Cue Competition in Human Associative Learning.

R.P. McLaren[1], F.W. Jones[2], F. Yeates[1] & I.P.L. McLaren[1] (i.p.l.mclaren@exeter.ac.uk)

[1]School of Psychology, University of Exeter, UK.
[2]Department of Applied Psychology, Canterbury Christ Church University, UK.

Abstract

There is a question as to whether cue competition effects can be observed in incidental learning paradigms in humans. The SRT and other location prediction tasks fall into that group often considered to show associative learning under incidental conditions. We used a standard SRT task in which the preceding two trials of a run of three predicted the third 2/3 of the time, and added another predictive cue, a colored square, which could also stochastically predict the next response required. The question was to what extent would these two cues compete in terms of incidental learning to make the next response faster and more accurate than controls? We assessed this by comparing the dual cue group to a color only control and a sequence only control. Our results showed that all three groups learned, and that the dual group learned about both cues at least as well as the individual controls, but that when switched to a test phase where each cue could be assessed independently, the dual group showed a marked decline in performance relative to the color control. We interpret this as evidence for overshadowing occurring between the two predictive cues in the dual group, such that when combined their performance is equivalent or superior to either control, but when assessed independently, the color cue actually has a weaker association to the outcome than the equivalent cue in the control group.

Keywords: SRT, color prediction, overshadowing, associative.

Introduction

Cue competition is one of the characteristic features of associative learning. Mackintosh (1976) showed that when a rat is trained with a combination of a light and a tone presented as a compound CS that predicts a shock US, then, while the animal learns to expect the shock US when the compound is presented, the animal shows less control of responding to the elements of the compound when they are presented individually in a test phase. We say that the two cues overshadow one another, and compete for associative strength to the US, such that the amount of associative strength accruing to either cue is less than if it had been trained on its own. Many studies have shown overshadowing between different CSs in animals such as the rat and the pigeon. The question addressed here is whether we can be confident that a similar phenomenon occurs in humans. This is an important point for those who wish to argue that humans and infra-humans share associative processes in common, and that we can understand this component of human learning and memory by studying other animals.

On the face of it, the answer to this question is a resounding "yes", based on the research currently available. To take a simple and rather common example, if we ask people to learn whether or not an allergic reaction occurs when a hypothetical patient eats meals made up of various food combinations, then they will happily do so by trial and error. If we then ask them to rate the foods for their propensity to bring about an allergic reaction in the patient, those foods trained individually will tend to be given higher ratings than those trained in compound with another food (e.g. see Le Pelley and McLaren, 2001), thus demonstrating a basic overshadowing effect. But there are a number of reasons why this result (which has been reported by many laboratories many times) is not as secure as it may seem. The first of these is that, in many cases, it is difficult to be sure that the result is due to associative processes rather than cognitive inference based on a heuristic of the type "if there are two cues predicting the outcome, then credit for this prediction must be shared between them". Researchers who subscribe to the view that humans possess a dual system capability when it comes to learning (as we do, see McLaren, Green and Mackintosh, 1994, for a summary), must take seriously the need to ensure that associative processes are being studied in a relatively pure form if statements about associative processing are to be made on the basis of that study. The Le Pelley and McLaren result already mentioned is perhaps one of the cases where this criticism might not apply, because in those experiments the authors were at pains to use conditions (high memory load due to using many cues and trial by trial presentation) that had been shown to encourage associative processing (see Le Pelley, Oakeshott and McLaren, 2005 for a discussion of this issue and a demonstration that these procedures are effective). But in many other cases, where few cues are used and memory load is low, the rating given may well owe more to cognitive inference than associative learning.

A second issue is that the stimuli that serve as the CSs in these experiments are too similar in kind, in that they are both foods. The analogy would be to an animal experiment in which the overshadowing was demonstrated to two tones, rather than a tone and a light. The former might give rise to concerns that the two tones when played together interacted in some way so as to change their stimulus quality, and that this interaction was lost when presented individually, so that the reduction in rating that occurred on test could be explained by some change in the perceived stimulus. No such process would apply when the stimuli were trained alone. It would clearly be better if the two CSs were different in kind so that this type of potential confound could be avoided.

