Implicit Learning: A Demonstration and a Novel SRT Paradigm

Fayme Yeates[1] (fy212@exeter.ac.uk)

F.W. Jones[2]   A.J. Wills[1]   M.R.F. Aitken[3]   I.P.L. McLaren[1]

[1]School of Psychology, College of Life and Environmental Sciences,
University of Exeter, UK.
[2]School of Psychology, Canterbury Christ Church University
[3]School of Psychology, University of Cambridge.

Abstract

Evidence for human learning without awareness of what is learned has been sought in serial reaction time (SRT) tasks in which, unknown to participants, the locations of stimuli follow a particular rule or sequence (Willingham, Nissen & Bullemer, 1989). A number of criticisms have been levelled at such tasks, including a lack of adequate control for sequential effects and a discrepancy in sensitivity between measures of implicit and explicit knowledge about the task (Jones & McLaren, 2009; Shanks & St. John, 1994). In this study we provide a novel, two-choice SRT paradigm whereby the locations of the response stimuli are sometimes predicted by a separate set of stimuli on screen. A color-filled square appears before each stimulus requiring a response, with participants informed this is simply a fixation point to prepare for the next trial. Two out of eight colors are predictive on 80% of trials, and performance on these consistent trials was faster than on the other six colors that were equally likely to result in either of the two possible responses. All these trial types were faster and more accurate than the remaining inconsistent 20% of trials for the predictive colors, which also produce more errors than control colors. A prediction task and interview followed the task, on which participants performed at near (slightly below) chance levels. We suggest that this task is a useful tool for studying associative learning in humans, as it provides reliable effects that appear to demonstrate implicit learning with relatively brief training.
Keywords: Associative learning; implicit learning; SRT task

Introduction

In 1994 Shanks and St. John published their seminal review of implicit learning in Behavioral and Brain Sciences. Their conclusion was that implicit learning in humans had not yet been conclusively demonstrated. One paradigm that seemed to hold some promise in this regard was the SRT task pioneered by Willingham, Nissen and Bullemer (1989). In this task participants had to simply respond to one of four possible stimulus locations by pressing a spatially compatible key. Unknown to them, the sequence in which the stimulus locations were presented was not random but instead repeated a particular ordering. Willingham et al found that participants trained on such a sequence learned to respond faster than controls, yet were unable to say much about the sequence or, importantly, to predict the next stimulus location given a number of preceding trials. Shanks and St. John analysed this result, and others like it, and concluded that whilst the prediction task was potentially a sufficiently sensitive method of assaying participants' knowledge of the sequence, the evidence suggested that the non-significant findings on the prediction task had more to do with a lack of power than a lack of knowledge. Subsequent reviews of more recent evidence (e.g. Lovibond and Shanks, 2002; Shanks and Lovibond, 2002; Mitchell, De Houwer and Lovibond, 2009) have seen no reason to change this conclusion.

In this paper we attempt to provide the evidence needed to prove that learning can take place without awareness of what is being learned. We do this using a variant of the SRT task employed by Jones and McLaren (2009) developed by Aitken (1996) which, instead of using the preceding trials to determine the later ones, uses a separate stimulus to predict which of the location stimuli will be presented next. A colored square presented between two stimulus locations is the predictor stimulus, and which of the two stimulus locations changes from an open to a filled circle is what determines the response required. If the left circle fills, then a left key press is needed and if the right-hand circle fills then a right key press is appropriate. The colored square is presented just before a response is required (while both circles are unfilled) so that it can predict which response will shortly be needed. Rather than run a separate control group, we follow Cleeremans (1993) in using a within-subject control. Thus some colors are never predictive (they have a random relationship with the stimulus location) whereas others are correlated with the response needed on the next trial.

The detailed design of this paradigm was dictated by our assumption of a dual-process mechanism for human learning. Following McLaren, Green and Mackintosh (1994), we conceptualized learning as being driven by two quite different sets of processes, one Cognitive in nature, employing conscious, controlled, rule-based symbolic processes, the other Associative in nature, automatic in operation and based on simple algorithms that capture the correlations between events. The challenge, then, was to arrive at a paradigm that would readily allow learning on an associative basis but would be much less amenable to rule-based cognition. The Aitken SRT paradigm allows for a parameterization that meets these requirements. We used 8 colors in total, but made only 2 of them predictive (one left, one right). Thus, 6 out of the 8 bore no relationship to the next stimulus location to occur. Following the technique used by Posner and Snyder (1975), the 2 predictive colors were themselves only 80% reliable. Hence, if one of them typically predicted a left response would be the next needed, 20% of the time it was followed by a stimulus location that

