Title
Rules and Associations

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 21(0)

Authors
Jones, F. W.
McLaren, I. P. L.

Publication Date
1999

Peer reviewed

Rules and Associations

F.W. Jones (FWJ1000@CUS.CAM.AC.UK)
I.P.L. McLaren (IPLM2@CUS.CAM.AC.UK)
University of Cambridge, Psychological Laboratory,
Downing Street, Cambridge CB2 3EB. U.K.

Abstract

Two-process theories of human cognition, that state that learning can occur by both associative and rule-based processes, are currently popular. We report two experiments which support such a view. Both employed a set of six stimuli which varied along a luminance dimension, and followed the same general design. That is, participants were trained to discriminate between the two stimuli in the middle of this set, before being tested on the whole set. In Experiment I, the length of training was varied. Following short training, participants' performance on test exhibited a peak-shift, and therefore may be explained in associative terms. After longer training, however, their behavior was consistent with rule-based learning. In Experiment II, the contingency during the training phase was varied. Participants in the 'Full Contingency' group performed in a manner consistent with rule-learning, while the 'Reduced Contingency' condition produced a peak-shift. These results are discussed in terms of McLaren, Green & Mackintosh's (1994) version of the associative/rule-based distinction.

Introduction

The idea that human cognition comprises both associative and rule-based processes has a long history, that stretches at least as far back as William James. Moreover, its current popularity is illustrated by the volume of literature devoted to the subject (e.g. the entire issue of Cognition 65). One recent incarnation of this 'hybrid' view can be found in McLaren, Green & Mackintosh (1994). Their two process model of human learning comprises: (i) an associative system that is sensitive to the statistical structure present in the surface features of the environment, and operates through the establishment and alteration of connections between representations; and (ii) a 'cognitive' process capable of rule abstraction, whose behavior resembles that of a symbolic logic machine.

There already exists a considerable body of evidence in support of such a dichotomy (see Shanks & St John, 1994; Shanks, 1995; and Sloman, 1996 for reviews). Briefly, in some categorization experiments participants' performance on novel transfer items is dependent on their similarity to the training exemplars (e.g. Perruchet, 1994), which suggests that the knowledge acquired is encoded in terms of surface features. Alternatively, under different conditions, the results from transfer tests and verbal reports are more consistent with participants having abstracted rules (e.g. Regehr & Brooks, 1993). Furthermore, this dissociation can be observed within the same experiment (e.g. Nosofsky, Clark & Shin, 1989), with some participants responding on the basis of similarity and others abstracting rules.

The aim of the research presented here was to further investigate the viability of this associative/rule-based distinction, and examine some of the conditions under which each process could dominate performance. Like the previous work described above, we differentiated between rule-based and associative learning by examining how participants' performance generalized to novel stimuli. To make this more concrete, consider a set of stimuli that vary along a hypothetical dimension, with a midpoint 'd'. Further, suppose that we create a category structure, such that all those stimuli to the left of d form one group, while those to the right form another; and train participants on one example from each category -namely R_{TRAIN} and L_{TRAIN}. An examination of participants' performance, when they are subsequently required to categorise stimuli spanning the entire dimension, should then allow these two forms of learning to be distinguished.

If their learning is rule-based, then we might expect them to abstract the rule: 'greater than d respond category one, less than d respond category two'. As a result, unless their performance is at asymptote, their accuracy may well be dependent on the stimulus' distance from the category boundary, with more extreme stimuli being classified more accurately. Alternatively, if their learning is just encoded in terms of surface-features, we might expect performance on the transfer stimuli to be dominated by their similarity to the training examples, and to exhibit a pattern known as 'peak-shift'. That is, as we pass along the dimension from the category boundary, we might expect accuracy to increase to a maximum and then drop off, with the peak being positioned further along the dimension than the training stimuli -hence the term peak-shift. Such a pattern of results was first demonstrated by Hanson (1959), in pigeons trained on a wavelength discrimination, and may be explained in associative terms.[1]

Consider the simple associative network illustrated in Figure 1B. The input layer comprises a bank of feature detectors, the pattern of activation across which represents the current stimulus being presented to the network. The activation of each of these detectors is a Gaussian function of stimulus' position on the dimension, with each unit

