UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Spatial Variability and Peak Shift: A Challenge for Elemental Associative Learning?

Permalink
https://escholarship.org/uc/item/3402b04n

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN
1069-7977

Author
Lo, Lap-Yan

Publication Date
2005

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Spatial Variability and Peak Shift: A Challenge for Elemental Associative Learning?

E. J. Livesey (el253@cam.ac.uk)
L. S. Pearson (lsp25@cam.ac.uk)
I. P. L. McLaren (iplm2@cam.ac.uk)
Department of Experimental Psychology, University of Cambridge
Downing Street, Cambridge CB2 3EB. UK.

Abstract

Recent evidence from animal research suggests that peak shift effects produced with complex pattern stimuli are influenced by the level of spatial variability over successive presentations of the stimuli. Initial attempts to explain this have relied on postulating a shift between elemental and configural representations supporting learning. Here we outline an elemental model based on contemporary associative theory that can accurately simulate this effect, and furthermore can do so without any built-in assumptions of dimensionality and without requiring some strategic shift from elemental to configural processing. Further evidence in support for the model is provided from a categorization experiment examining the effect of spatial variability with human subjects.

Peak shift with an artificial dimension

Peak shift is a well-documented consequence of stimulus discrimination and generalization and continues to stimulate interesting research nearly half a century after its discovery (for recent reviews see Ghirlanda & Enquist, 2003; Honig & Urcioli, 1981). Typically, subjects will be trained to discriminate between two very similar stimuli, usually by learning to respond to one (S+) and not the other (S-). Subjects are then presented with a range of stimuli that lie at successive points along the dimension to test for stimulus generalization. Peak shift takes place when the peak response rate (or accuracy of response) occurs not for S+ but for a similar stimulus that is further away from S-.

By its very nature, peak shift can only be observed over a series of stimuli that have a systematic relationship with one another. While very simple stimuli with characteristics that lie along a physical dimension such as wavelength of light obviously fit this requirement, some experiments have used much more complex stimuli including morphed faces (Spetch, Cheng, and Clifford, 2004) and patterns of small abstract shapes commonly referred to as 'icons' (Wills and Mackintosh, 1998; Oakeshott, 2002). The latter generally employ a range of icons organized so that a logical sequence of patterns is produced, akin to a series of points along a dimension. The frequency of occurrence of each type of icon is systematically varied from one stimulus to the next so that shifting one 'step' along the artificial dimension means replacing some icons with new ones but still retaining a proportion from the original. Figure 1 demonstrates one hypothetical set of four stimuli of this nature (see Table 1 below for further explanation of the dimensional design).

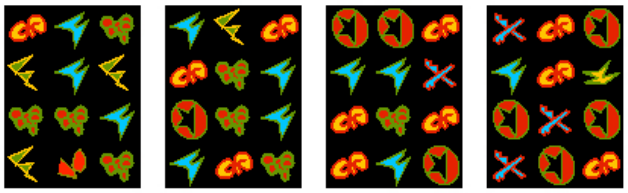

Figure 1. Examples of 'icon' stimuli. A possible sequence of generalization test stimuli used here.

Peak shift along an artificial dimension has been of particular interest to proponents of an elemental associative learning account of generalization and discrimination. Associative theories such as that proposed by Blough (1975) have been used to model peak shift effects with great success using several assumptions about the pattern of graded activation across hypothetical sensory units which are stimulated by the presentation of training and test stimuli. While the details of Blough's (1975) model and others like it are beyond the scope of this paper, it is worth noting that one of the key reasons for studying peak shift along an artificial dimension was to gain more control over the similarity between stimuli and the pattern of activation across units that they might stimulate (Wills and Mackintosh, 1998). Indeed, peak shift effects with icon stimuli have now been shown with both human and pigeon subjects under a variety of experimental conditions (Jones and McLaren, 1999; Livesey, 2004; Oakeshott, 2002; Wills and Mackintosh, 1998), and several elemental associative models are able to predict peak shift and related generalization effects with impressive accuracy. Generally speaking, these models rely on assumptions about the underlying dimensionality of the stimuli in order to calculate how the patterns of activation of two similar stimuli might overlap. This seems plausible when modelling generalization between stimuli that have measurable, physical similarities, even though the units are

hypothetical and to a large degree untestable. However, relying on dimensional assumptions when the dimension itself is artificial may be less tenable. Alternatively, for the purpose of simulating icon experiments, some models have assumed that different units are solely activated by a single type of icon but that the activation of each unit is directly proportional to the number of copies of each icon. This too is somewhat problematic. There is no way to gauge whether two copies of an icon would indeed generate twice the activation of a single copy, for instance. In any case, a recent finding by Oakeshott (2002) has highlighted the inadequacies of both of these approaches.