required a right response. Our rationale for doing this was that it would make conscious detection of the contingencies in play much harder. Due to 6 of the 8 colors being non-predictive, the overall prediction rate possible in this experiment is a mere 57.5% if complete knowledge of the contingencies is assumed, and no stimulus is an entirely reliable predictor of anything. In these circumstances, we expected it to be very difficult to notice which colors had some predictive value. Any attempt to "work out" what was going on in this fast paced task, with many different colors involved, would overtax working memory and severely impair performance on the main, SRT task, which should lead to this type of strategy being quickly rejected.

We also used the fact that there were 6 non-contingent colors to configure a design in which half the blocks were control blocks, and the other half experimental. The latter included the 2 contingent colors and 2 of the non-contingent, whereas the control blocks had the other 4 non-contingent colors in them. Control and Experimental blocks were alternated. This meant that for much of the time our participants were not experiencing any contingency between the colored square and the stimulus location / response required. It also allowed us to define two, somewhat different, within subject controls for our experimental manipulation. The first are the control colors in the experimental blocks: these have the advantage of sampling our participants' behavior under similar circumstances to those in force for the contingent, experimental colors. The second are the colors in the control blocks: these have the advantage of sampling behavior over an entire block without any distortions caused by the contingent colors. The control blocks were given exactly the same sequence of right and left responses as the corresponding experimental blocks. Thus, any sequential effects, caused by a particular run of right and left responses, will be controlled for by these blocks. As we shall see, this proved to be a useful aspect of our design.

Despite our efforts to prevent conscious rule-based learning, the parameters chosen for this design should nevertheless support strong associative learning. The colors occur just before the stimulus location is presented and the response required, and feedback is available shortly afterward. Thus, given that there is a reliable contingency between a clearly perceived stimulus and stimulus location / response, it should be easily learned. We ensured that the colors were presented boldly at fixation and gave our participants good reason to be looking there at the start of each trial. If this type of learning is automatic, and stimulus specific, then there should be little difficulty in learning the association between the two predictive colors and their correlated stimulus locations / responses.

Method

Participants
The study was conducted with 32 participants, randomly divided into two groups (n=16 in each) who performed the task with slightly different parameters (different distances

between the stimuli on screen). The participants were all first year psychology students at the University of Exeter, aged from 18 to 22 years old, who were rewarded with a course credit in return for their participation.

Materials
An Apple iMac computer was used to run the experiment, with participants seated roughly 50cm from the screen. The display during blocks contained two white, outline circles and a white, outline square on a black background. The circles were 1.9cm in diameter, and the square was 1.9cm in width. The square was positioned directly in the centre of the screen for both groups. One group saw the circles positioned 2.2cm to the right and left of the centre of the screen, consistent with the distance separating the two-choice SRT task stimuli in Jones & McLaren (2009). The other group saw the same size white circle outlines placed 7.5cm from the right and left of the centre of the screen, following the stimulus locations specified in Aitken (1996).

The predictor stimulus was a colored filled square 1.9cm in width that filled in the white square outline. A choice of eight possible colors: red, green, blue, yellow, pink, orange, brown and teal were used. The response signal was a white filled circle 1.9cm in diameter that replaced either the right or left outline circle during the trials. The participants were instructed to press the spatially compatible "x" key on a QWERTY keyboard if the target stimulus appeared on the left, and the ">" key if the stimulus appeared on the right.

Design
The experiment consisted of a two-choice SRT task conducted over one session that lasted approximately an hour. The task consisted of 20 alternating blocks, half of which were control blocks and half experimental, each with 120 trials. The ordering of the blocks was counterbalanced across participants. On both experimental and control blocks participants received a random ordering of equal amounts of right and left response stimuli. In the experimental blocks participants saw four colors (red, green, blue, yellow) as the filled square predictor stimulus. One of these colors preceded a right circle fill with an 80% contingency, and another color appeared before the left circle filling 80% of the time. The other two colors in the experimental block occurred with equal probability before a right or left circle fill. The colors were randomized for each participant and the experimenter was blind to which colors were predictive, and which response they predicted. During control blocks, all four colors (pink, orange, brown, teal) were equally likely to occur before a right or left response signal. Each color that could appear in a block appeared on an equal number of trials in that block, and all trials were constructed so that a color would not repeat itself on the next trial.

