1989). McLaren et al. pointed out that preexposure to AX and BX will allow the formation of associations between the various elements of these compound stimuli. In both the intermixed and blocked schedules, excitatory within-compound (A-X and B-X) associations can be expected to form. Additionally however, standard principles of associative learning imply that the intermixed procedure should allow the development of inhibitory associations between the unique features (A and B) of the preexposed stimuli. When it comes to testing BX after conditioning with AX, for subjects that have received blocked preexposure, responding will be partly determined by the ability of X to contact a representation of the unconditioned stimulus (US) by way of the associative chain X-A-US. This indirect source of responding will not be available to subjects that have been given intermixed preexposure, for whom the presence of B on the test will serve to inhibit activation of A. The result will be a less extensive generalization after intermixed than after blocked preexposure.

When I first came across this proposition it evoked a range of reactions; the first was to think how silly of me not to have come up with that myself; second was delight at this demonstration of the explanatory power of a very simple basic notion; third was the feeling that this explanation was too good to be true – and so it proved. Perhaps the most convincing evidence comes from experiments using a within-subject version of the experimental design depicted in Figure 5. This is shown as Experiment 1 in Figure 6. 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 generalise less readily to BX (the stimulus presented intermixed with AX during preexposure) than to the blocked stimulus, CX. This constitutes a within-subject example of the effect shown in Figure 5. This result in itself is compatible with the suggestion that the presence of B on test inhibits activation of the representation of the (conditioned) A element. But the inhibition explanation cannot apply to the modified design shown as Experiment 2 in the table. In this, 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. None the less, conditioning to AX was still found to generalize less well to BX than to CX on the test. The experiment presented as Experiment 3 in the figure makes the same point in a different way. Here conditioning was given with a novel stimulus, but the same basic result

(less generalization to a compound containing B than a compound containing C) was obtained. Since A is not conditioned with this procedure, any ability that B might have to inhibit activation of the A representation would be irrelevant to the outcome.

	Preexposure	Conditioning	Test
Experiment 1	AX/BX _ CX	AX+	BX and CX
Experiment 2	X/BX _ CX	X+	BX and CX
Experiment 3	AX/BX _ CX	Y+	BY and CY

