Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

______________________________ _______________
Thomas C. Hassett Date

Change in the relative contributions of habit and working memory contributes to expertise in serial reversal learning

By

Thomas C. Hassett
Master of Arts

Psychology

Robert R. Hampton
Advisor

Jocelyne Bachevalier
Committee Member

Patricia J. Bauer
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Change in the relative contributions of habit and working memory contributes to expertise in serial reversal learning

By

Thomas C. Hassett

BS, University of Wisconsin-River Falls, 2013

Advisor: Robert R. Hampton, PhD

An abstract of
A thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Arts
in Psychology
2015

Abstract

Change in the relative contributions of habit and working memory contributes to expertise in serial reversal learning

By Thomas C. Hassett

Multiple memory systems likely evolved to account for learning across functionally distinct environmental demands. Despite their independence, multiple memory systems simultaneously aid in learning so long as each system is able to account for learning on its own. Multiple memory systems may aid in the learning-to-learn phenomenon, and more specifically, in serial reversal learning acquisition. In serial reversal learning, subjects differentially respond to the same two stimuli across all trials of the task. At any given time, only one of the two stimuli is positively reinforced when selected. Once a preference for the positive stimulus is developed, the reinforcement properties of the two stimuli reverse (i.e., S+ to S- and S- to S+). Interestingly, performance on the task improves as subjects gain experience with the task demands. Naïve subjects reverse gradually while experts exhibit rapid reversals through win-stay, lose-shift responding. In the current study, we assessed whether the development of serial reversal learning expertise is facilitated by a shift in memory control, from habit to working memory. In Experiment 1, we alternated the duration of the inter-trial-interval between 1 and 30-seconds to dissociate the relative contributions of habit and working memory on responding. Our results suggest that responding in naïve and expert reversers is under habit and working memory control, respectively. In Experiment 2, we determined that working memory control facilitates the transfer of expertise to a novel image set. In Experiment 3, we used two alternating concurrent cognitive demands to disrupt working memory during win-stay, lose-shift responding. Together our results provide converging evidence that working memory control is integral to the development of expertise in serial reversal learning.

Change in the relative contributions of habit and working memory contributes to expertise in serial reversal learning

By

Thomas C. Hassett

BS, University of Wisconsin-River Falls, 2013

Advisor: Robert R. Hampton, PhD

A thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Arts
in Psychology
2015

Table of Contents

Figures and Tables

1. Introduction

The diversity of learning in animals is thought to be facilitated by multiple, functionally distinct memory systems, which evolved to combat an assortment of pressures encountered across developmental and environmental conditions (Sherry & Schacter, 1987). A cortico-striatal habit memory system appears early in development and facilitates the gradual learning of habits and skills which are robust and inaccessible to conscious monitoring (Bachevalier, 1990; Hay & Jacoby, 1996; Mishkin, Malamut, & Bachevalier, 1984; Squire, 1993). In contrast, working memory facilitates the rapid acquisition of a limited amount of information through active, effortful maintenance (Baddely, 1992; Shettleworth, 2010, chapter 6; Stern, et al., 2000; Squire, 1993). Definitions of working memory often differ between human and nonhuman researchers, however, both definitions include active cognitive control of information (Basile & Hampton, 2013).

Despite their independence, multiple memory systems simultaneously aid in learning of a particular task demand so long as each system is able to provide its own solution (Poldrack & Packard, 2003). In such cases, behavior is determined by the system that exerts the greatest relative strength. For example, the extent to which habit and one-trial memory control behavior has can be experimental adjusted by either decreasing the likelihood that a stimulus predicts reward or increasing the retention interval, respectively (Tu & Hampton, 2012). The relative contributions of two memory systems on behavior may also shift during the process of learning. Control rats tested on a plus maze will initially respond under the control of a hippocampal system, but with additional training, they respond under the control of a caudate-based habit system (Packard and McGaugh, 1996). The temporal shift in memory control may also be true during the formation of learning sets. In Harlow's seminal *Formation of Learning Sets* (1949), he used the term *learning-to-learn* to describe a shift in responding from gradual to rapid learning as subjects were exposed to more discrimination learning sets. Similar, multiple memory system learning may be true for similar learning-to-learn tasks, such as the serial reversal learning task.

In the serial reversal learning task, subjects are repeatedly presented with discrimination trials containing the same two stimuli. At any given time, only one of the two stimuli is rewarded when selected. Within every reversal, the positive stimulus (S+) will be rewarded if selected, and will remain positive until a performance criterion is met. Upon reaching criterion, the contingencies of reinforcement reverse (S+ to S- and S- to S+). Subjects are then required to meet criterion by repeatedly selecting the formerly non-reinforced stimulus. This process occurs over many reversals.

Performance on the serial reversal learning task improves with reversal experience. Naïve reversers exhibit a sub-optimal, gradual reversal response pattern, characterized by perseverative selection of the previously rewarded response following a reversal (Shettleworth, 2010). Expert reversers, on the other hand, show flexible, win-stay, lose-shift responding, where very few errors occur following a reversal (Shettleworth, 2010). The appearance of win-stay, lose-shift occurs in the absence of any change in task demands, suggesting that expertise is facilitated by underlying mechanisms associated with learning.