Procedure
Participants were instructed to fixate on the colored square and then to simply respond as quickly and accurately as possible to the white circle fills. They were told to use the

colored square to bring their focus back to the centre of the screen and therefore improve their overall performance on the task by avoiding bias to one side or the other. They were informed that the colors changed to attract their attention back to the centre of the screen: no mention was made of any contingencies or the role of the colored square as a predictive or instructive stimulus.

On each trial the square would fill with the stimulus color and remain on screen until the computer detected a response. A variable interval in the range of 250-500ms would then occur before the white circle response signal appeared on screen. Reaction time was measured from this stimulus' appearance on screen until a key press was detected. A 250ms response-stimulus interval (RSI) then followed before the appearance of the next colored square stimulus during which the screen was blank except for the square and circle outlines. On each trial the stimuli remained on the screen until the participant had responded or was timed-out for not having pressed a key within 4.25s.

If participants pressed an incorrect key or were timed-out then the trial terminated and the computer issued a short 'beep' sound. Following each block, participants rested during a 30 second break in which they were shown their average reaction time (in milliseconds) and their accuracy (as a percentage) for the block of stimuli just completed. They were also informed whether these scores were better or worse than those from the previous block. At the end of the 20th block, participants were instructed to fetch the experimenter.

A short verbal interview asking general questions about the task was given, and participants were then assessed to determine the extent of their knowledge about the contingencies during the experiment. They were first asked to describe any relationship that they had noticed between the colors and the circle fills. Then a prediction task was conducted that closely followed the procedures employed in the main task they had just experienced. Thus, participants were instructed to attend to a colored square at fixation in the middle of two outline circles, but now had to make a prediction about which circle they thought would fill by pressing the appropriate key. They were told that the circles would not fill, and that no response would be considered an error, apart from pressing keys other than the two response keys, and that there was no time limit on making their choice. The prediction task was presented as two short blocks of 16 trials per block, with each color from both control and experimental blocks occurring twice in each block in a randomized order.

Results

Results were computed for both errors and RTs, with four Stimulus Types compared across Blocks and at the different Distances between the two response signals. Those colors that had an 80% contingency with a right or left circle fill (the Experimental colors) are split into both their Consistent and Inconsistent presentations. For example, if a red square preceded a right circle fill on 80% of the trials, those

responses would be Consistent, whilst when a red square was followed by a left circle fill on the other 20% of trials this would be Inconsistent with the trained contingency. The two other colors in the experimental block form one control set (denoted Control – Experimental Block) for within-subject comparison, as do the four control colors from the control blocks (Control – Control Block). Trials following an error were excluded from both RT and error analyses. The RT and error data are shown in Figures 1 and 2, respectively.
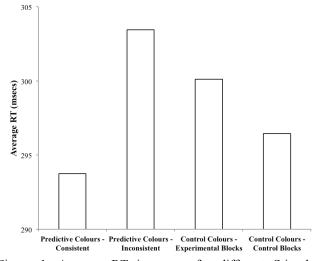

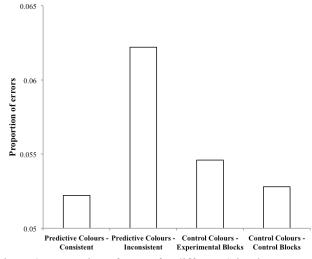Figure 1. Average RT in msecs for different Stimulus Types.

Figure 2. Proportion of errors for different Stimulus Types.

ANOVAs were conducted on both RT and proportion of errors, with Block, Distance and Stimulus Type as the independent variables. Blocks were collapsed into fours for the analyses, so each Block comprised two Experimental blocks and two Control blocks. The Distance of the stimulus did not have a main effect on the RTs, $F(1,30) = .27$, $p = .61$, nor on the errors, $F(1,30) = .56$, $p = .46$. Distance did not interact with any other variable, and therefore the location of the response stimuli was not a significant factor in determining learning of the color contingencies.

Figure 1 shows the average RT and proportion of errors for the four Stimulus Types, the main effect of Stimulus Type was significant for both the RTs, $F(3,90) = 17.45$, $p < .001$, and errors, $F(3,90) = 3.72$, $p = .014$. Both RTs and errors follow the same ordinal pattern, which was examined in more detail with a set of planned contrasts. Faster responses and fewer errors were made on Experimental Consistent trials. The difference between Experimental Consistent and Inconsistent trials is significant for both RTs, $F(1,30) = 33.15$, $p < .001$, and errors, $F(1,30) = 5.463$, $p = .026$, with Inconsistent stimuli responded to more slowly and with lower accuracy. Experimental Consistent stimuli are also responded to faster than both Control stimuli from Experimental, $F(1,30) = 34.46$, $p < .001$ and Control blocks, $F(1,30) = 6.173$, $p = .019$, this trend is also apparent in the errors, but it is significant only in the RTs.