[1] The associative explanation offered here is an extension of the work of Spence (1932), Blough (1975) and Wills & Mackintosh (1998), respectively.

responding maximally to a different point on the dimension (see Figure 1A). The output layer comprises two units, corresponding to the two categories, and the two layers are fully interconnected. During training, on each trial either R_{TRAIN} or L_{TRAIN} is presented as input to the network, and the activation target for the appropriate category unit is set to one, while that for the other unit is zero. The weights are then updated using the delta-rule (e.g. McClelland & Rumelhart, 1985).[2]

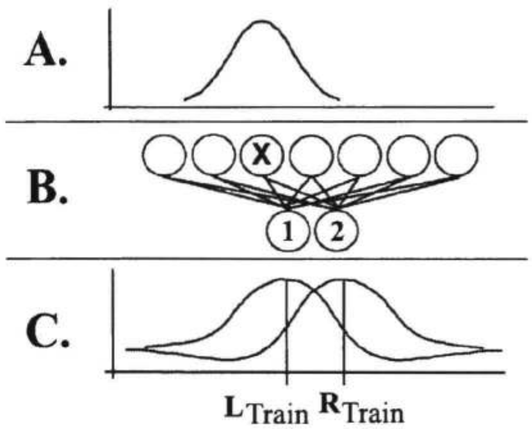

Figure 1: (A): The activation function for feature unit X. (B): A delta-rule network, with feature units on top and category units below. (C): Category units' activations on test.

On test, the network's output activations, in response to stimuli from different points on the dimension, can be measured. The results of just such a simulation are illustrated in Figure 1C. It can be seen that, if participants' responses are dependent on the difference in activation between the two category units, then they will be most accurate at positions outside the training examples, because that is where the difference between the curves is the greatest.[3] Thus, peak-shift can be understood in associative terms.

The differing predictions of associative and simple rule-based accounts are illustrated in Figure 2. Only half the dimension is illustrated, because the pattern should be symmetrical either side of the category boundary. As shown, a monotonically increasing trend indicates rule-based learning, and a peak-shift is diagnostic of associatively based performance. Thus, we have a way of distinguishing between rule-based and associative learning.

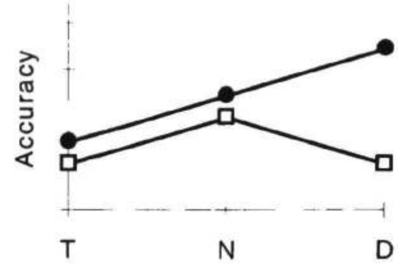

Figure 2: The predictions of associative (open squares) and simple rule-based (filled circles) accounts. T refers to training stimuli, while N and D denote test stimuli that are further out on the dimension.

Wills & Mackintosh (1998, Experiments 3a & 3b) have already performed experiments, on people, following the general design outlined above. For stimuli, they employed green rectangles that varied along a luminance dimension. In total, six different shades of green were used, divided equally between two categories -'dark' (D) and 'light' (L). When arranged in order of increasing brightness these stimuli can be referred to as $D_{DISTANT}$, D_{NEAR}, D_{TRAIN}, L_{TRAIN}, L_{NEAR} and $L_{DISTANT}$, respectively (see Figure 3). During training, on each trial participants were either presented with D_{TRAIN} or L_{TRAIN}, which they had to classify using one of two keys, before receiving feedback. Having learned this discrimination, they were transferred to a test phase. On each trial of this, any one of the six stimuli could appear and participants were again required to categorize them, using the same two keys. No feedback was given.

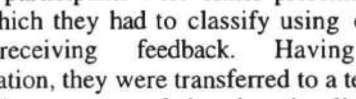

Figure 3: From left to right: $D_{DISTANT}$, D_{NEAR}, D_{TRAIN}, L_{TRAIN}, L_{NEAR} and $L_{DISTANT}$, respectively.