The effect of spatial variability

Recent research suggests that the level of spatial variability across successive presentations of these complex patterns of icons significantly affects the post-discrimination generalization gradient. This poses an interesting problem for elemental associative models which assume that the relative position of the icons should make little or no difference to the overall pattern of generalization. Icon experiments that have produced strong peak shift effects have typically allowed the positions of each icon within the pattern to vary randomly from one trial to the next, and many experiments have also allowed the exact number of each type of icon to vary stochastically, averaging to a given frequency over multiple trials. Looking at peak shift with pigeon subjects, Oakeshott (2002) demonstrated that when spatial and frequency variation are removed so that the position and number of each icon for a given stimulus remains fixed throughout the experiment, the peak shift effect seems to disappear. Using the design shown in Table 1, Oakeshott trained pigeons to discriminate between S+ and S- (highlighted), under either fixed or variable spatial organization. When tested across several stimuli, pigeons in the variable group showed evidence of a peak shift, with highest response rates for N+, one step removed from S+, and declining response rates further along the dimension at F+. However, pigeons in the fixed group did not show any evidence of peak shift, with highest response rates at S+, declining steadily to N+ and F+. Oakeshott (2002) found that the varied generalization gradient mirrors the predictions of an elemental model such as Blough (1975) while the fixed gradient resembles that predicted by a configural model such as Pearce (1987; 1994). As Oakeshott (2002) noted, the result taken as a whole appears to be irreconcilable with either type of model. Hence it poses a significant problem for the associative models that seek to explain post-discrimination generalization gradients in experiments with artificial dimensions. However, as discussed below, a more sophisticated elemental model can in fact account for both of these results.

Table 1. Dimension used by Oakeshott (2002)

Stimulus	Frequency of occurrence of icons A to I in each stimulus								
	A	B	C	D	E	F	G	H	I
F+	1	3	4	3	1				
		1	3	4	3	1			
N+			1	3	4	3	1		
S+				1	3	4	3	1	
S-					1	3	4	3	1

Modeling the fixed vs varied result

While Oakeshott (2002) concluded that these gradients appeared to support conflicting models of learning and generalization, there remains a strong case for developing a model that can accommodate both patterns of responding in appropriate circumstances. Presented below is just such a model derived from the elemental associative theory proposed by McLaren and Mackintosh (2000; 2002), which has a wide scope of application to animal conditioning, human learning, perceptual learning, and discrimination and generalization effects such as peak shift. The model predicts the differences in post-discrimination generalization gradient between the fixed and varied conditions in both the current experiment and by Oakeshott (2002). Several aspects of this model diverge from past attempts to simulate peak shift effects. Unlike previous models of peak shift, there is no attempt to assume or impose any form of dimensionality in the form of graded activations or differential generalization based on the position of the stimulus within the artificial dimension. For simplicity the stimulus, represented by activation of input units or nodes, takes on a relatively simple, local coding whereby each icon in each specific location activates a single input node.

The mechanism through which learning proceeds is the delta rule (for instance see Widrow-Hoff, 1960) linking a large number of representational units to one or more output units. The complexity that allows the model to predict different generalization gradients for the fixed and varied cases lies in the representational coding of the stimuli themselves. As the spatial location of the icons is important for the presence of peak shift, each icon in each different location is given a separate input unit which is fully activated ($A_i = 1$) when that location-specific icon is present and not activated at all ($A_i = 0$) otherwise.

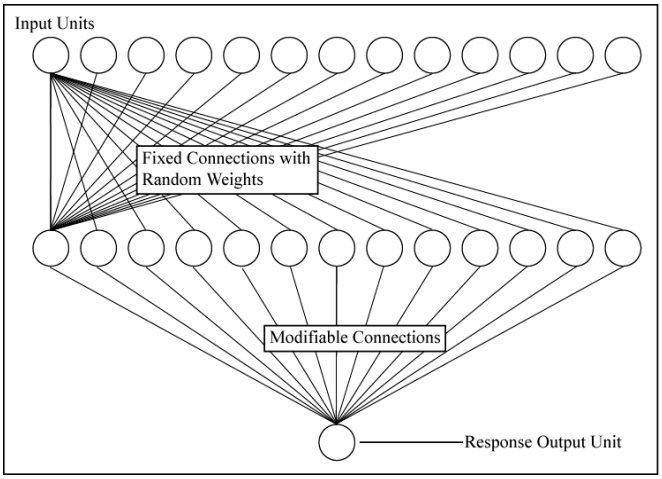

Figure 2. A schematic representation of a model based on McLaren and Mackintosh (2000; 2002)

Thus to model a single subject, one needs 144 input units (12 different icons are used and the stimulus has 12 different locations in which an icon can appear).