Experimental Inconsistent trials are slower than Control stimuli from Experimental blocks in RTs, $F(1,30) = 4.85$, $p = .035$, and show a trend toward more errors, $F(1,30) = 3.34$, $p = .078$. Experimental Inconsistent trials are also slower and less accurate than Control stimuli from the Control blocks, significantly so in RTs, $F(1,30) = 14.51$, $p = .001$, and again demonstrating a trend in the errors, $F(1,30) = 4.05$, $p = .053$. It seems therefore that learning of the contingencies occurred, with performance averaged across the experiment showing faster and more accurate responses to those trials that were predicted 80% of the time by the preceding color. This is further highlighted by the decrement in performance for those trials inconsistent with the learnt relationship, causing participants to make more errors and respond slower to these stimuli than both the Consistent stimuli and the Control stimuli.

The Control stimuli are interesting, in that they significantly differ from one another in the RTs, $F(1,30) = 9.16$, $p = .005$, although not in the errors, $F(1,30) = .62$, $p = .44$. Control stimuli from the Experimental blocks are responded to more slowly than Control stimuli from the Control block. We suspect that this is due to the presence of the Consistent and Inconsistent stimuli in the Experimental Block. It is not that more errors are being made (overall) in the Experimental blocks, instead, it may be that the conflict that occurs on inconsistent trials engages mechanisms that produce more cautious responding, though we note that this cannot simply be a speed-accuracy trade-off as the error rate is, if anything, lower for the Control block stimuli than for Control stimuli from the Experimental Blocks.

Block has a significant main effect in RTs, $F(4,120) = 9.26$, $p < .001$, and errors, $F(4,120) = 19.02$, $p < .001$. With practice, participants got faster at the task but exhibited something of a speed-accuracy trade-off, making more errors progressively throughout the experiment. Block did not interact with Stimulus Type in the RT analysis, thus Experimental Consistent stimuli in Experimental Blocks were responded to faster than Control stimuli in Control Blocks, which were in turn responded to faster than Control stimuli in Experimental Blocks, and all are faster than Experimental Inconsistent stimuli. The relative positions of Inconsistent and Consistent stimuli are reversed in Block 1 in the errors, however, and the interaction between Block and Stimulus Type for errors is significant, $F(12,360) = 2.68$, $p = .002$. Thus, in this case, the Consistent / Inconsistent difference emerges over blocks.

The structured questionnaire and prediction task were completed by half of the participants, eight at each Distance value. The other participants were simply asked to pick which colors were predictive: they were invited to select two colors from the eight possible candidates. 14 of the 16 felt able to make the attempt, and out of the 28 responses that were generated 6 were correct (7 would be expected by chance).

The results of the structured questionnaire for the other 16 participants who completed the prediction task show that of the 32 responses, 8 were correct. If we take it that there are two predictive colors, and do not require the participant to remember which predicted left and which predicted right, then if we ask the participants to pick two colors out of the eight, the expected number of colors selected correctly is 0.5. With 16 participants taking part the total number of colors correctly selected is expected to be 8 by chance, and it is exactly 8, approximately evenly distributed across the different responses (right or left) and Distances.

Turning now to the prediction task itself, each participant experienced 32 trials in total. 8 of which involved the two predictive colors. Taking into account the requirement to give the appropriate response for the color, we can expect 4 correct responses by chance to these 8 presentations. Inspection of Figure 3 shows that the distribution of correct responses across participants is approximately binomial at each Distance, and is centered at a mean of 4 or less.
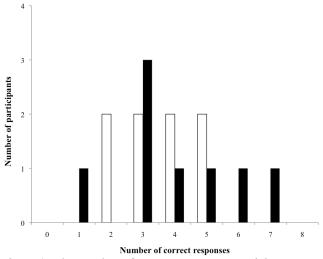

Figure 3. The number of correct responses participants gave for Consistent – Predictive colors. Data presented from 8 participants who performed the task with 2.2cm Distance (open bars) and 8 participants who performed the task with 7.5cm (solid bars) between response stimuli.