Wills & Mackintosh found that, during the test phase, participants responded more accurately to the 'Distant' stimuli than to the 'Near' ones, and were more accurate on the 'Near' stimuli than on those presented in training. In other words, moving along the dimension away from the category boundary, participants' performance improved monotonically. This, in conjunction with the fact that all the participants were able to verbalise a variant of the 'bright respond key one, dark respond key two' rule, suggests that their learning was rule-based. Wills & Mackintosh's procedure provides the basis for the two experiments reported here.

[2] This may be formally expressed: $dw_{CF} = S\, a_F\, (t_C - a_C)$ where 'dw_{CF}' is the change in the weight connecting feature unit 'F' to category unit 'C', 'a_C' is the activation of category unit c (which is equal to the weighted sum of the activation from the feature units), 'a_F' is the activation of feature unit 'F', 't_C' is the activation target for category unit 'C', and S is a constant that determines the rate of learning.

[3] In fact, a peak-shift may also be obtained if probability of classifying a stimulus into a particular category is dependent upon the ratio of that categories' activation to the total output activation, provided that noise is added to the system. Alternatively, if the output activations are first transformed using an exponential function, then a ratio rule will again produce the desired pattern. In addition, a 'winner-take-all' decision network (e.g. see Jones, Wills & McLaren, 1998) acting on these activations can also produce a peak-shift.

Experiment I

The aim of the first experiment was to establish whether varying the amount of training participants received would determine whether associative or rule-based processes dominated performance on test. According to McLaren, Green & Mackintosh (1994), it might be expected that a short training period would produce an associative pattern of results, with rule-based performance only emerging after greater experience of the contingencies. They argue that, initially, the stimuli-category associations will be too weak to support rule abstraction; but that, as training progresses, these traces will become sufficiently strong to enable the 'cognitive' process to use them as the basis for the development of rules.

In order to test this hypothesis, we employed a similar method to Wills & Mackintosh. There was, however, at least one important difference. That is, during both the training and test phases, the green stimuli were only presented on even numbered trials. On odd numbered trials, participants were required to perform a filler task. Specifically, they had to classify stimuli, that comprised a set of colored icons (see Figure 4), using the same two keys. Aside from the stimuli being different, the filler task was identical to that involving the greens. The inclusion of this additional task served to increase the difficulty of the initial discrimination, with the purpose of increasing the likelihood of obtaining an associative pattern of learning after short training. Without the filler task the discrimination would have been far easier, because the next green stimulus would have appeared immediately after the previous one had disappeared, allowing participants to compare them more directly.

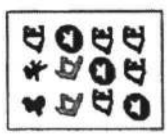

Figure 4: An example of an icon stimulus.

Given space constraints, and the fact that the results from the filler task do not bear directly on the question under investigation, details concerning the construction of the filler stimuli, and participants performance on them, will not be reported here. However, aside from making the discrimination harder to learn, we do not believe that the filler task affected the results for the greens, especially since its stimuli were identically generated for both conditions. Suffice to say that, their construction was identical to that in Wills & Mackintosh's Experiments 2a & 2b, and that the results for the filler task do not substantially differ from their findings.

The experiment comprised two conditions, namely, 'Short' and 'Long' training, which differed only in the number of trials of discrimination training the participants received. As already discussed, on the basis of McLaren, Green & Mackintosh we might expect associative learning to occur in the Short Training condition, with rule-based learning being manifest by Long Training participants. If this were the case, then we would expect the results for the former group to exhibit a peak-shift, compared to a monotonically increasing trend for the latter. Therefore, if these predictions are born out, performance on the Near and Distant test stimuli should differ between the two conditions, with Near being responded to more accurately than Distant after short training and vice versa for long training.

Method

Participants and Apparatus The participants were 58 Cambridge University undergraduates, whose ages ranged between 18 and 35. They were randomly divided equally between the two conditions, and did not receive payment for their help. The experiment was run on a RISC PC 700 computer, situated in a quiet room. Illumination was provided by a small desk lamp. The low light level was employed because pilot work suggested that some participants would be unable to discriminate between some of the shades of green under normal illumination.