These units are then connected to a second layer of units with connection strengths that vary randomly between -1 and +1 and are fixed throughout the experiment (this system of representation is assumed to exist prior to the experiment and does not change as a result of exposure to the stimuli or the outcomes involved). This second layer can be thought of as the "elements" described in the McLaren and Mackintosh model. Importantly, there is no assumption of dimensionality inherent in this system of representation – the elements are activated via a series of random weights with all of the activated inputs such that input $i = \sum A_i W_{ij}$ where A_i is the activation of input unit i and W_{ij} is the random fixed weight between input unit i and second layer unit j. If this input is negative then unit j has no activation ($A_j = 0$), otherwise if i is positive then activation of unit j, $A_j = E.i / (D + E.i)$ where E and D are constants. In effect, this means that on presentation of any specific stimulus, roughly half of the second layer units will have an activation of zero as their summed input will be negative, while the activation of the other half will vary between 0 and 1.

For the purposes of modeling Oakeshott's (2002) animal-based experiment, each of these second layer units is connected to just a single output unit corresponding to the key-peck response required for reinforcement. These weights are initially set to zero and vary according to the delta-rule. The change in weight between second layer unit i and output j, $W_{ij} = S.\Delta.A_i$ where S is a learning rate parameter, A_i is the activation of unit i and Δ is the difference between external input, e, and summed internal input, i. The summed input i is calculated in much the same fashion as for the second layer units ($i = \sum A_i W_{ij}$) . During training, the output unit was given an external activation of 1 for reinforced S+ trials and an external activation of 0 for S- trials.

Since Oakeshott (2002) reported proportional response rates to each of four test stimuli (F+, N+, S+, S-), for this model the summed input to the single response output was also converted to proportional values of e^{kI} for the same four stimuli, where k is a constant held at a single value for all subjects and I is the summed input to the response output unit. The results of these simulations can be seen below in Figure 3, plotted against the proportional response rates reported by Oakeshott (2002). It is quite clear that the model predicts a peak shift for the varied condition but not for the fixed condition.

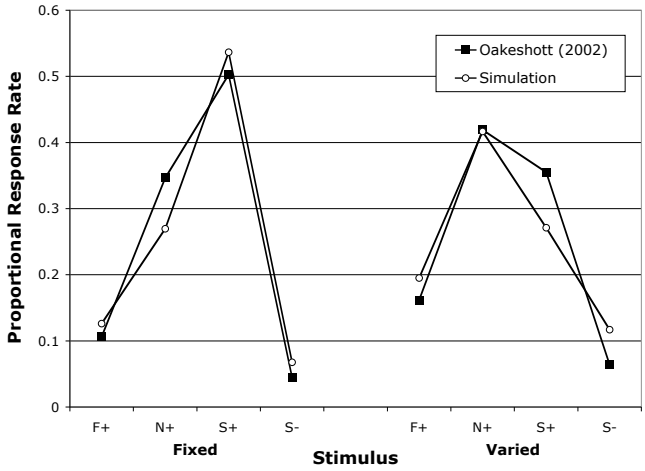

Figure 3. Experimental data from Oakeshott (2002) and results from simulation.

The rules governing weight change and activation from summed inputs are deliberately relatively simple and chosen as a matter of convenience rather than being favoured for any theoretical reason. The key aspect of the model that allows it to capture the difference between the two conditions is the interaction between the pattern of activation of the second layer units and weight change based on a summed error term. The system of representation, with random fixed weights between the inputs and hidden units allow the expression of both general frequency of occurrence of icons and also location specific configurations of icons. In the fixed condition, the units that are activated by configurations of position-specific icons present in S+ but not S- gradually accrue the strongest weights to the response output unit. As soon as some icons change, as happens when presenting N+, the configurations change and the activation of the units governing the discrimination may change considerably. In contrast, in the varied condition, the hidden units that will gradually come to govern the discrimination are those that are *generally* activated by icons occurring more frequently in S+ than S-, irrespective of the position of these icons. This involves a more gradual and variable process of weight change, which seems to fit with the observation that varied conditions are typically more difficult and are learnt more slowly than fixed conditions (Livesey, 2004;

Oakeshott, 2002). In this case, N+ will contain many of the icons predictive of S+ (and relatively few of those predicting S-) and will result in a high level of activation of the units with associative control, leading to peak shift.