The actual means are 3.5 and 4 for Distances of 2.2 cm and 7.5 cm respectively, neither of which differ significantly from chance, and neither of which exceed mean chance

expectation. Our conclusion, based on these data, is that our participants do not know which colors are predictive, and do not know what they predict. This is in accord with their self-reports at the end of the experiment: they claimed to be unaware of any contingency between the color of the central square and the response that followed.

General Discussion

Our basic claim is that our participants' behavior was affected by the contingencies between color and stimulus location / response in the absence of any conscious knowledge on their part about these contingencies. That their behavior was affected is beyond doubt; the effect is statistically highly significant and, by the standards of these things, rather easy to obtain after relatively little training. We believe that these characteristics make this paradigm a good candidate for studying associative learning in humans. In particular, the highly reliable difference between Consistent and Inconsistent responses to the contingent colors is noteworthy, and our results suggest that this difference is made up of a disadvantage for Inconsistent trials and an advantage for Consistent trials.

Our use of two different types of within-subject control enables us to be confident that the advantage for Consistent trials and the disadvantage for Inconsistent trials are real. They also allow us to speculate about the effect of having these types of trial in a block. In some ways the most appropriate controls are those colors used as controls in the Experimental blocks. They are subject to the same local conditions, in terms of motivation, mood etc., and so provide an appropriate baseline for comparison with Consistent and Inconsistent responses to the contingent stimuli. We included the separate control blocks because we suspected that simply having some predictable trials in a block might alter our participants behavior, and also to provide a control for sequential effects. This last was accomplished by using exactly the same sequence of left and right responses in paired Control and Experimental blocks, the only difference between the two was in what colors were used in the block, and how they were paired with responses. Thus, in a Control block our participants would have exactly the same sequence of responses to make, but no useful information from the colored squares that were presented to help (or hinder) them. The fact that our contingent stimuli also differ from these controls on Consistent and Inconsistent trials gives us considerable confidence that there is no hidden artifact responsible for our results.

The difference between the Experimental control stimuli and the Control block control stimuli is also revealing. We hypothesize that this difference, with Experimental block controls performing worse than Control block controls, is due to the conflict experienced on Inconsistent trials leading to a greater cautiousness on some trials (perhaps best described as some kind of post-stimulus presentation checking process) that does not shift performance in terms of any speed – accuracy trade-off. By this we mean that the

slower RTs do not lead to higher accuracy, but instead reflect an additional "cost" that is incurred in task performance. It is difficult to speculate any further on what the exact nature of this cost might be, but one possibility is that perceptual information specifying which stimulus location had occurred might continue to be gathered post-decision on some proportion of trials as part of the checking process. If this mostly occurred on trials where the decision was the correct one (which given the low absolute error rates of our participants is quite likely), then it would have relatively little impact on the error rate but a significant impact on RTs. This analysis can only apply if the chance of engaging the checking process is unrelated to the probability of making an error, and low enough to result in a low frequency of occurrence on trials that result in an error. In these circumstances our RT measure would be much more sensitive to such a process than the error estimate (because the latter would be based on so few data points), so a result in the RTs alone would be quite possible. All speculation aside, this result is, as far as we are aware, a novel one, and opens up a new line of enquiry that may well be related to recent work by Verbruggen et al (in press).

Now we turn to some consideration of the arguments that may be raised to suggest that our prediction test results do not prove that our participants had no conscious knowledge of the contingencies. The most obvious is that the prediction test results apply to only half the participants. This objection is of no concern, as the effects we report, i.e. the advantage for Consistent vs. Inconsistent, the advantage for Consistent trials relative to both types of Control trial and the disadvantage for Inconsistent trials are all significant if only the 16 participants who were given the prediction test are included. We've included all our data because this gives the best estimate of the pattern of performance on the main task, and because it gives the best estimate of the ability of participants to pick the contingent colors during the structured interview. Another objection that might be raised is that the prediction task itself is flawed, in that it a) takes place after the main task by which time participants have forgotten which are the contingent colors and b) is not sufficiently like the main task to help them recall this information. Our response is that our participants appear to have learned to respond faster to Consistent trials relatively early in the experiment, certainly by the time they reached the 6[th] block, and appear to have had no trouble remembering this information (if that was how they were doing it) from block to block from then on. And our prediction task used the same room, same computer, same display, same stimuli presented in a very similar fashion, it was simply the response requirements and consequent presentation rate (and feedback) that were different to the main task. If our participants consciously knew which were the contingent stimuli, this should have acted as a salient reminder of that information.