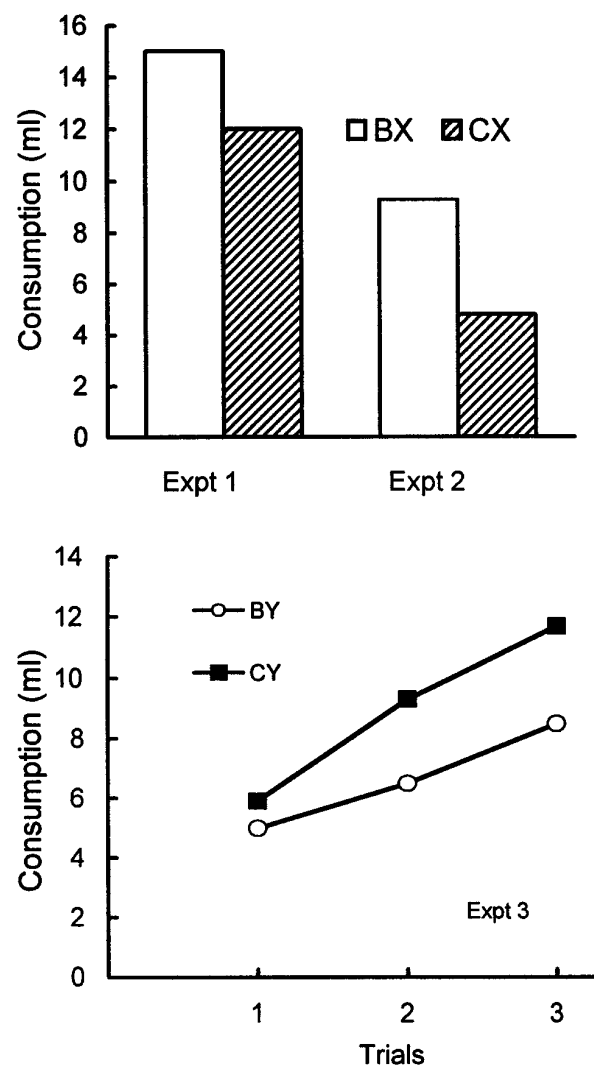


Figure 6. A, B, C, represent flavours that could be presented in compound with flavour X. In preexposure, AX was presented in alternation with BX (or X). The CX compound was presented on a separate block of trials. In the conditioning phase flavours were presented along with (+) and aversive reinforcer. The test phase measured the extent of the aversion shown to the test stimuli. Experiments 1 and 3 were reported by Blair and Hall (2003), and Experiment 2 by Hall et al. (2006).

These new results fit with Gibson's (1969) account as it was outlined above. The opportunity for comparison afforded by alternating presentations of AX and BX (or X and BX) will enhance the perceptual effectiveness of the distinguishing features of the stimuli (A and B in this case). Performance on the generalization test will be largely determined by the response that was established during the conditioning phase (to X in Experiments 1 and 2; to Y in Experiment 3). To the extent that the presence of another element (such as B or C) detracts from the ability of the animals to perceive stimulus the conditioned stimulus, 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 or Y to evoke its response on test and generalization will be restricted – the result obtained.

All this may seem a lot of work to establish what some seemed to know from the outset. Gibson (1969) was very clear that perceptual learning was not to be explained associatively and involved a separate learning process that heightened the perceptual effectiveness of the unique features that distinguish between stimuli. But the journey really was necessary. First it allowed us to establish that associative processes might well play an important role in perceptual learning effects (the acquired distinctiveness effect seems evidence of that). Second, it is not enough simply to assert that perceptual learning involves changes in perceptual effectiveness; this needs to be demonstrated experimentally. And finally, the challenge now becomes that of specifying the nature of the learning process responsible for these changes – we can hope that the experimental procedures described here will be useful in that regard.

The Present

I will now outline two lines of work that I am currently following that build on what has just been described. The second deals directly with the mechanisms responsible for changes in perceptual effectiveness; but before that I want to describe our recent attempts to extend the generality of our standard procedure for demonstrating a perceptual learning effect.

Same-Different Judgments

The examples of perceptual learning described above (Figure 6) all used the same basic test procedure – conditioning followed by a generalization test. My explanation of the outcome of these tests (that preexposure produces changes in the perceptual effectiveness of features of the stimuli) implies that the effect of preexposure should be evident with other test procedures. I have attempted to devise alternative tests in a series of experiments with rats

as the subjects (Blair, Wilkinson, & Hall, 2004); but an alternative strategy is to make use of humans whose discrimination can be tested by procedures that do not involve explicitly reinforced training or generalization testing – they can simply be instructed to report whether or not they can perceive a difference between AX and BX.

I have just completed such a study (Mitchell, Kadib, Nash, Lavis, & Hall, in press). The experimental subjects were Australian undergraduates and the stimuli the checkerboards shown in Figure 7. All had the same background (shown at the bottom of the figure); this constituted the X stimulus. Stimuli AX, BX, CX, and DX were prepared by adding a (different) unique feature, shown outlined in black in Figure 7. (The black outline is for illustrative purposes and was not present during experimental trials). With these stimuli we could use the within-subject version of the intermixed/blocked design. The students saw 60 presentations of each stimulus (each of 470 ms, with an interval between trials of 2 s). In the intermixed phase two of the stimuli (AX and BX) were presented on alternate trials; in the blocked phase, presentation of CX were given consecutively, followed by presentations of DX. (The order of these phases was counterbalanced across subjects, as was the particular version of the checkerboard designated as AX, BX, CX, or DX).