A final point worth raising is that in all the experiments (that we are aware of) that have studied overshadowing that

come close to meeting our first two objections, the comparison has been between CSs trained in compound and tested individually, and a group or groups trained with the individual CSs and then tested. The problem with this procedure is that one group experiences a major change from training to test (the compound group) whereas the other does not. This, on its own, may be enough to depress responding in the compound group if they come to believe that circumstances have changed and deliberately alter their responses as a consequence (something that seems intuitively less likely to be the case in a rat or a pigeon). Once all these three objections are taken into account, we are unaware of any study that can be said to establish the existence of overshadowing as a characteristic of associative processing in humans.

Experiment

Shortly, we will describe a design that answers the three objections we have raised to existing demonstrations of overshadowing in humans. We start by considering what might be expected under compound training. If cues A and B are trained to predict an outcome, O, then an associative analysis analogous to that offered in animal learning studies states that both A and B will acquire associative strength for the outcome, O, but that this will be such that, when tested individually, responding to either element of the compound will be weaker than to the compound itself. We could add that responding to elements that had been trained individually would be expected to be stronger, on test, than that for elements trained in compound, and that the drop in performance from training to test would be much less (if any occurred at all) for these individually trained cues than for those trained in compound. With these characteristics of overshadowing clear in our minds, we can now turn to considering the details of a design that would enable us to test some of these predictions.

Our two classes of cue are chosen to have quite distinct characteristics. We will employ a basic SRT paradigm similar to that of Willingham, Nissen and Bullemer (1989), in which there are two circles that define two stimulus locations, left and right. The circles are outlines at the start of a trial, then one of them fills, and the corresponding key has to be pressed. This gives us a fast-paced choice RT task that allows little time for reflection. Unknown to our participants, in those groups that are given sequential information, there is a 2/3 chance of a trial being predicted by the two preceding trials. The rule is that if the two preceding trials are both the same, then that trial is likely to be a left, whereas if they are different, it is likely to be a right, with these response assignations counterbalanced across participants. Thus, the first type of cue is provided by the sequence of locations that occur / responses required. Our second cue type is provided by a colored square that flashes up before the circle fills in, at fixation between the two circles. Participants that receive color information have a 3/4 chance that the color will predict the response location on half the trials. On the other half of trials different colors are used that are not predictive and so can be used as color control trials. We settled on these parameters for the tasks as a result of extensive piloting and prior work, to ensure that both the sequential information and the color information were capable of supporting learning.

There are three groups in this experiment. Group Dual has both sequence and color information programmed in. Group Color has the same type of color information as Group Dual, and Group Sequence has the same type of sequence information as Group Dual. These last two groups will serve as our controls. Group Sequence are still shown a colored square just before the response location is indicated, but the color bears no relation to that location; equally Group Color experiences sequences of trials in just the same way as Group Dual, but these are not predictive. The point is that all groups experience a fast-paced sequence of trials cued by a colored square during both training and test, and so there should be no difference in their experience, except that brought about by the contingencies, and no obvious difference between training and test.

Method

Participants

90 University of Exeter students participated in this experiment, with ages ranging from 18-35. They were either given 1 course credit or paid £5 for a one hour session.

Stimuli

These were displayed on an iMac computer with participants seated about 70cm from the screen. The display consisted of two outline white circles (the two choice SRT stimuli) and a white outline square (which would become the color stimulus) on a black background. The circles had a diameter of 1.9 cm, and the square was also 1.9cm on a side and positioned in between the two circles, which were 2.2 cm to the right and left of its center respectively so as to be separated by the same distance as in Jones and McLaren (2009). The color stimuli were 1.9 cm square blocks of color that replaced the outline square just before one of the circle outlines filled. The colors used could be any one of red, green, blue or yellow. The signal to respond was given by a white filled circle 1.9 cm in diameter replacing one of the outline circles. Participants had to press a key on the side that the filled circle appeared, using either the 'X' (for left) or '.' (for right) keys on a QWERTY keyboard.