Perhaps more troublesome for our analysis is the objection that, whilst our prediction test does attempt to assess the correct information (Shanks & St. John, 1994), it

does not do so sensitively enough. We only gave two blocks of 16 trials each, compared to the 20 blocks of 120 trials each for the main task. Is this really a fair test of our participants' knowledge about the color-response contingencies? To which question our answer has to be "yes". The point here is that, in order to be effective in determining RT to respond to the stimulus location, participants' knowledge must be of the categorical variety if the idea of "conscious knowledge" is to mean anything different to "there's been some change in behavior". They must know that this color predicts this location / response, even allowing for a caveat to the effect that this prediction will not always be valid. In other words, they must know, when faced with the appropriate stimulus input, that some colors are predictive and that others do not seem to be. Only then can conscious knowledge produce the desired pattern of results. If this were the case, then our 8 trials for the contingent colors should be more than adequate to detect this type of categorical knowledge. If they get all 8 correct, then this is significantly better than chance on an individual basis. Even allowing for only some of the participants being aware of the contingencies, the means over all 16 participants who took the test should produce overwhelming evidence of such knowledge if only as many as 4 were aware of the contingencies. But we do not even have means that are greater than chance to report for our prediction test, let alone means approaching significance.

Another possibility is that the effect observed in the RTs might be due to our participants being aware of the contingencies on very few trials during the experiment (and then they forget), and that we have simply not given enough trials in the prediction task to capture any of these trials. This is implausible. Firstly, once the contingencies are explicitly coded – then this should in itself increase the probability of noticing them again until it becomes a frequent occurrence. Secondly, even a rate as low as say 1 trial in every 8 of the predictive colors (i.e. 2.5 on average a block in training) would produce a mean on our prediction test of 4.5 that is significantly different to that observed (Z=2.27, p<.05). And the effect on these individual trial RTs would have to be large (approx 40msec). There is no evidence for such a bimodal distribution in our data.

It would seem, then, that we have good evidence for implicit performance on our task. It is still possible that we may have not demonstrated implicit learning, if we adopt a position that results from a stringent application of the Shanks and St. John criteria (our thanks to Tony Dickinson for pointing this out). Imagine that during training on some trials our participants became aware of the contingencies, but then forgot this prior to being asked. They would have learned the contingencies, but not implicitly. This learning might then give rise to performance that reflected the learning, but still without conscious knowledge of the contingencies so that they passed our prediction test. This type of implicit performance, based on a form of implicit memory, would explain our results. In the limit, it is almost impossible to discount this type of critique, but we think

that it lacks plausibility when applied to our paradigm. There is not much to learn, just that two colors predict responses, and if participants were to become repeatedly aware of this (so that learning could occur), it's hard to see why they would forget this so easily.

Our conclusion, then, is that we have demonstrated learning in human participants without conscious awareness of what has been learned. This robust effect has been accomplished in a single, one-hour session with relatively few participants. These considerations suggest that this paradigm is ideally suited to the study of implicit processes, and we hope to report the results of further investigations with this paradigm in the near future.

REFERENCES

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

Cleeremans, A. (1993). *Mechanisms of Implicit Learning: Connectionist Models of Sequence Processing.* Cambridge, MA: MIT Press.

Jones, F.W., & McLaren, I.P.L. (2009). Human sequence learning under incidental and intentional conditions. *Journal of Experimental Psychology: Animal Behavior Processes, 35(4),* 538-553.

Lovibond P.F., & Shanks, D.R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes, 28(1),* 3-26.

McLaren, I.P.L., Green, R.E.A., & Mackintosh, N.J. (1994). Animal learning and the implicit/explicit distinction. In N.C. Ellis (Ed.) *Implicit and explicit learning of languages* (pp. 313-332). New York, NY: Academic Press.

Mitchell, C.J., De Houwer, J., & Lovibond, P.F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences, 32,* 183-246.

Posner, M.I. & Synder, C.R.R. (1975). Attention and Cognitive Control. In Robert L. Solso (ed.), *Information Processing and Cognition: The Loyola Symposium.* Lawrence Erlbaum.

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

Shanks, D.R. & Lovibond, P.F. (2002) Autonomic and eyeblink conditioning are closely related to contingency awareness: Reply to Wiens and Öhman (2002) and Manns et al. (2002). *Journal of Experimental Psychology: Animal Behavior Processes, 28,* 38-42.

Verbruggen, F., Adams, R., & Chambers, C.D. (in press). Proactive motor control reduces monetary gambling. *Psychological Science.*

Willingham D.B., Nissen M.J. and Bullemer P. (2009). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 1047-1060.