The fixed vs varied effect in human categorization

Peak shift with icon stimuli has been found with human subjects in a categorization paradigm (Wills and Mackintosh, 1998) and the current model can quite readily be applied to human categorization as well as animal conditioning. The results of the model and supporting data from a human categorization experiment are discussed below.

Using the same architecture but with two response output units rather than one, we simulated a very similar experiment in which the task was to learn to respond with a left key press for one training stimulus (S_L) and with a right key press for another training stimulus (S_R). The modifications to the model were kept to a minimum. The system of representation expressed through the input units and random fixed weights to the second-layer units was identical to the animal model. Each of the second layer units was connected to both output units, with weights modified by the same mechanism as before. During training, on each trial the output unit corresponding to the correct response was given an external activation of 1 while the incorrect response output had an external activation of 0. During the test phase, response probabilities for each test stimulus were calculated using a log form of the ratio rule. For instance, the probability of a left response, P(Left Response) = $e^{kL} / (e^{kL} + e^{kR})$ where L and R are the summed inputs to the units corresponding to left response and right responses respectively, k is a constant held at a single value for all subjects.

Not surprisingly, this model predicted very similar results to the animal model, with a peak shift evident for the varied condition but not for the fixed. An experiment with human subjects was run as further confirmation of the predictions of the model and to replicate Oakeshott's (2002) effect of spatial variability. Participants in the fixed and varied conditions were first given discrimination training in which they were required to respond differentially to intermixed presentations of two training stimuli, S_L (requiring a left key response) and S_R (requiring a right key response). A transfer test then followed in which both training stimuli and a further three stimuli from each side of the dimension were presented in successive trials. In terms of the position of each of the stimuli along the artificial dimension, the transfer test was symmetrical with respect to the training stimuli. Therefore, the two 'sides' of the dimension were collapsed for the purposes of analyses, with results to be expressed in terms of the distance to the nearest training stimulus (combined in this way, S_L and S_R will be referred to as S). Therefore positions 1 and 8 are 3 steps from S and are equivalent to F+ in the Oakeshott (2002) experiment, while positions 3 and 6 are 1 step from S, and are equivalent to S+. It is worth noting that under this symmetrical design any observed peak shift cannot reasonably be attributed to simple range effects as had been argued for some other peak shift experiments using human subjects (see for instance Thomas, 1993). We predicted that the varied group would show a peak shift similar to that found in past experiments while the fixed group would most likely exhibit the same generalization gradient as that found by Oakeshott (2002).

Table 2. Dimension used in the current experiment, based on Oakeshott (2002).

Stimulus no.	distance from nearest S	Frequency of occurrence of icons A to L in each stimulus											
		A	B	C	D	E	F	G	H	I	J	K	L
3	1	1	3	4	3	1							
2	2		1	3	4	3	1						
1	3			1	3	4	3	1					
0	4				1	3	4	3	1				S_L
0	5					1	3	4	3	1			S_R
1	6						1	3	4	3	1		
2	7							1	3	4	3	1	
3	8								1	3	4	3	1

One concern with using this method with human subjects was that it is relatively easy for participants to form rules about a particular icon and how often it appears in each training stimulus. Similar experiments in the past have demonstrated that the formation of rules can have a large impact on post-discrimination generalization gradients (Jones & McLaren, 1999). For instance, icon F always appears three times in S_R and four times in S_L, and one could formulate a simple rule based solely on icon F which would still solve the discrimination. A subject responding according to this rule might well respond with a right key press to stimuli corresponding to position 3 of the dimension since these stimuli also have three copies of icon F. Icon counting rules are applicable to both the fixed and varied conditions. As such, they could significantly reduce any peak shift in the varied condition and also render the lack of peak shift in the fixed condition fairly meaningless since the generalization gradient may be due more to the application of a rule than to any other process of generalization, associative or otherwise. Rules also tend to break down quickly during the generalization test, given the variation in the number of each icon at each position along the dimension, and hence their use produces a tendency towards chance level responding as soon as the test stimulus appears even slightly dissimilar to the training exemplars. Experiments using stimuli of this size with