Design

The experiment consisted of three equal sized groups (Dual, Color alone and Sequence alone) of participants who each experienced 18 blocks (16 training and 2 test) of a two-choice SRT task over the course of an hour. For Group Color, during training the sequences were pseudo-random. This was accomplished by taking the 8 different triplet subsequences possible (RRR, RRL, RLR, RLL, LRR, LRL, LLR, LLL) and concatenating them in a random order to make a block of trials with equal numbers of each triplet in it. Two of the four colors were 75% accurate in predicting

the response location (the other two colors were 50:50; see Yeates, Jones, Wills, Aitken and McLaren, 2012 for details). This part of our paradigm was based on that developed by Aitken (1996). For Group Sequence, during training the four colors were 50:50 in predicting the response location, but the sequences that could occur were now constrained so that only four of the possible eight triplets were used to construct any given training block. The four triplets used had to conform to the rule, "if the first two trials are the same then the third is X, but if they are different then it is Y" where X could be one of L or R, and Y was the complementary response. Thus, one participant in this group might receive training blocks constructed by randomly concatenating eight of each of the triplets LLL, LRR, RLR and RRL to make a block of 96 trials. This has the effect of making the response location required predictable 2/3 of the time on average (see Jones and McLaren, 2009 for more details on this method). Group Dual simply had both sequence and color information, arranged so that when the sequential information was guaranteed to be predictive (on every third trial) so was the color information, thus maximising the correlation between the two cues. This arrangement ensured that the two cues were in agreement and predictive on 71% of trials on which one of them occurred, which compares well with their overall contingent relationship to the response location (sequence, 67%; color, 75%). Note that we ensured that the colors that were predictive in Group Dual occurred on the third trials in Group Sequence (but split 50:50 by response location so that they were not predictive), so that we could pick out these dummy "predictive color trials" and compare them to Group Dual to allow an estimate of color learning uncontaminated by sequence learning. In a similar fashion, we also ensured that the predictive colors for Group Color occurred on third trials (but the sequences were in a 50:50 relationship with the response location), so that we could use these trials as controls to assay a relatively pure measure of sequence learning in Group Dual. Previous work from our laboratory suggests that participants do not learn about the special status of third trials (e.g. Jones & McLaren, 2009).

After 16 blocks of training, participants went straight into a two-block test phase with no further instructions or indication that anything had changed. Both color and sequence information was now entirely non-predictive (50:50 in all cases) and uncorrelated with one another..

Procedure

Participants were asked to fixate on the square outline, and to respond as quickly and accurately as possible to the circle fills. They were told that the square would fill in with a color to warn them that the next response location was about to be signalled. No mention of any contingent relationships was given to any of the groups in this experiment. Thus, the order of events for a trial was that a colored square appeared centre-screen, and then after a variable interval of 250-500 msec one of the circles filled. These stimuli remained until a response was made then

cleared, leaving the circle and square outlines for a 250 msec RSI until the next trial started. If participants pressed the wrong key, a beep sounded to inform them of their error. At the end of the experiment participants were interviewed to determine if they had noticed any relationship between the colored squares and the response location, or that any of the sequences were predictive, then paid and thanked.