We hypothesize that the development of serial reversal learning expertise is facilitated by a shift in the relative control of behavior by two independent memory systems, habit and working memory. We hypothesize that the behavior of naïve reversers is under greater relative control of a habit system, while the behavior of expert reversers is under greater relative control of a working memory system. Inflexible, habit controlled responding would explain the perseverative, gradual reversing observed early in a serial reversal learning task. Additionally, cognitively controlled responding would explain the flexible, rapid reversing observed late in the serial reversal learning task.

2. Experiment 1

We tested whether the development of serial reversal learning expertise is facilitated by a shift in the relative control of behavior by habit and working memory systems. Habit and working memory systems are differentially affected by time intervals, with the latter showing susceptibility to long intervals, ostensibly due to a decay in the memory trace (Grant & Roberts, 1973). Using this difference in time-interval sensitivity, we measured accuracy on discrimination trials following short and long inter-trial-intervals (ITIs) in order to determine the extent to which habit and working memory control choice. For purposes of the current study, the ITI can be reinterpreted as the amount of time separating when information was last given about the current positive stimulus (i.e., the outcome of the previous trial), and the current discrimination trial. We compared performance on discrimination trials following 1-second (short) and 30-second (long) ITI types to determine whether the relative contributions of habit and working memory systems differentially influenced responding during the early and late phases of a serial reversal learning task. If naïve reversers respond under habit control, there should be no difference in discrimination performance following 1 or 30-second ITIs. Furthermore, if serial reversal expertise is under working memory control, then discrimination performance should be significantly better following 1-second ITIs than 30-second ITIs.

2.1 Methods

2.2 Subjects and apparatus

Six adult, male rhesus monkeys (*Macaca mulatta*; mean age = 9.16 years) were used. All subjects were pair-housed when not testing, received daily full food rations, and ad libitum access to water. Testing occurred for up to seven hours a day, six days a week. Subjects were tested in their home cages using portable testing rigs. Each testing rig was equipped with 15-inch color LCD touch-sensitive screen (Elo TouchSystems, Menlo Park, CA), running at a resolution of 1024 X 768 pixels. Testing rigs were equipped with two automatic food dispensers (Med Associates, Inc., St. Albans, VT) which delivered nutritionally balanced primate pellets (Bio-Serv, Frenchtown, NJ). Tests were controlled by a personal computer running a custom program written in Presentation (Neurobehavioral Systems, Albany, CA). All six subjects had previous experience with touchscreen tasks, including image discrimination, however, none of the six had previous experience with any type of learning-to-learn task.

2.3 Procedure

Subjects initiated all trials by touching a 100 x 100 pixel green square (FR 2) on a black background (Figure 1). Following initiation, subjects saw two images (350 x 350 pixels), each placed 250 pixels left or right of the central point on the touch screen. The same two picture stimuli were used throughout all reversals. Stimuli were placed on either the left or right side of the screen using a pseudorandom assignment schedule. At most, a stimulus could be placed on the same side four times in a row. Subjects needed to select one of the two stimuli (FR 2) before receiving either a 1 or 30-second ITI. The presentation of 1 and 30-second ITIs followed an alternating schedule, where every trial separated by a 1-second ITI

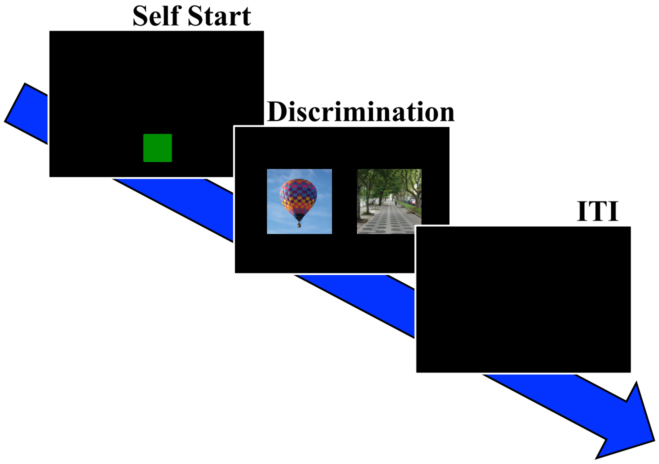

Figure 1: The order of events for all trials in Experiment 1: self-start, discrimination, and ITI. The ITI alternated every trial, between 1 and 30-seconds.

was always followed by a 30-second ITI, and vice-versa, regardless of discrimination trial outcome.

Monkeys completed discrimination trials under a given S+/S- arrangement until they reached a performance criterion of 15 out of 16 correct discrimination trials (or 93.75%). A subject's performance was calculated on every 16th trial. If a subject reached criterion, a reversal occurred, but if criterion was not met, 16 trials additional trials were given using the same S+/S- arrangement. The first trial following a reversal, which will be referred to as *trial 0*, was not included in the performance criterion. Thus, every reversal contained at least 17 trials: trial 0, followed by blocks of 16 trials, 15 of which had to be correct to trigger a reversal.

Responses following 1 and 30 second ITIs were averaged separately to determine the relative contributions of habit and working memory systems to responding. Proportion correct scores were collected and arcsine transformed before analysis (Aron & Aron, 1999). We hypothesized that if naïve reversers responded under habit control, discrimination performance would not differ following 1 or 30-second ITIs. Furthermore, we hypothesized that if expert reversers responded under working memory control, they would perform significantly better on discrimination trials following 1-second ITIs compared to 30-second ITIs.