On the test trials the subjects saw stimuli presented in pairs (one for 800 ms, a 550-ms gap, and then the second for 800 ms). Their task was to simply to say “same” or “different”. There were four types of trial: Those on which the stimuli, AX and BX, that had been intermixed during preexposure, were presented (*intermixed different* trials); *intermixed same*, in which AX and AX (or BX and BX) were presented; *blocked different*, in which CX and DX were presented; and *blocked same*, in which CX and CX (or DX and DX) were presented. The results (accuracy of performance on these trials) are presented in Figure 8. It is evident that accuracy in responding “same”, when identical stimuli were presented, was high, making it impossible to detect any difference between the intermixed and blocked conditions, both for the intermixed and blocked stimuli. But a difference was evident on the “different” trials. The subjects continued to respond “same” to these (admittedly rather similar) blocked stimuli, but responded “different” to those given intermixed preexposure. Discrimination as assessed by the same-different test is better after intermixed than after blocked preexposure, an outcome readily explained by the proposal that perceptual effectiveness of unique features of the stimuli is higher after intermixed than after blocked preexposure.

Previous studies of perceptual learning in humans have used varied (and often exceedingly specialised) training procedures (see Hall, in press), and these have tended to give rise to ad hoc mini theories, devised solely to explain the effects obtained within a particular experimental paradigm. But what all these studies have in common (and share with the example just described) is that the procedure




Figure 7. The stimuli used in the experiment by Mitchell et al. (in press). The display at the bottom is the common background (X) used for all stimuli. In the other four a distinctive feature has been added (a pattern of 6 adjacent coloured squares). These features are outlined in black for the purpose of illustration.

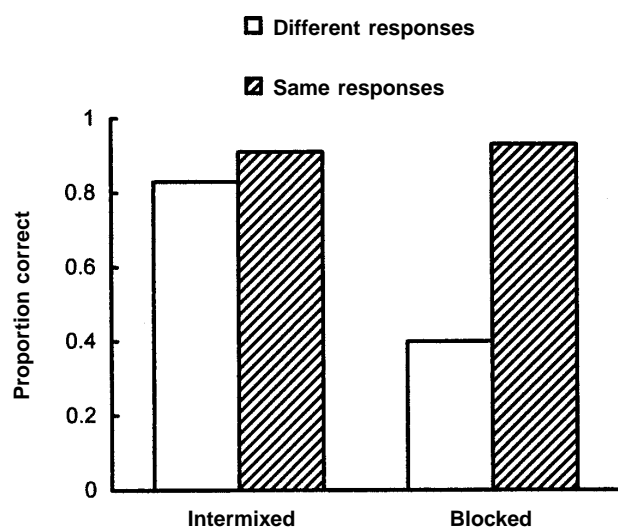


Figure 8. Proportion of correct responses on the same-different task used in the test phase of the experiment by Mitchell et al. (2007). All participants had received preexposure to one pair of stimuli presented on an intermixed schedule and to another pair presented according to a blocked schedule.

involves exposing people to similar stimuli that have a range of common features (such as X) and distinctive features (such as A and B). When discrimination is improved after such experience —when the radiographer becomes able to distinguish between the normal and abnormal X-ray, the chicken-sexer between cocks and hens, or the experimental participant between one checkerboard and another—this is because their behaviour has come to be controlled by the unique features of the stimuli rather than the common features. It is time now to consider the mechanisms involved.

Learned Changes in Salience

So far I have talked glibly of changes in the “perceptual effectiveness” of stimuli (or features of stimuli), offering no definition or explication of the term. This has been enough for our purposes up to this point, but when it come mechanisms, more precision is needed. Specifically, we want to be able to integrate some account of the learning process that produces changes in effectiveness with our existing accounts of learning. My own attempt to do this (Hall, 2003) took as its starting point the observation that formal theories of learning, such as the Pearce-Hall (1980) model had two parameters associated with the stimulus (or at least with the CS): The associability parameter (discussed previously) and a salience parameter (S), the value of which was set by the intensity of the stimulus. Although our theories have usually supposed that the salience of a given stimulus is fixed, I proposed that it might change with experience and that changes in perceptual effectiveness should be equated with changes in the value of S.