Results

Both RTs (for correct responses) and errors were recorded and analysed for both the training and test phases of the experiment. Sequence learning was analysed by comparing performance on the trained sequences to performance on the other set of four sequences. Thus, the sum of the mean RTs (or errors) for the trained sequences was subtracted from the sum of the mean RTs (or errors) of the complementary set. This measure has an expected value of 0, where higher positive scores indicate learning. It controls for sequential effects because both sets of four sequences contain the same set of transitions, simply transposing left and right responses, and so gives an unbiased measure of sequence learning. Our measure for color learning was to compute the mean RT (or error rate) for the predictive color trials and subtract this from the mean RT (or errors) of the control color trials. We did not use the inconsistent predictive color trials (the 25% that were "wrong") as our comparison, because there were relatively few such trials, making any such measure rather variable.
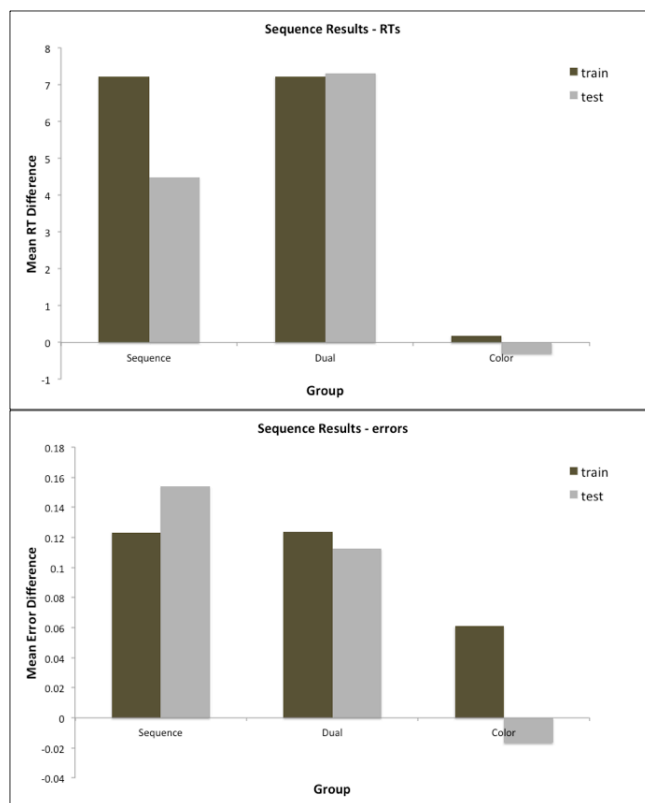

Figure 1. The top panel displays the mean differences (untrained sequences–trained sequences) for RTs. The bottom panel shows the equivalent error differences.

Sequence learning

We begin with an analysis of the sequential information. Figure 1 shows the mean difference between trained and untrained sequences in RTs (top panel) and errors (bottom panel) for each of the three groups over training and on test. Higher scores indicate more learning, so we can see that over the course of training both the Sequence and Dual groups show good evidence of sequence learning as measured by RTs and errors, whereas there is little learning in the Color control group as far as sequences are concerned. This impression is confirmed by planned comparisons against an expected value of 0, for Group Sequence RT differences, $F(1,29)=20.54$, $p<.001$, for the error differences, $F(1,29)=36.32$, $p<.001$; for Group Dual, $F(1,29)=28.62$, $p<.001$ (RTs), $F(1,29)=37.79$, $p<.001$ (errors); and for Group Color, max $F=2.87$, p=ns. Turning now to the Test phase, a similar pattern emerges, though reliability is weaker. Group Sequence, $F(1,29)=2.17$, p=ns (RTs), and $F(1,29)=5.35$, $p=.028$. (errors). Group Dual has $F(1,29)=12.74$, $p<.005$. (RTs), $F(1,29)=2.03$, p=ns. (errors), and Group Color has max $F(1,29)<1$. The reduction in reliability is not entirely unexpected given that the test phase is, in fact, an extinction treatment that will degrade the learning that has already taken place. Nevertheless, the evidence for sustained performance on the basis of what has been learned during training in Group Sequence and Group Dual, is something that we shall return to.

Color learning

Turning now to the color data, Figure 2 shows plots of the difference scores obtained by comparing performance for the predictive colors with the control colors for RTs and errors. In the case of Group Sequence, the 'predictive colors' are the same colors as those used in matched participants in Group Dual, such that the two groups are identical except that color has no predictive value in Group Sequence. The color assignments in Group Sequence were set by swapping occurrences of one predictive color (say red which predicted a left response) in a given Group Dual participant, with the other predictive color (say green which predicted a right response) until they no longer predicted a left or right response.