2.4 Results and Discussion

Five subjects required an average of 5.2 testing sessions and 120.8 reversals to complete the serial reversal learning task. The total number of reversals performed ranged between 84 and 166. On the 4th testing session, 4 of the 6 subjects received an incorrect variation of the original serial reversal learning program which presented 1-second ITIs only. The four subjects performed one session's worth of reversals under this alternate condition, which averaged to 56.5 reversals completed. Three of the four subjects completed at least 35 reversals before receiving the incorrect version of the program. The fourth subject completed fewer than 10 reversals before receiving the incorrect variation, and was dropped from Experiment 1. Following the incorrect program administration, all 5 subjects, including the two that did not receive the incorrect variation, were given an additional 30 reversals of the correct, alternating ITI program.

We compared the first and last 10 reversals performed by each subject to determine whether discrimination performance had significantly improved across the intervening reversals. Performance scores from trials 1-16 were taken from each of the 20 reversals. Only trials 1-16 were used because every subject was guaranteed to have completed at least 16 trials within every reversal. Monkeys made significantly fewer errors to criterion in the last 10 reversals than in the first 10 reversals, suggesting that subjects had developed serial reversal expertise (first 10: M=13.54, SD=7.278; last 10: M=4.66, SD=3.274).

To determine whether the control of habit and working memory on behavior changed as a function of reversal experience, we compared the proportion of correct discrimination trials following 1 and 30-second ITIs during the first and last 10 reversals. Performance following 1 and 30-second ITIs did not differ during the first 10 reversals, but a significant difference in performance following 1 and 30-second ITIs was observed during the last 10 reversals (Figure 2, two-factor repeated measures ANOVA; reversal experience: $F_{(1, 4)}=39.3$, $p < 0.005$; ITI type: $F_{(1,4)}=101.7$, $p < 0.005$;

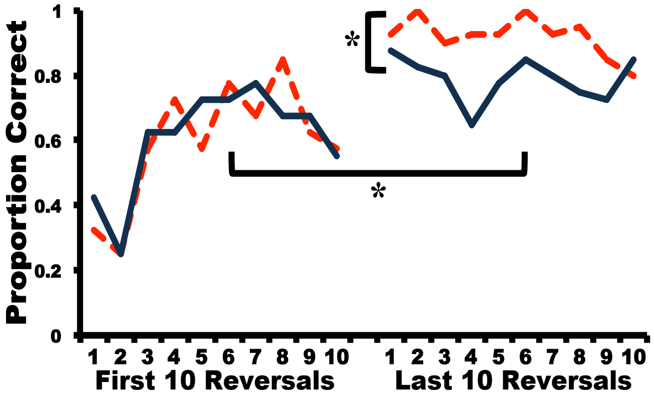

Figure 2: Proportion correct scores on discrimination trials following 1 (red dashed line) and 30-second (solid blue line) ITI lengths for the first and last 10 reversals each subject performed. Asterisks designate significant main effects of ITI type and reversal experience. There was also a significant interaction between ITI and experience.

interaction: $F_{(1,4)}=73.5$, $p<0.005$). Follow up analyses confirmed that performance was significantly better following 1-second ITI, compared to 30-second ITIs in the last 10 reversals, and this difference did not occur during the first 10 reversals (paired samples t-tests; first 10 reversals: $t_4=-.801$, $p > 0.05$; last 10 reversals: $t_4= 11.706$, $p<0.001$).

As naïve reversers develop expertise, responding shifts from a habit-like gradual reversal response pattern, to a flexible win-stay, lose-shift response pattern. These two patterns manifest differently when looking at accuracy during the first 4 trials of a reversal. Performance significantly improves between trials 0 and 1 during win-stay, lose-shift, and only moderately improves during gradual reversal. If behavior is under habit and cognitive control during gradual reversal and win-stay, lose-shift, respectively, then the two response patterns should be differentially affected by 1 and 30-second ITI types. Specifically, naïve reversers should show no difference in performance for trials 0 through 4 following either ITI duration. In contrast, expert reversers should perform significantly better following 1-second ITIs, compared to 30-second ITIs, on trials 1 through 4.

We determined whether naïve and expert responding was under habit and cognitive control, respectively, by averaging discrimination accuracy on trials 0 through 4 across the first and last 10 reversals each subject performed. Accuracy within the first 10 reversals did not differ following 1 and 30-second ITIs, while performance was significantly better following 1-second ITIs compared to 30-second ITIs within the last 10 reversals (Figure 3, three-factor

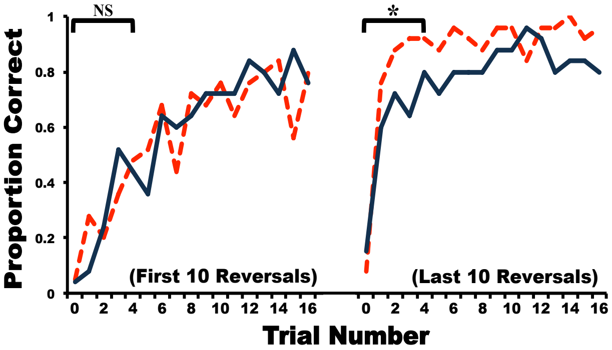

Figure 3: Proportion correct scores on discrimination trials following 1 (red dashed line) and 30-second (blue solid line) ITI lengths.

repeated measures ANOVA; ITI type: $F_{(1,4)}=18.147$, $p < 0.05$; trial number: $F_{(4,16)}=27.715$, $p < 0.005$; reversal experience: $F_{(1,4)}=171.222$, $p < 0.005$; trial number * reversal experience: $F_{(4,16)}=3.843$, $p < 0.05$; ITI type * reversal experience & ITI type * trial number * reversal experience: ns).