Stimuli Both types of stimuli occupied a rectangle measuring 3.6 cm wide by 2.8 high, that was surrounded by a thin grey border. For the greens this rectangle was entirely filled in green. The luminance of these stimuli was determined by the value of a computer parameter, that ranges between 0 and 255. This was set to 50, 108, 137, 166, 195 and 253, for $D_{DISTANT}$, D_{NEAR}, D_{TRAIN}, L_{TRAIN}, L_{NEAR} and $L_{DISTANT}$, respectively. The filler stimuli comprised 12 icons arranged in a 4 by 3 grid, within the rectangle. See Figure 4 for an example, and Wills & Mackintosh (1998), Experiments 2a & 2b, for further details concerning their construction.

Design In the Short Training condition, training lasted for 48 trials, compared to 96 for the Long Training group. The order of stimulus presentation followed a pseudo-random sequence, such that L_{TRAIN} and D_{TRAIN} appeared 3 times during every set of 12 trials and only on even numbered trials. The test phase was identical for both conditions and comprised 120 trials. On test, the greens appeared in a random order, within the constraints that during every batch of 24 trials each of the six stimuli had to appear twice and only on even numbered trials. While the orders of presentation were designed in batches, there was no actual batching of the stimuli. In both training and test, filler stimuli appeared on odd numbered trials. The key assignments were counterbalanced, such that, for a random half of the participants in each condition, the 'x' key equalled 'light' and the '.' key equalled 'dark'. The remaining participants had the mapping reversed.

Procedure Participants sat approximately 1m away from the computer monitor, which was positioned roughly at eye-level. Some general instructions and ones specific to training phase were explained by the experimenter, before he left the room. Participants were informed that they would be performing two unrelated categorization tasks and that the computer would switch between the two on alternate trials. They were told to use the feedback to help them learn which key went with which stimulus. Participants initiated training by pressing the Spacebar. On each trial, the appropriate stimulus appeared in the centre of the screen, and, after 3 sec, the words 'Please respond now' appeared below it. Participants then had to respond using either the 'x' or '.' key, as quickly as possible, whilst avoiding errors. If they pressed the wrong one of these two keys, then the computer beeped. In addition, if they pressed a key other than 'x' or '.', or had not responded within 5 secs of the prompt's onset, then the stimulus was replaced by error messages. Respectively, these read: 'You have pressed an invalid key' and 'You did not respond in time'. Key presses prior to the appearance of the prompt were ignored. The next trial followed immediately after the response.

At the end of training, the computer displayed a message requesting the participants to find the experimenter, who then explained the instructions for the test phase, before again leaving the room. Participants were told to use whatever they had learned in training to classify the new stimuli they would see. The test phase commenced at their initiation. Test trials were identical to those in training, except there was now no prompt. Participants were informed that they could respond as soon as the stimulus appeared, and were again asked to be as fast as possible, whilst avoiding errors. No feedback, concerning the accuracy of their responses, was given. Since there was no prompt, the time-out occurred 5 sec after the beginning of the trial.

Following testing, participants again had to fetch the experimenter, who administered a structured questionnaire. This comprised a series of increasingly specific questions, designed to determine what strategies the participants had employed during the task and whether they could verbalise the underlying rule.

Results and Discussion

During training, unsurprisingly, participants in the Long Training condition responded significantly more accurately than those in the Short Training group (means 71.7 and 55.5 % correct, respectively; $t(56)=2.96$, $p<0.05$). The results for the test phase are shown in Figure 5. It can be seen that the accuracy of the Long Training group follows a monotonically increasing trend, while the means for the Short Training condition exhibit a peak shift. As already discussed, if our predictions are correct, then it is performance on the Near and Distant stimuli that should distinguish between the groups. Therefore, in order to assess whether the group differences were significant, the mean accuracy on Distant stimuli was subtracted from the value for Near stimuli, for each participant, and a planned contrast performed on the resulting scores. This demonstrated that the two conditions did, indeed, significantly differ ($F(1,56)=5.03$, $p<0.05$).