Analysis of these scores against an expected value of 0 reveals that, over the course of training, Group Sequence shows significant evidence of learning in the RTs (see later for an explanation), F(1,29)=10.06, p<.005, and also in the errors $F(1,29)=12.33$, $p<.005$; Group Dual shows learning in the RTs, $F(1,29)=4.32$, $p<.05$, and an effect in the errors, $F(1,29)=25.07$, $p<.001$. Group Color has a more complex pattern of effects. The error differences show significant learning, $F(1,29)=12.84$, $p<.005$, but the RT differences show an effect in the opposite direction, $F(1,29)=20.95$, $p<.001$. This suggests some form of speed/accuracy trade-off, with participants in this group slowing down on the predictive colors and making fewer errors. The reason they slow down on the predictive colors, might be that it is on

these colors that they tend to make errors, as they are inconsistent on 25% of their occurrences. We can check by looking at performance on the inconsistent color trials.
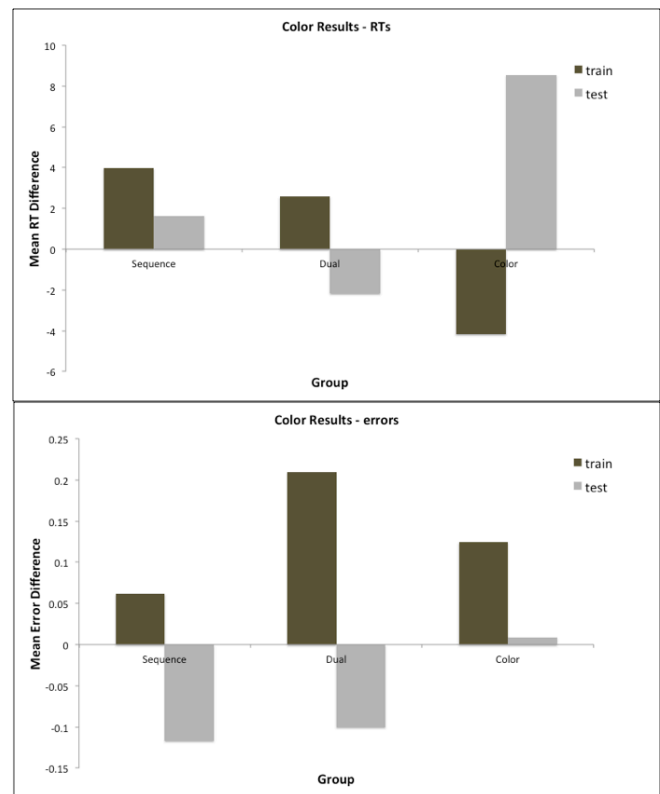

Figure 2. The top panel displays the mean differences (control colors–predictive colors) for RTs. The bottom panel shows the equivalent error differences.

If we do this then we find no evidence for slowing relative to controls, however, although there is a high error rate compared to consistent trials. Given this, we at present have no explanation for this speed/accuracy trade-off.

The reason why a similar slowing effect was not observed in Group Dual could be because in that group the effect is a joint one, as sequence and color information are highly correlated. Hence, the analysis of the color effect just given for Group Dual is highly likely to be contaminated by sequence learning. We can control for this by making use of the fact that the Group Sequence data has exactly the same sequence information as Group Dual, but none of the color information. This is why there is some evidence of learning during training in the color information analysis for that group. It simply reflects the correlated effect of sequence learning. If we contrast Group Sequence against Group Dual using the color scores shown in Figure 2 for training, then we find that whilst there is no significant effect in the RTs, there is one in the errors $F(1,58)=9.74$, $p<.005$. Hence, we can conclude that there has been learning of the color information in Group Dual over and above any effect of learning about the sequences.