human subjects have produced a peak shift in the past (e.g. Wills and Mackintosh, 1998), but have generally used a procedure whereby the exact number of each icon varies from trial to trial during training, and have thus overcome rule-based icon counting in a way that cannot be employed in this study. The only human study to date that has looked at the effect of spatial variability on generalization gradients with icon stimuli used much larger and more complicated patterns made up of 36 icons, which were less susceptible to the influence of such rules (Livesey, 2004). To combat this potential problem, participants' data were only used in the analysis if their mean accuracy across all test stimuli was significantly above chance level. Pilot testing also suggested that the fixed condition was significantly easier to solve than the varied condition. To equate the mean test accuracy of the two groups and to prevent ceiling effects in the fixed condition, the duration of training in the fixed condition was reduced relative to the varied group. Matching the two groups for accuracy was necessary so that the accuracy criteria did not exclude a greater range of subjects in one group than the other. For instance, if the varied group training was of the same duration as the fixed then all but the fastest learners would fail to show performance above chance, which may well serve to bias results from the test phase. There is a possibility that the difference in training duration may affect more than just overall accuracy, however it is not clear that the changes in duration of training used here could produce a qualitative change in the form of the generalization gradient. In any case, the differences in training duration were also incorporated into the training regime used in the model. Post-experiment questionnaires were used to identify participants who admitted to solving the discrimination by counting a specific icon. It was hoped that by employing the accuracy criteria and this questionnaire, the remaining subjects would represent a group who were attending to the whole stimulus and were not basing their response decisions on an explicit assessment of the number of copies of a specific icon.

Method

Participants and Apparatus 32 undergraduate students from the University of Cambridge participated in the experiment. They were randomly allocated into the two groups Fixed and Varied. Participants were tested individually in a dimly lit room. The experiment was programmed using REALbasic software and run on a Power Macintosh.

Stimuli All stimuli appeared in a rectangular region in the centre of the computer screen measuring 2.8cm wide by 3.6cm high, surrounded by a thin white border. Each pattern consisted of an array of 12 icons, 3 icons wide and 4 icons high. The icons with which each stimulus was endowed depended on its position on the dimension, as shown in Table 2. For each subject, 12 icons were randomly assigned a position A to L on the dimension. By way of counterbalancing the icon distributions, each randomly ordered set of icons was allocated to one subject from each group.

In the Varied condition, the position of each icon within the array was randomized from trial to trial. For group Fixed the position of each icon was randomized on the first trial, but then remained constant for all subsequent presentations of that stimulus. For any pair of adjacent stimuli in the fixed condition (e.g. S_L and S_R) the icons common to both stimuli remained in a fixed location for both stimuli. Presentations of the icon stimuli were interspersed with filler trials consisting of uniform coloured rectangles that differed in hue. The colours varied in similar fashion to the artificial dimension, with two very similar shades of green presented during training, S_L and S_R, and a broader range of colours, from blue-green to yellow-green, presented during the transfer test. The primary purpose of the colour trials was to negate any effect of immediate contrast between patterns that would make the icon discrimination too easy.

Procedure Participants sat approximately 1m from the computer screen and were given verbal and written instructions pertaining to the training and test phases and the responses they would be required to make. For every trial, participants were required to make either a left or right key response by pressing either 'x' or '.' on the keyboard. During training, feedback appeared after every response in the form of the words "correct" or "wrong" flashing in the centre of the screen. During the test phase no feedback was given. If no response was given within 4 seconds of the stimulus appearing then the trial timed out and the message "no response" appeared on the screen.

As shown in Table 2, the stimuli at positions 4 and 5 were used as S_L and S_R respectively during the training phase. Trial order was randomized within blocks of 12 trials, containing three presentations each of S_L and S_R and three each of the corresponding filler trials, with the condition that trials alternate between icon stimuli and filler stimuli. To equate the groups for test accuracy, a total of 48 trials were presented in the fixed condition and 192 in the varied condition. During the transfer test a total of 8 stimulus positions were tested, with an equal number of interleaved colour filler trials. Trial order was randomized within blocks of 16 trials, containing one trial of each of the 8 transfer stimuli and 8 filler stimuli. There were 6 blocks, or 96 trials in total.