Our findings suggest that responding in naïve and expert reversers is under greater relative control of habit and working memory systems, respectively. Naïve and expert reversers were differentially affected by ITI duration, and 30-second ITIs affected the expression of win-stay, lose-shift responding. The deleterious effects of ITI duration on serial reversal learning has been previously reported in pigeons, where performance was significantly worse when reversals contained only long ITIs than when they contained only short ITIs (Ploog & Williams, 2010; Williams, 1976). Similar ITI-sensitive performance has been reported in rhesus monkeys performing on Object Discrimination Learning-Set task (Deets, Harlow, & Blomquist, 1970). However, an alternative explanation for our results may be that performance following 1-second ITIs is significantly better because 3 subjects had substantially more experience discriminating the image pair following 1-second ITIs. If our findings can be explained in terms of having disproportionately more experience discriminating a specific image set under 1-second ITI conditions, then we should not expect serial reversal expertise to transfer to a new image set.

Experiment 2

In Experiment 2, we tested subjects on the same serial reversal learning task, using two new image stimuli. The substitution of two new images served two purposes. First, to determine whether the significant difference in reversal performance following 1 and 30-second ITIs was due to 3 subjects having substantially more discrimination experience with 1-second ITIs. Second, to test whether serial reversal expertise is transfers to a new pair of images. The successful transfer of serial reversal learning expertise to two novel images would suggest that expertise is dependent on working memory control, instead of increased familiarity with the image pair. Conversely, non-transfer of expertise would suggest that serial reversal learning improvement was developed by having substantially more exposure to trials separated by a 1-second ITI.

3.1 Methods

3.2 Subjects and Apparatus

All 6 subjects used in Experiment 1 were also used in Experiment 2. In addition, all testing apparatuses were identical to those described in Experiment 1.

3.3 Procedure

Testing procedures used in Experiment 2 were identical to those described in Experiment 1. The two picture stimuli used in Experiment 1 were replaced with two new 350 x 350 pixel image stimuli.

3.4 Results and Discussion

All 6 subjects completed 60 reversals in an average of 3.5 testing sessions. Monkeys transferred serial reversal learning expertise to new images, showing superior performance following 1-second ITIs from the first block of 10 reversals (Figure 4, two-factor repeated measures ANOVA; ITI type: $F_{(1,5)}=15.577$, $p < 0.05$; reversal experience: $F_{(1,5)}=3.531$, $p > 0.05$; interaction: $F_{(1,5)}=.513$, $p > 0.05$). Our results suggest that

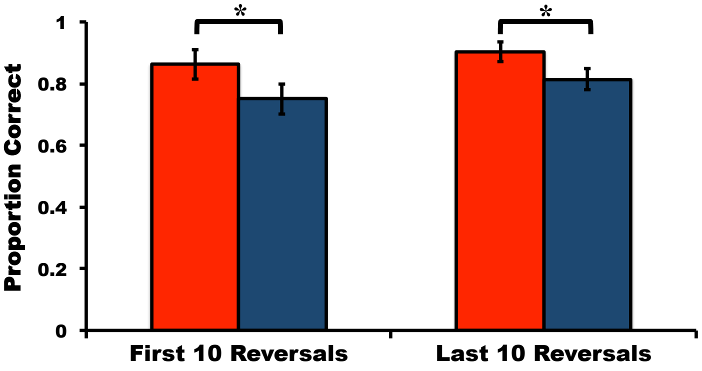

Figure 4: Proportion accuracy scores following 1 (red) and 30-second (blue) ITI lengths. Monkeys exhibited working memory across reversals 1-10 and 51-60.

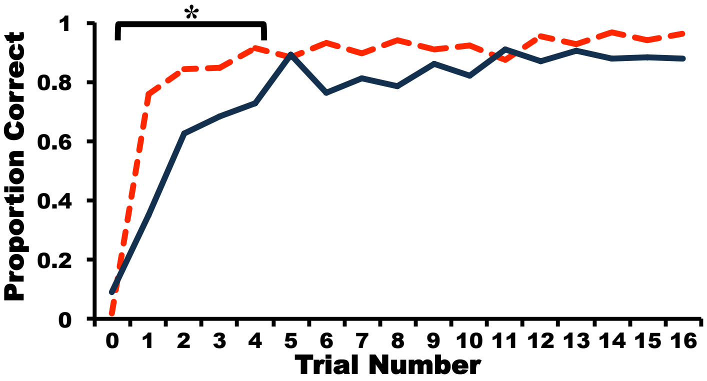

Figure 5: Discrimination accuracy following 1 (red dashed line) and 30-second (solid blue line) ITI lengths on trials 0-16, across all 60 reversals.

working memory control is transferable across stimulus sets, and that the development of expertise in Experiment 1 was not due to increased experience with trials separated by 1-second ITI lengths.

Monkeys exhibited a win-stay, lose-shift across all 60 reversals of Experiment 2 (Figure 5, two-factor repeated measures ANOVA; ITI type: $F_{(1, 5)}=8.807$, $p < 0.05$; trial number: $F_{(4,20)}=59.561$, $p < 0.005$; interaction: $F_{(4,20)}=5.809$, $p < 0.005$). As in Experiment 1, trials 0 through 4 were analyzed because they show the most critical distinction between win-stay, lose-shift and gradual reversal. Since there was no main effect of experience between the first and last 10 reversals, we averaged accuracy on trials 0 through 4 across all 60 reversals.