Given this, and in order to allow a more detailed examination, the data from the two groups was then analysed separately. Planned contrasts were used to compare performance on Training and Near stimuli, and performance on Near and Distant stimuli. Since the hypothesis we were testing made clear the directions of the expected effects, these contrasts were one-tailed. For the Short Training group, the contrasts revealed that Near stimuli were responded to significantly more accurately than to both Training and Distant ones ($F(1,28)=8.05$ and 3.29 respectively, one-tailed $p<0.05$ for both). Therefore, the peak-shift was reliable. With regards to the Long Training condition, accuracy on Near stimuli was significant higher than that on Training ones ($F(1,28)=28.37$, one-tailed $p<0.01$). However, while participants responded more accurately to Distant stimuli than to Near ones, this difference was only marginally significant ($F(1,28)=1.75$, one-tailed $p<0.10$). Nevertheless, the results for this condition are consistent with application of a rule, and clearly do not exhibit a peak-shift. Moreover, when questioned, all the Long Training participants were able to verbalise a variant of the underlying rule, as compared to none in the Short Training group.

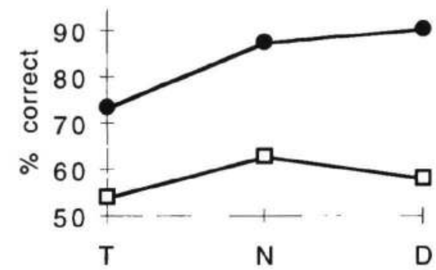

Figure 5: Participants' mean accuracy on test, for Experiment I. Open squares=Short Training, Filled circles= Long Training. T=Training, N=Near, D=Distant.

In summary, as predicted, the short training participants produced a pattern of performance, on test, that is explicable in associative terms. Following more lengthy exposure to the contingencies, participants were able to abstract the underlying rule, and their behavior was consistent with its application.

Experiment II

McLaren, Green & Mackintosh (1994) argue that as we increase the complexity of the contingencies between stimuli, so also we increase the amount of information that needs to be stored in working memory in order for a rule to be abstracted. Thus, as the difficulty of the to-be-learned mappings is increased, so the likelihood that performance will be associatively-based, rather than rule-based, increases.

The second experiment sought to test this prediction, using the same rationale as the first. It comprised two groups, namely 'Full Contingency' and 'Reduced Contingency'. The former was identical to the Long Training condition in the first experiment, while the latter differed only in that, in the training phase, the contingency between the two green stimuli and their respective categories was reduced. The was achieved by reversing the keys assigned to L_{TRAIN} and D_{TRAIN} on 25% of the training trials. Participants were not told about this manipulation.

If McLaren et al.'s predictions proved accurate, we would expect the Full Contingency condition to exhibit rule-based performance, and the Reduced Contingency group to produce a peak-shift. Finally, given that the Full Contingency group was a straight replication of the Long Training condition, and that in the previous experiment this did not produce a completely reliable monotonically increasing trend, a larger number of participants were run in this group. It was hoped that the resulting extra power would lead to a significant monotonically increasing trend.

Method

Participants, Apparatus, Stimuli and Procedure 80 new participants were drawn from the same pool as in Experiment I, split 60:20 between the Full and Reduced Contingency conditions, respectively. The stimuli, apparatus and procedure were identical.

Design The Full Contingency condition was identical to the Long Training condition in Experiment I. The Reduced Contingency condition differed from this only in that during training, on a randomly selected quarter of the L_{TRAIN} trials and arbitrary quarter of the D_{TRAIN} trials, the key assignments were reversed.

Results and Discussion

The data from Experiment II was analysed in exactly same way as that from the first experiment. A comparison of the training scores revealed that participants in the Full Contingency condition responded significantly more accurately than those in the Reduced Contingency group (means 67.4 and 54.9 % correct, respectively; $t(78)=2.87$, $p<0.01$). Figure 6 shows the mean accuracy results from the test phase. From this it is clear that, the trend in the Reduced Contingency group follows a peak-shift, while that produced by the Full Contingency condition monotonically increases. As previously, the reliability of these group differences was assessed by performing a planned contrast on the differences between Near and Distant stimuli. This demonstrated that the conditions were significantly different ($F(1,78)=5.28$, $p<0.05$).