There is a similar analysis we can do to check that participants in Group Dual actually learned the sequence

information and were not simply relying on color to predict the next response location. This time we compare the sequence scores in Group Dual to the color scores in Group Color, as the latter give the maximum effect that could be expected if color information was all that was driving the effect in Group Dual. There is no significant difference between the error differences (the difference is zero to two decimal places), but a highly significant effect in the RTs, $F(1,58)=48.92$, $p<.0001$. We can safely conclude that Group Dual learns something about the sequence structure in addition to what it learns about the ability of the colors to predict the next response location.

Finally, we can assess whether Group Dual learned more, overall, during training than the two control groups. If we contrast learning on the sequence measure for Dual vs. Sequence then we find that there is a numerical trend in that direction in both RTs and errors, but neither is significant (Fs<1). Turning to the color data, if we contrast Dual vs. Color then there is a significant advantage for Dual in the RT differences, $F(1,58)=15.02$, $p<.005$, and a trend in the same direction in the errors, $F(1,58)=3.26$, $p=.076$.

Analysis of the Test phase reveals a somewhat different pattern for the color data to that seen in the sequence measure, as there is a precipitate drop in performance by Group Dual, but this time maintained performance by Group Color. Only Group Color shows a significant effect, in the RT differences, $F(1,29)=6.61$, $p=.016$, this time with the error trend in the same direction. In fact, there is no evidence of any deterioration in effect in Group Color from training to test, but there is a substantial drop in performance in Group Dual so that it is numerically (though not significantly) in the wrong direction for learning (i.e. slower and more error prone on what were the predictive color trials).

The question that we must address now is whether there is evidence that Group Dual shows less learning on test than the appropriate controls. Some of the trends are certainly in this direction: For sequential information, contrasts of Group Dual against Sequence on the test data give Fs<1 for both RT and the error differences, so here performance in Group Dual is indistinguishable from that in the control group. But contrasting the color scores for Group Dual against Group Color on test gives an $F(1,58)=5.34$, $p=.024$ in the RT differences, and a numerical trend in the same direction in the error differences. Thus, we have significantly less learning exhibited on test in Group Dual than in the controls, even though there was strong evidence for learning of both color and sequence information in Group Dual during training.

General Discussion

We can summarize our findings by noting that there is good evidence that Group Dual learned to use both color and sequence information to predict the next response location. Group Sequence learned to use the sequence information available, and Group Color the color

information available (though here the evidence is somewhat mixed, and only resolved by looking at the test data). Overall, Group Dual was better than one of our two control groups (Group Color) at predicting the next response location during training, but worse on test compared to this control and indistinguishable from the other. What are we to make of this pattern of results?

We would argue that this pattern is one that is consistent with associatively-based cue competition in all three groups. The corollary of this position is that a cue does not have to be predictive to enter into cue competition, but does so automatically if it is present. In what follows we offer a simple error correcting analysis based on elemental representations of the type posited by McLaren, Kaye and Mackintosh (1989) and further developed in McLaren and Mackintosh (2000, 2002), in conjunction with the SRN architecture developed by Elman (1990) and further refined by Cleeremans and McClelland (1991). But a similar analysis follows from the use of Rescorla and Wagner type models (Rescorla and Wagner, 1972), or configural models such as Pearce (1987), or hybrid models such as McLaren, Forrest and McLaren (2012). In short, we believe that our analysis is generally applicable and speaks to something fundamental about associative learning in humans.

If we begin with both control groups, and take Group Color as an example, then what happens is that an association builds up from the representation of the predictive colors to a representation of the contingent stimulus location, over the course of training. This is simple enough, and a similar thing happens in Group Sequence as the sequential contingencies are learned, though here a recurrent architecture is needed to allow our model to capture these contingencies. But the non-contingent sequence information (in Group Color) and color information (in Group Sequence) is also present on each trial. Because of this, it will also form essentially spurious associations with representations of stimulus locations, and act as "noise", slowing learning of the contingent information. Group Dual does not suffer from this problem. Both types of information are predictive, and are learned. Other things being equal, we can expect learning to proceed more rapidly to this combination of cues, but the association to either cue alone will be less than is developed in the group where that is the only contingent cue because of overshadowing. This simply reflects the idea that there is an asymptotic net associative strength needed to predict an outcome in an error correcting system. If two cues are involved, then they have to split this maximum associative strength between them. But if only one is present, then, even if learning is slower, it can eventually achieve an associative strength that is nearer that maximum, which is in excess of that achieved in the dual cue case. Hence, on test, when each cue's associative strength is independently evaluated, the prediction is that Group Dual cues will show less evidence of learning than the control groups, as was the case in our study.