The findings from Experiment 2 support our hypothesis that win-stay, lose-shift responding is under working memory control. Furthermore, working memory controlled responding was transferred to two novel images. Additionally, our findings indicate that the development of expertise in Experiment 1 was due to a shift in the contributions of habit and working memory, instead of increased experience with 1-second ITIs.

An alternative way to disrupt working memory is by introducing a cognitively demanding task when subjects are actively maintaining information in working memory. In rhesus monkeys, the active maintenance of a highly familiar image is disrupted when subjects are required to perform a cognitively demanding task during the retention interval of a matching-to-sample task (Basile & Hampton, 2013). If monkeys are actively maintaining the outcome of the previous trial in working memory during the serial reversal learning task, then performance should be attenuated if a cognitively demanding task is introduced between discrimination trials. Using a concurrent cognitive task to disrupt serial reversal expertise would provide converging evidence that win-stay, lose-shift is under working memory control.

4. Experiment 3

We assessed the role of working memory in serial reversal learning expertise by alternating two concurrent cognitive load levels across all trials of a serial reversal learning task. The ability of rhesus monkeys to actively maintain a highly familiar image in memory can be attenuated by requiring them to concurrently perform a cognitively demanding task. Specifically, performance on small-set memory tests is significantly worse if the retention interval contains a categorization task compared to an empty time interval (Basile and Hampton, 2013). We compared performance on the serial reversal learning task when discrimination trials were preceded a classification task (e.g., high concurrent cognitive load) or an empty time interval (e.g., no concurrent cognitive load). We hypothesized that if expertise is facilitated by working memory control, monkeys will perform significantly better in the absence of a concurrent cognitive demand than when the concurrent cognitive demands are high.

4.1 Methods

4.2 Subjects and apparatus

Experiment 3 used the same testing equipment described in Experiment 1. The same six subjects used in Experiment 2 were also used in Experiment 3.

4.3 Classification Training

All subjects used in Experiment 3 had previous experience with classifying images as birds, fish, flowers, and people (Basile & Hampton, 2012; Basile & Hampton, 2013; Gazes, Brown, Basile, & Hampton, 2012). However, all subjects were retrained on the classification task before classification and reversal tasks were combined. The stimulus set contained 425 unique images from each of the four categories, resulting in a total of 1,700 images. All images were collected from the online photograph repository Flickr (Yahoo!, Sunnyvale, CA). The entire stimulus set was screened for duplicates using DupDetector (Prismatic Software, Anaheim, CA) and visual inspection. The stimulus set was manually screened to ensure that no image contained more than one category (Gazes, et al., 2012).

Monkeys initiated all training trials by touching (FR2) a green start square (Figure 6). Following trial initiation, monkeys received a centrally located 400 x 300 pixel image corresponding to one of the four categories. Subjects needed to touch the sample image (FR2) before progressing to test. Four 100 x 100 pixel classification icons, each corresponding to one of the four image categories, appeared in the same four corners of the touchscreen. Incorrect classifications resulted in a correction trial containing the same to-be-classified image. Incorrect correction trials were followed by a second correction trial. Second correction trials included the same image, however, only the correct category icon was presented on the screen. This ensured a correct response would occur. All correct classification and

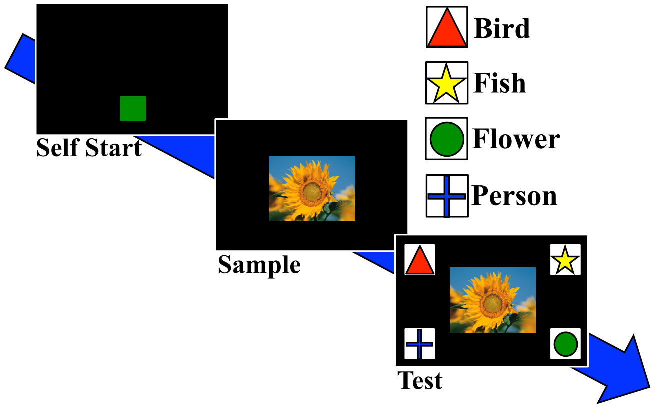

Figure 6: The order of events across all trials of classification training.

correction trials were paired with positive auditory feedback and food reinforcement. All incorrect classification and correction trials were paired with negative auditory feedback and 5-second time-out interval.

All subjects received at least two classification sessions consisting of 600 trials. Stimuli from each of the four classification groups were presented pseudo-randomly, and each group was represented equally within each session. Correction trials did not contribute to the maximum number of trials, thus, every subject viewed 150 images from each category within a session. Monkeys trained until they completed two consecutive classification sessions at 80% or better.

4.4 Procedure

Subjects performed a serial reversal learning task which alternated two concurrent cognitive load types across all trials. All trials followed the same order: self-start, concurrent cognitive load, discrimination, and ITI. All trials were identical with the exception of the concurrent cognitive demand event, which alternated between two conditions: a classification task and a empty time interval yoked to the amount of time it took to complete the classification on the previous trial (Figure 7).