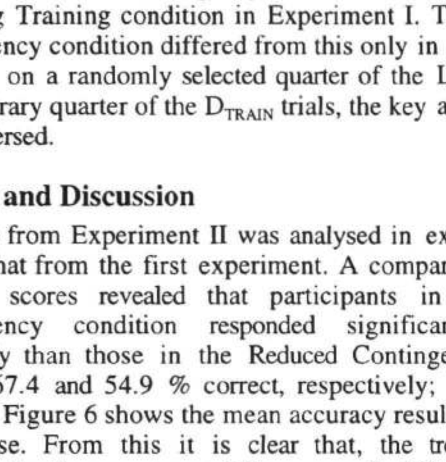

Figure 6: Participants' mean accuracy on test, for Experiment II. Open squares=Reduced Contingency, Filled circles= Full Contingency. T=Training, N=Near, D=Distant.

For each group, planned contrasts were used to compare performance on Training and Near stimuli, and performance on Near and Distant stimuli. Again, since the direction of the predicted effects was pre-specified, these contrasts were one-tailed. For the Full Contingency condition, these revealed that Distant stimuli were responded to significantly more accurately than to the Near ones, which in turn were responded to more accurately than the Training stimuli ($F(1,59)=5.51$ and 25.85 respectively, one-tailed $p<0.05$ for both). Therefore, the monotonically increasing trend was reliable, suggesting rule-based performance. Since this condition was a straight replication of the Long Training condition in Experiment I, this result compensates for the failure to find a completely reliable trend in that experiment, and reinforces the conclusion that Long Training participants were rule-learners. With regards to the Reduced contingency condition, participants were significantly more accurate on Near stimuli than Training ones ($F(1,19)=15.69$, one-tailed $p<0.01$), but the difference between Near and Distant stimuli was not significant ($F(1,19)=1.51$, one-tailed $p=0.12$). While this means that the peak-shift observed in this experiment was not completely reliable, it is clear that performance in this condition was more consistent with the associative predictions than the rule-based ones.

The structured questionnaire revealed that 63% of the participants in the Full Contingency group reported learning a version of the underlying rule, compared to 40% in the Reduced Contingency condition. A chi-squared performed on these scores demonstrated that, as predicted, this proportion was significantly higher for the Full Contingency group ($\chi^2(1)= 3.34$, one-tailed $p<0.05$).

To summarise, the Full Contingency group produced a rule-based pattern of performance, with the majority of participants also being able to verbalise a rule. This reinforces the findings of Experiment I. The results of the Reduced Contingency condition significantly differed from this, and exhibited a peak-shift trend. Moreover, significantly fewer participants in this condition were able to report the rule. This is consistent with the McLaren, Green & Mackintosh's prediction that reducing the contingency decreases the likelihood of rule abstraction, leaving associative learning to dominate performance.

General Discussion

We have argued that associative and rule-based learning, in this type of discrimination task, may be distinguished by examining how participants' performance generalises to test stimuli: rule abstraction being indicated by a monotonically increasing trend and associative learning by a peak-shift. In two experiments, we have shown that, by these criteria, both short training and a reduced training contingency produce associatively based performance, with rule abstraction only emerging after longer training with a 100% contingency.

These findings can be understood in terms of McLaren et al.'s (1994) cognitive/associative dichotomy: After little training, or exposure to a reduced contingency, stimulus-category associations will be too weak to support rule induction, but nevertheless strong enough to produce above chance test performance. With further training, the strength of these associative traces will increase sufficiently to enable rule abstraction, and the resulting rule-based knowledge will then dominate responding on test. Moreover, a similar peak-shift/rule-based dissociation has been found using a different task and type of stimuli (Aitken, McLaren, & Mackintosh, in preparation), suggesting that these findings apply more generally.

We will now address three possible criticisms of this work. First, perhaps apparent rule learning participants correctly guessed the rule on test, rather than learning it during training. This is a real possibility, since a light-dark rule is an obvious way to divide a set of stimuli varying in luminance. However, if the participants had just guessed the rule, then their performance as a group would not have been significantly above chance, because they would not have known which key to assign to bright and which to dark.