Results and Discussion

As discussed above, only subjects who performed above chance based on an average across all test stimuli were used in the analysis. A criterion was set at a mean test accuracy of 61.87% based on the 1-tail .05 level of

significance for a score out of 12 for each of the four (collapsed) stimulus positions. Only subjects who performed above this criterion were used in the analysis. In addition, participants who reported a specific icon counting strategy were also excluded. Only 8 participants in the varied group and 11 participants in the fixed group scored above the chance level criteria and did not report icon counting. The post-discrimination generalization gradients for these subjects are shown in Figure 4 below.

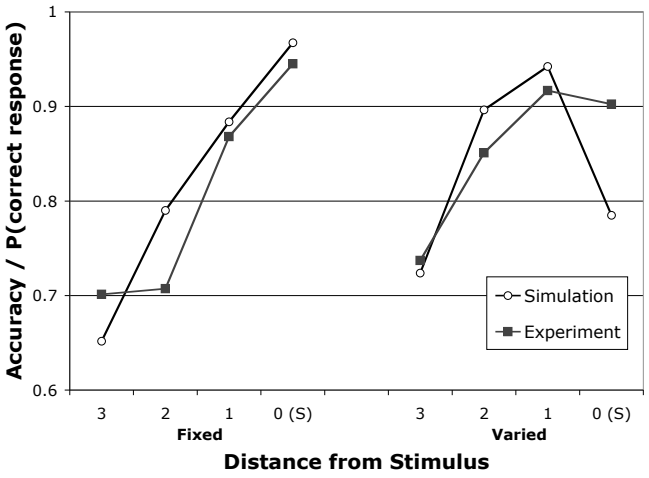

Figure 4. Experimental data and simulation results.

The model gives a reasonably good fit to the data. In the fixed group, post-discrimination generalization decreases rapidly as the test stimuli become systematically more dissimilar to the training stimuli. In the varied group, the generalization gradient resembles a negatively accelerating curve as accuracy peaks for test stimuli that are one step removed from S.

General Discussion

The simulation data predicts a very similar difference in post-discrimination generalization gradients between fixed and varied groups, under the conditions run in the current experiment and under conditions more closely resembling the pigeon experiments reported by Oakeshott (2002). While past models of icon-based peak shift have been unable to account for the differences in post-discrimination generalization gradients produced under fixed and varied conditions, this adaptation of the McLaren and Mackintosh model does so without resorting to the assumption that the stimuli have dimensional properties of any sort. Instead, the dimension-like qualities of the generalisation gradients and the peak shift effect emerge as a consequence of the process of associative learning. There is also no need to posit any shift from elemental to configural processing of the stimuli to explain the different post-discrimination gradients observed in the fixed and varied cases, an explanation suggested by Oakeshott (2002).

References

Blough, D. S. (1975). Steady state data and a quantitative model of operant generalization and discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 1*, 3-21.

Ghirlanda, S., & Enquist, M. (2003). A century of generalization. *Animal Behaviour, 66*, 15-36.

Honig, W. K., & Urcuioli, P. J. (1981). The legacy of Guttman and Kalish (1956): 25 years of research on stimulus generalization. *Journal of the Experimental Analysis of Behavior, 36*, 405-445.

Jones, F., & McLaren, I. P. L. (1999). Rules and associations. In *Proceedings of the Twenty-First Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Livesey, E. J. (2004). *Representation and Discrimination: An analysis of the curvature of post-discrimination generalisation gradients.* Unpublished 1st year report, University of Cambridge, Cambridge.

McLaren, I. P. L., & Mackintosh, N. J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior, 28*(3), 211-246.

McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning & Behavior, 30*(3), 177-200.

Oakeshott, S. M. (2002). *Peak shift: An elemental vs a configural analysis.* Unpublished PhD, University of Cambridge, Cambridge.

Pearce, J. M. (1987). A model of stimulus generalisation for Pavlovian conditioning. *Psychological Review, 94*, 61-73.

Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review, 101*, 587-607.

Spetch, M. L., Cheng, K., & Clifford, C. W. G. (2004). Peak shift but not range effects in recognition of faces. *Learning and Motivation, 35*(3), 221-241.

Thomas, D. R. (1993). A model for adaptation-level effects on stimulus generalization. *Psychological Review, 100*, 658-673.

Widrow, G., & Hoff, M. E. (1960). Adaptive switching circuits. In *Institute of Radio Engineers: Western electronic show and convention* (Vol. Convention Record, Pt 4, pp. 96-104).

Wills, S., & Mackintosh, N. J. (1998). Peak shift on an artificial dimension. *Quarterly Journal of Experimental Psychology Section B- Comparative and Physiological Psychology, 51*(1), 1-32.