Images from each category were pseudo-randomly presented so that each category was represented twice every 16 trial. Since concurrent cognitive loads alternated every trial, only 8 category presentations occurred within 16 trials. Subjects viewed a centrally located 400 x 300 pixel image, with the four category icons in each corner. Correct classifications were paired with positive auditory feedback and allowed subjects to progress to the discrimination trial. Incorrect classifications were paired with negative auditory feedback and resulted in the immediate presentation of a different to-be-classified image. This same process occurred until an image was correctly classified. On the following trial, a yoked empty interval replaced the

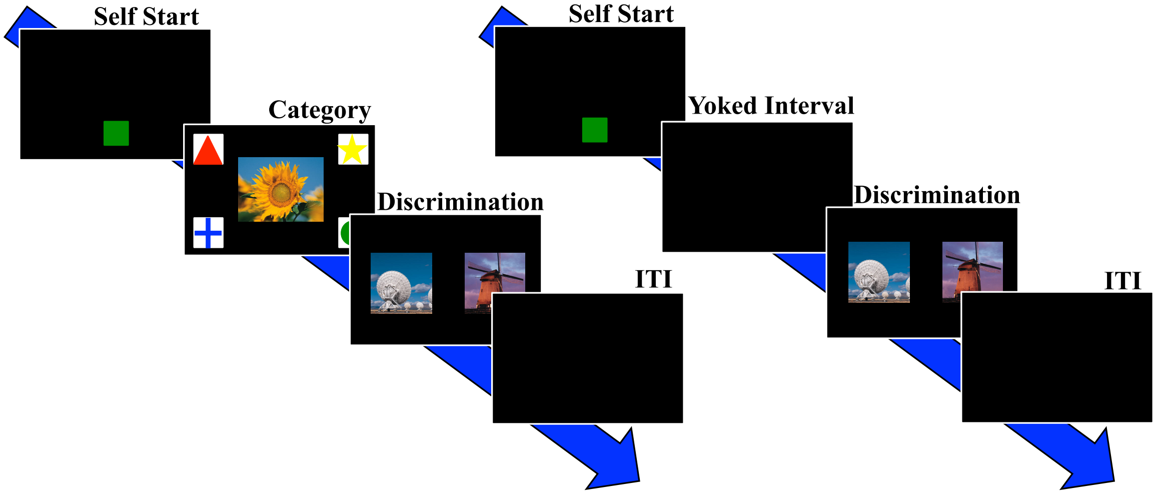

Figure 7: Order of events for the two trial types in Experiment 3. All trials had the same basic structure of: self-start, concurrent cognitive load, discrimination, and a 1000ms ITI. The concurrent cognitive load type alternated every trial so that a high concurrent cognitive load (i.e. category task) trail was always followed by a low concurrent cognitive load (i.e. yoked interval) trial, and vice versa.

classification task and served to control for time across both concurrent cognitive loads. During the yoked empty interval, subjects viewed a black screen which was yoked to the duration of time it took to complete all correct and incorrect classification attempts from the previous trial.

Following the concurrent cognitive load presentation, subjects received a discrimination trial. Two new discrimination images were used for Experiment 3. Following all discrimination trials, regardless of whether the discrimination was correct or incorrect, subjects received a 1000ms ITI.

4.5 Results and Discussion

Monkeys required significantly more classification corrections during the first 10 reversals compared to the last 10 reversals, suggesting that a practice effect occur when combining classification and reversal procedures together (paired samples t-test: t_5=2.670, $p < 0.05$). Because empty intervals were yoked to the all of the classification attempts from the previous trial, the practice effect may have increased the difficulty of actively maintaining information across the entirety of the empty interval. Due to this finding we used two methods for analyzing discrimination accuracy on trials 1-16 during the first and last 10 reversals every subject performed.

First, we analyzed all trials, regardless of whether a classification trial was needed. Using this method, monkeys were significantly more accurate when concurrent cognitive conditions were low, compared to high, and also, they performed better overall during the last 10 reversals (two-factor repeated measures ANOVA; cognitive load: $F_{(1,5)}=21.464$, $p < 0.01$; reversal experience: $F_{(1,5)}=6.606$, $p = 0.05$; interaction: $F_{(1,5)}=3.642$, $p > 0.05$). In the second method, we eliminated all category correction trials and the subsequent yoked interval trials from our analysis. Using this method, monkeys performed significantly better under low concurrent cognitive differences, and did not show an effect of experience (two-factor repeated measures ANOVA; cognitive load: $F_{(1,5)}=7.266$, $p < 0.05$; reversal experience: $F_{(1,5)}=1.183$, $p > 0.05$; interaction: $F_{(1,5)}=2.765$, $p > 0.05$). However, a problem with our second approach is that many trials from the first 10 reversals required classification corrects, and thus, were eliminated from our analysis.

Since both ways of analyzing the first and last 10 reversals were affected by a practice effect, we decided to look at performance following concurrent cognitive conditions for the last 10 reversals only. For this analysis, we looked trials 1-16, regardless of whether a

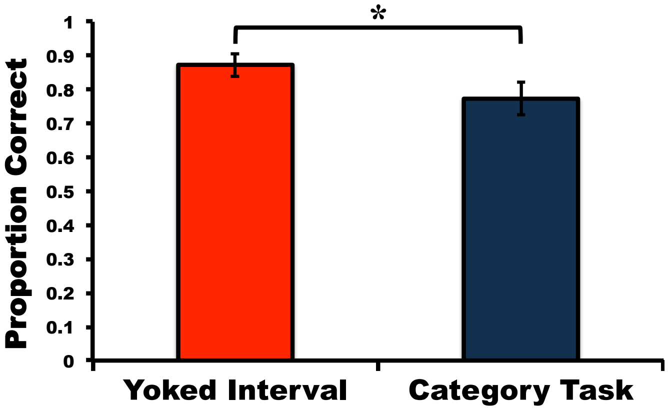

Figure 8: Discrimination accuracy was significantly better following when concurrent cognitive demands were low (red) compared to high (blue).

category correction was needed. Monkeys performed significantly better when discrimination trials followed an empty interval than a classification task (Figure 8, paired samples t-test; $t_5=14.055$, $p < 0.001$).