Second, maybe participants classified as rule-learners would in fact have shown a peak-shift had we tested them with more extreme stimuli from the dimension. We did not use such stimuli because they no longer appeared green in color, a fact which could have introduced some new artifact into the data. However, we are confident in our conclusion that they were rule learners, because the majority of them were able to verbalise the rule. Moreover, previous work suggests that increasing the amount of training should produce a peak-shift with its peak closer to the training examples (Aitken, et al., in preparation). This makes it unlikely that the Long Training/Full Contingency groups were showing a peak-shift, with its peak further along the dimension than we were testing.

Third, a more complex rule-based account can predict a peak-shift, enabling both patterns of performance to be explained in rule-based terms. Suppose that after short training, or exposure to a reduced contingency, the rule that develops is highly context dependent. If this were the case, then Distant stimuli might evoke the rule less than Near stimuli, resulting in a peak-shift pattern. However, it would be wrong to assume that this 'single' mechanism is more parsimonious than separate cognitive and associative processes, because it also comprises two processes -namely, the similarity based context activation of the rule, and the application of the rule itself. Neither is it clear that such a context sensitive rule-based account can explain the large body of evidence consistent with the cognitive/associative distinction (e.g. Sloman, 1996). Moreover, if we adopted this explanation then we would lose the ability to account for peak-shift in pigeons using the same mechanism, since we probably do not wish to ascribe rule-learning capabilities to them. In short, we believe that considering both pigeons and people together, the cognitive/associative explanation is the more parsimonious.

Finally, it should be made clear that we are not suggesting that learning must always be initially, purely associatively driven. Existing evidence suggests that this would be too simplistic a view for the real world (for a discussion see Keil, Carter Smith, Simons & Levin, 1998). Rather, we would argue that, in everyday life, learning occurs via some complex interaction between cognitive and associative processes. Clearly, attention now needs to be focused on further specifying these two processes, and the way in which they interact.

Acknowledgements

This research was funded by a BBSRC grant awarded to I.P.L. McLaren and an MRC research studentship awarded to F.W. Jones. The authors would thank Chris Ingram, Rob Leech and Daniel Reisel, who ran the participants for Experiment II; and an anonymous reviewer, for their helpful comments.

References

Aitken, M.R.F., McLaren, I.P.L., & Mackintosh, N.J. (In preparation).

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

Hanson, H.M. (1959). Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology,* 58, 321-34.

Jones, F.W., Wills, A.J., McLaren, I.P.L. (1998). Perceptual categorization: Connectionist modelling and decision rules. *Quarterly Journal of Experimental Psychology*, 51B (1), 33-58.

Keil, F.C., Carter Smith, W., Simons, D.J., & Levin, D.T. (1998). Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge. *Cognition*, 65, 103-135.

McClelland, J.L., & Rumelhart, D.E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, 114, 159-188.

McLaren, I.P.L., Green, R.E.A., & Mackintosh, N.J. (1994). Animal learning and the explicit/implicit distinction: Or why what we think of as explicit for us can be implicit for them. In N. Ellis (Ed.), *Implicit and Explicit Learning of Languages.* Academic Press.

Nosofsky, R.M., Clark, S.E., & Shin, H.J. (1989). Rules and exemplars in categorization, identification, and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 282-304.

Perruchet, P. (1994). Learning from complex rule-governed environments: On the proper functions of nonconscious and conscious processes. In C. Umilta & M. Moscovitch (Eds.), *Attention & Performance XV: Conscious and nonconscious information processing* (pp. 811-835). Cambridge, MA: MIT Press.

Regehr, G., & Brooks, L.R. (1993). Perceptual manifestations of an analytic structure: the priority of holistic individuation. *Journal of Experimental Psychology: General*, 122, 92-114.

Shanks, D.R., & St John, M.F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17, 367-447.

Shanks, D.R. (1995). *The Psychology of Associative Learning. Cambridge*, UK: Cambridge University Press.

Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119 (1), 3-22.

Spence, K.W. (1937). The differential response in animals to stimuli varying within a single dimension. *Psychological Review,* 44, 430-441.

Wills, S., & Mackintosh, N.J. (1998). Peak shift on an artificial dimension. *Quarterly Journal of Experimental Psychology*, 51B (1), 1-31.