In fact the notion that salience can change is not all that novel (at least for stimuli categorised as USs). Repeated presentation of such a stimulus will produce habituation an outcome has been taken as reflecting a reduction in the sensitivity of the central representation of the stimulus, or, equivalently, a loss of stimulus salience. For USs this loss will show as reduction in the UR; for CSs the behavioural change may be less evident but we may assume that the central representation of the stimulus will undergo the same loss of sensitivity. It is a straightforward matter to propose, therefore, that the various preexposure procedures described above (such as those summarised in Figure 6) might produce a loss of salience. Less obvious is how such a loss might be reversed. To take a specific example, consider the design listed as Experiment 2 in Figure 6, which involved preexposure to alternating BX and X trials plus a separate block of CX trials. The test result showed that B was effectively more salient than C. Loss of salience by C as a result of CX presentations is no problem; but why does not B suffer the same loss during BX trials? Its preexposure treatment differed from that given to C only in that trials were alternated with presentations of X alone. How could presentation of X alone influence the sensitivity of the representation of B?

My answer to this latter question (Hall, 2003) was based on the observation that the only obvious way in which an X presentation might affect the B representation was by way of associative learning. Trials with BX can be expected to establish an excitatory association between X and B. On X-alone trials, therefore, the representation of B will be activated associatively. My further suggestion was that associative activation of a representation has a different effect from that produced by direct activation (i.e., by the presentation of the stimulus itself). The latter produces habituation; my suggestion was that associative activation reverses the habituation process, restoring lost salience. Stimulus C will simply lose salience during preexposure; for stimulus B, by contrast, the presence of the alternated X trials will oppose this effect, leaving B more salient than C

I was pleased with this account – it had a nice symmetry and appeared to fit the facts – but I was well aware that it lacked empirical support. Recently I have been trying hard to supply some. The design of one of our latest experiments (Hall, Blair, & Artigas, 2006, Experiment 3a) is shown at the top of Figure 9. There were two groups of rats, both given essentially the same treatment as that depicted for Experiment 2 in Figure 6; that is they received trials in which presentations of the BX compound alternated with X and a separate block of trials with the CX compound. The critical difference was that in this experiment *serial* compounds were used. The group labelled forward received access to X followed immediately by access to B (and to X followed by C on the blocked trials); the backward group received B (or C) *before* X. The hypothesis under investigation holds that the salience of B is restored on X-alone trials when X is able to activate the representation of B associatively. This will be

true for the forward group of the present experiment; that is the forward pairings in X precedes B should establish X as an excitatory CS for B. But it will not be the case for the backward group; trials in which B precedes X cannot be expected to establish the necessary association. The hypothesis predicts therefore all the stimuli should lose salience, with the exception of stimulus B for the forward, in which case effective salience will be maintained.

The results of the experiment are shown at the bottom of Figure 9. After flavour-aversion conditioning with X as the CS, all the rats received two test trials with each of the compounds, BX and CX. As conditioning had established an aversion to X, levels of consumption on test were low (consumption was less suppressed on the second test trial, presumably reflecting extinction of the aversion). Critically, suppression of consumption was found to be attenuated in just one of the test conditions—that for rats in the forward group given the BX compound. Only in this condition, we concluded, had the preexposure procedure left the B stimulus with enough salience that it could interfere with expression of the aversion conditioned to X—just what the theory had predicted

	Preexposure	Conditioning	Test
Forward group	X/X-B _ X-C	X+	BX and CX
Backward group	X/B-X _ C-X	X+	BX and CX