Trial analysis provided further support that monkeys were responding under working memory control (Figure 9, repeated measures ANOVA; concurrent cognitive load: $F_{(1,5)}=27.603$, $p < 0.01$; trial number: $F_{(4,20)}=16.916$, $p < 0.005$; interaction: $F_{(4,20)}=4.744$, $p < 0.01$). Our trial-by-trial analysis assessed performance on trials 0-4, for the last 10 reversals only. All trials, regardless of whether or not category was correct on the first attempt were included in this analysis. Again, we only analyzed the last 10 reversals because of the reported practice effect.

Experiment 3 varied the difficulty of discrimination trials by alternating the concurrent cognitive load. Our results showed that subjects performed significantly better when the concurrent cognitive load was low (i.e., yoked empty time interval) compared to when it was high (i.e., categorization). This effect was demonstrated in a time-independent manner, suggesting that optimal performance on the serial reversal learning task is facilitated by active maintenance of the outcome of the previous trial, rather than interference effects caused by PI. The results from Experiment 3 provide converging evidence that serial reversal learning expertise is facilitated by working memory control.

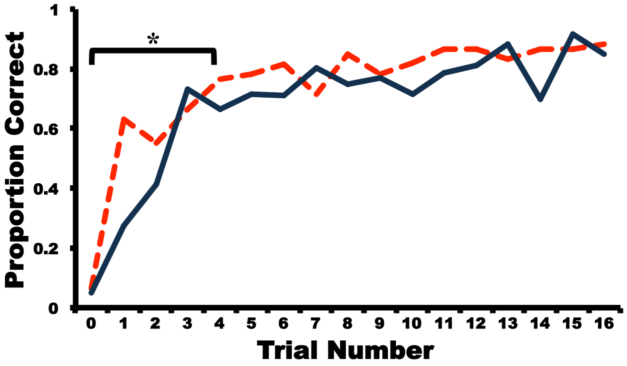

Figure 9: Discrimination accuracy following a yoked empty interval (red dashed line) and discrimination task (solid blue line).

5. General Discussion

Our findings suggest that the development of serial reversal expertise is facilitated by a shift in the relative contributions of habit and working memory on responding. Our data further suggest that expertise is robust and transfers across different experimental procedures. However, while our findings stress the importance of working memory in the manifestation of expertise, reference memory may also be involved. Expertise likely is also dependent on learning task rules, such as, that a subject will never be reinforced for selecting a stimulus that was not rewarded on the last trial. If this is true, then reference memory likely guides working memory towards the goal of remembering the outcome of the last trial. Thus, while working memory likely does not act alone in the expression of expertise, we suggest that it is critical for win-stay, lose-shift responding both across and within reversals.

It should be noted that a memory systems approach is not the first account for serial reversal learning improvement. Beginning in the late 1960s, researchers began to emphasize the importance of either proactive interference or response-outcome cues in serial reversal learning expertise (Clayton, 1966; Gonzalez, Behrend, & Bitterman, 1967; Williams, 1976). Under a proactive interference account, as subjects develop a history for selecting either of the two stimuli, proactive interference increases the difficulty of remembering the previously rewarded response following a reversal. Performance is improved, then, by decreasing perseverative responding at the beginning of a reversal (Clayton, 1966; Gonzalez et al., 1967). Alternatively, under a response-outcome cue perspective, as subjects experience more reversals, the outcome of the previous trial becomes an increasingly salient source of information for guiding choice on the upcoming trial (Williams, 1976). More simply stated, proactive interference improves performance through forgetting, while response-outcome cues improve performance by remembering.

A few weaknesses have been associated with using proactive interference and response-outcome cues to explain the development of serial reversal expertise. Proactive interference cannot explain performance improvements when reversals are experienced in rapid succession (Mackintosh, McGonigle, Holgate, & Vanderver, 1968). Specifically, when reversals are separated by short time intervals, proactive interference is abated by recency. Our findings illuminate this weakness of the proactive interference account, because all of our monkeys significantly improved performance under conditions of successively experienced reversals.

Proponents of the response-outcome cue perspective, suggests that the outcome of the previous trial becomes more informative as reversal experience is accrued. However, it has yet to be determined whether subjects develop a tendency to attend to the outcome of previous trial, or whether it is important throughout the entire experiment, but only behaviorally expressed after many completed reversals (Williams, 1976). Another weakness of the response-outcome cue perspective is that it does not explain why the outcome of the trial becomes more salient with reversal experience. Our results address this weakness by suggesting that the outcome of the previous trial may become a more salient as responding gradually comes under the control of working memory.