Of course, this is not quite the pattern of results we have obtained in our experiment, and here we have recourse to the idea that the sequence information is a stronger, more salient cue than the color information in our experiment. Because of this, the sequence cues overshadow the color cues to some extent and acquire a disproportionate share of the associative strength available. This still leaves Group Dual with an overall somewhat higher associative strength compared to either control, but the decrement on test when the sequential and color-based information are decoupled is particularly sever for the latter (see Mackintosh, 1976 for a demonstration of the asymmetric effects obtained when cues of differential salience are trained in compound).

We do not believe that our demonstration of overshadowing is contaminated by any kind of cognitive inference because the task was fast paced, allowing little opportunity for reflection, and our post-experiment questionnaire indicated that none of our participants were aware of the contingencies in play. We do not believe that there was any cognitive or sensory interference between our two types of cue because they are quite different in nature and occupy different domains of representation. We also do not believe that participants changed the way they responded on test, because there was no indication of entering a test phase, and participants in all groups experienced a similar degradation in contingency. For all these reasons, we believe that our results give good reason to take the view that overshadowing occurs in human associative learning, and that it occurs automatically.

Conclusion

One of the fundamental questions about human learning is whether it can have an associative component similar to that found in other animals. Our results indicate that learning under incidental conditions, learning that is automatic and not based on verbaliseable knowledge, is subject to the phenomenon known as overshadowing in the animal learning literature. This strengthens the case for people possessing associative processes in addition to other rule-based processes as part of their cognitive machinery.

REFERENCES

Aitken, M.R.F. (1996). *Peak shift in pigeon and human categorisation.* Unpublished PhD Thesis, University of Cambridge.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General, 120,* 235–253.

Elman, J. L. (1990). Finding structure in time. *Cognitive Science, 14,* 179 –211.

Jones, F.W. and McLaren, I.P.L. (2009). Human Sequence Learning Under Incidental and Intentional Conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, 35, 538-553.

Le Pelley, M. E., & McLaren, I. P. L. (2001). Retrospective revaluation in humans: Learning or memory? *Quarterly Journal of Experimental Psychology,* 54B, 311-352.

Le Pelley, M. E., Oakeshott, S. M. and McLaren, I. P. L. (2005). Blocking and unblocking in humans. *Journal of Experimental Psychology: Animal Behavior Processes*, 31

Mackintosh, N.J. (1976). Overshadowing and stimulus intensity. Animal Learning and Behavior, 4, 186-92.

McLaren, I. P. L., Forrest, C.L., and McLaren, R.P. (2012). Elemental representation and configural mappings: Combining elemental and configural theories of associative learning. *Learning & Behavior*, 40 (3), 320-33.

McLaren, I.P.L., Green, R.E.A. and Mackintosh, N.J. (1994). Animal learning and the implicit/explicit distinction. In N.C. Ellis (Ed.). *Implicit and explicit learning of languages.* Academic Press.

McLaren, I.P.L., Kaye, H. and 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.* Oxford. OUP.

McLaren, I.P.L. and Mackintosh, N.J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning and Behavior, 38(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.

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

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non- reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.

Willingham D.B., Nissen M.J. and Bullemer P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 1047-1060.

Yeates, F., Jones, F.W., Wills, A.J., Aitken, M.R.F. and McLaren, I.P.L. (2012). Implicit Learning: A Demonstration and a Novel SRT Paradigm. In: *Proceedings of the 34th Annual Conference of the Cognitive Science Society*, Sapporo, Japan.