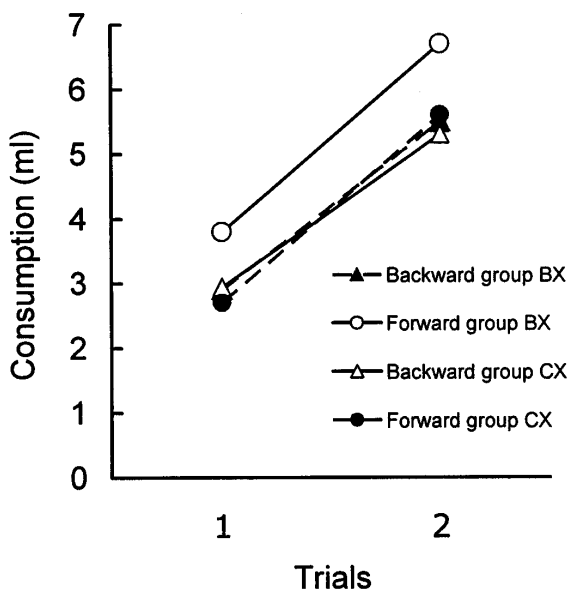


Figure 9. Design and results of an experiment by Hall et al. (2006). B, C, and X represent flavours. In preexposure X was presented in alternation with the serial compound X followed by B (Forward group) or B followed by X (Backward group). On a separate block of trials, X and C were presented either in the forward or the backward arrangement. Simultaneous compounds were tested after aversion conditioning with X as the CS (X+).

The Future

It goes without saying that much more work along the lines of the experiment just described will be needed to evaluate properly the particular theory of salience change being proposed here. But I want to move away from such details and look at a somewhat bigger picture in considering, in this final section of the paper, what research might be needed in the future. I have three specific issues that I want to mention, but each may be seen as a particular aspect of a more general point, that I will come to finally.

First, we have now reached a theoretical position that endows the CS with two parameters (associability and salience) each of which can change with experience. To some this may seem over-generous; but the notion has clear psychological plausibility – I may well want to limit some aspect of the attention I pay to an event whose consequences are well known (I no longer need it to be high in associability) while at the same time maintaining the attention needed to respond appropriately when it appears (i.e., maintaining its effective salience). It remains true, however, that with two parameters free to vary, the theory runs the risk of being dangerously flexible (and invulnerable to proof and disproof). What is needed for the future, therefore, is a precise specification of the (different) learning rules that govern changes in associability and in salience. I suspect that these will need to be developed as a formal theory, the worth of which will be established by its ability to make quite new predictions

The next issue concerns the need for the theory just outlined to confront the persistent trickle of experimental evidence taken to show that, in some circumstances, the attention paid to stimuli can be enhanced by training in which the cue is followed by a consistent consequence (for a recent example see Le Pelley & McLaren, 2003). Some of the experiments taken to show the acquired distinctiveness of cues can be explained away in other terms (something I attempted at length in my 1991 book); but there are other results (even some supplied by me; Bonardi, Graham, Hall, & Mitchell, 2006) that seem to demand an attentional explanation. Do we really need a third form of attention, in addition to the two just described? Clearly we need at least to consider the possibility when it comes to devising any new formal theory.

Third, almost all the discussion so far has been concerned with the events that we usually categorise as CSs. But what about USs? The possibility that the central representations of these stimuli might be changed by experience has already been acknowledged in our discussion of habituation. With repeated presentation, the properties of an electric shock appear to change – it hurts less; its effective salience declines. If a US can suffer changes in salience, why not also changes in other attentional parameters such as associability. And if such changes can occur what might be the learning rules that govern them – will they be the same as for CSs or different?

The general point that underlies the three just made is my feeling that we badly need an integrated theory that draws

together these various issues. They are not independent (they are often very similar) but we have tended to develop separate accounts for them. To give just one example, the US-preexposure effect is explained in one way (e.g., as context blocking or habituation) whereas latent inhibition (the CS-preexposure effect) gets a different set of explanations. Of course, several processes may be operating concurrently, which makes it all the more important to have a formal theory that deals with them all. It also means that developing such a theory will take a lot of work – but let's hope not another 40 years.

References

- Blair, C.A.J., & Hall, G. (2003). Perceptual learning in flavor aversion: Evidence for learned changes in stimulus effectiveness. *Journal of Experimental Psychology: Animal Behavior Processes*, 29, 39-48.
- Blair, C.A.J., Wilkinson, A., & Hall, G. (2004). Assessments of changes in the effective salience of stimulus elements as a result of stimulus preexposure. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 317-324.
- Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychonomic Bulletin & Review*, 12, 88-92.
- Channell, S., & Hall, G. (1981). Facilitation and retardation of discrimination learning after exposure to the stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, 7, 437-446.
- Gibson, E.J. (1969). *Perceptual learning and development*. New York: Appleton-Century-Crofts.
- Gibson, E.J., & Walk, R.D. (1956). The effect of prolonged exposure to visually presented patterns on learning to discriminate them. *Journal of Comparative and Physiological Psychology*, 49, 239-242.
- Hall, G. (1973). Overtraining and reversal learning in the rats: Effects of stimulus salience and response strategies. *Journal of Comparative and Physiological Psychology*, 84, 169-175.
- Hall, G. (1974). Transfer effects produced by overtraining in the rat. *Journal of Comparative and Physiological Psychology*, 87, 938-944.
- Hall, G. (1979). Exposure learning in young and adult laboratory rats. *Animal Behaviour*, 27, 586-591.
- Hall, G. (2003). Learned changes in the sensitivity of stimulus representations: Associative and nonassociative mechanisms. *Quarterly Journal of Experimental Psychology*, 56B, 43-55.
- Hall, G. (in press). Perceptual learning. In J. Byrne (Ed.), *Learning and memory: A comprehensive reference*. Elsevier.
- Hall, G., Blair, C.A.J., & Artigas, A.A. (2006). Associative activation of stimulus representations restores lost salience: Implications for perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 32, 145-155.
- Hall, G., Mitchell, C., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for associative mediation. *Journal of Experimental Psychology: General*, 132, 266-276.
- Hall, G., & Pearce, J.M. (1979). Latent inhibition of a CS during CS-US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, 5, 31-42.
- Hall, G., & Pearce, J.M. (1982). Changes in stimulus associability during conditioning: Implications for theories of acquisition. In M.L. Commons, R.J. Herrnstein, & A.R. Wagner (Eds.), *Quantitative analyses of behavior, Vol. III: Acquisition* (pp. 221-239). Cambridge, MA: Ballinger.
- Honey, R.C., & Hall, G. (1989). Enhanced discriminability and reduced associability following flavor preexposure. *Learning and Motivation*, 20, 262-277.
- Honey, R.C., & Hall, G. (1989). Acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 338-346.
- James, W. (1890). *The principles of psychology*. New York: Holt.
- Kamin, L.J. (1968). "Attention-like" processes in classical conditioning. In M.R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9-33). Coral Gable, FL: University of Miami Press
- Lawrence, D.H. (1949). Acquired distinctiveness of cues: I. Transfer between discriminations on the basis of familiarity with the stimulus. *Journal of Experimental Psychology*, 39, 770-784.
- Le Pelley, M.E., & McLaren, I.P.L. (2003). Learned associability and associative change in human causal learning. *Quarterly Journal of Experimental Psychology*, 56B, 68-79.
- Mackintosh, N.J. (1969). Further analysis of the overtraining reversal effect. *Journal of Comparative and Physiological Psychology, Monograph Supplement*, 67, 1-18.
- Mackintosh, N.J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- McLaren, I.P.L., Kaye, H., & Mackintosh, N.J. (1989). An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition. In: R.G.M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 102-130). Oxford: Clarendon Press.
- Mitchell, C., Kadib, R., Nash, S., Lavis, Y., & Hall, G. (in press). Analysis of the role of associative inhibition in perceptual learning by means of the same-different task. *Journal of Experimental Psychology: Animal Behavior Processes*.
- Pearce, J.M., & Hall, G. (1979). Loss of associability by a compound stimulus comprising excitatory and inhibitory elements. *Journal of Experimental Psychology: Animal Behavior Processes*, 5, 19-30.
- Pearce, J.M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532-552.
- Sutherland, N.S., & Mackintosh, N.J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- Symonds, M., & Hall, G. (1995). Perceptual learning in flavor aversion learning: Roles of stimulus comparison and latent inhibition of common elements. *Learning and Motivation*, 26, 203-219.

Received August, 10, 2007
Accepted September 11, 2007