Since the mid 1970s, debate regarding the validity for both proactive interference and response-outcome cues has largely subsided. Most experts from both sides of the argument agree that the role of proactive interference and response-outcome cues in reversal learning expertise appears to be species-dependent (Gonzalez et al., 1967; Williams, 1976). For instance, proactive interference appears to be the primary mechanism for reversal learning improvement in bumblebees and goldfish (Gonzalez et al., 1967; Strang & Sherry, 2013). Our results suggest that response-outcome cues may contribute to expertise in rhesus monkeys once responding comes under cognitive control. Future comparative studies may seek to address, specifically, how proactive interference and memory systems differ across species.

Serial reversal learning improvement may involve more than just a shift from habit to working memory controlled responding. Our findings suggest that a perseveration-reducing mechanism acts to facilitate improvements on the serial reversal learning task. In Experiment 1, performance following both 1 and 30-second ITIs improved equally across the first 10 reversals (see First 10 Reversals in Figure 2). No significant differences between performance following either ITI occurred during the first 10 reversals, suggesting that this improvement was independent of a working memory process. Perhaps, then, the development of expertise in serial reversal learning may involve two processes: first a perseverative reduction process, and second a working memory process.

Broadly, the serial reversal learning task requires subjects to flexibly adjust responding between two alternating positive stimuli. Our results suggest that monkeys use working memory processes in order to flexibly adjust responding following reversals. However, our work still has yet to address how memory systems interact within reversals. Our trial analyses show that when working memory is disrupted early in a reversal, subjects respond based on a habit that is incongruent with the current conditions. However, when working memory is disrupted late within a reversal, subjects are able to correctly respond. This may suggest that a new habit is formed every reversal. One possible explanation is that working memory controls behavior across all trials of a reversal. Alternatively, there may be a shift in memory control within reversals, from cognitively controlled, to habit controlled responding. If subjects are able to develop new habits within reversals, habit may exert greater relative control on behavior as subjects are reliably rewarded for selecting the same positive stimulus. This shift would likely be advantageous from an evolutionary perspective, since animals would benefit from ‘offloading’ behaviors on habit in order to allocate limited attentional resources elsewhere. Future research may benefit from studying the interactions of memory systems within reversals.

6. References

Aron, A., & Aron, E. (1999). Statistics for psychology. Upper Saddle River, NJ: Prentice Hall.

Baddeley, A. (1992). Working memory. *Science*, *255*(5044), 556-559.

Bachevalier, J. (1990). Ontogenetic development of habit and memory formation in primates. *Annals of the New York Academy of Sciences*, *608*(1), 457-484.

Basile, B. M., & Hampton, R. R. (2013). Monkeys show recognition without priming in a classification task. *Behavioural processes*, *93*, 50-61.

Basile, B. M., & Hampton, R. R. (2013). Dissociation of active working memory and passive recognition in rhesus monkeys. *Cognition*, *126*(3), 391-396.

Bessemer, D. W., & Stollnitz, F. (1971). Retention of discriminations and an analysis of learning set. *Behavior of nonhuman primates*, *4*, 1-58.

Clayton, K. N. (1966). T-Maze Acquisition and Reversal as a Function of Intertrial Interval. *Journal of Comparative and Physiological Psychology*, *62*(3), 409.

Deets, A. C., Harlow, H. F., & Blomquist, A. J. (1970). Effects of intertrial interval and Trial 1 reward during acquisition of an object-discrimination learning set in monkeys. *Journal of Comparative and Physiological Psychology*, *73*(3), 501.

Gazes, R. P., Brown, E. K., Basile, B. M., & Hampton, R. R. (2013). Automated cognitive testing of monkeys in social groups yields results comparable to individual laboratory based testing. *Animal cognition*, *16*(3), 445-458.

Gonzalez, R. C., Behrend, E. R., & Bitterman, M. E. (1967). Reversal learning and forgetting in bird and fish. *Science*, *158*(3800), 519-521.

Grant, D. S., & Roberts, W. A. (1973). Trace interaction in pigeon short-term memory. *Journal of Experimental Psychology*, *101*(1), 21.

Harlow, H. F. (1949). The formation of learning sets. *Psychological review*, *56*(1), 51.

Mackintosh, N. J., MCGONIGLE, B., & HOLGATE, V. (1968). Factors underlying improvement in serial reversal learning. *Canadian Journal of Psychology/ Revuecanadienne de psychologie*, *22*(2), 85.

McGregor, P. K., & Avery, M. I. (1986). The unsung songs of great tits *(Parus major):* Learning neighbours' songs for discrimination. *Be- havioral Ecology & Soclobiology, 18,* 311-316.

Ploog, B. O., & Williams, B. A. (2010). Serial discrimination reversal learning in pigeons as a function of intertrial interval and delay of reinforcement. *Learning & Behavior*, *38*(1), 96-102. doi:10.3758/LB.38.1.96

Poldrack, R. A., & Packard, M. G. (2003). Competition among multiple memory systems: converging evidence from animal and human brain studies. *Neuropsychologia*, *41*(3), 245-251.

Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, *94*(4), 439-454. doi:10.1037/0033-295X.94.4.439

Shettleworth, S.J. (2010). *Cognition, Evolution, and Behavior*. New York, NY: Oxford University Press

Strang, C. G., & Sherry, D. F. (2013). Serial reversal learning in bumblebees (bombus impatiens). *Animal Cognition*, doi:10.1007/s10071-013-0704-1

Squire, L. R., Knowlton, B., & Musen, G. (1993). The structure and organization of memory. *Annual review of psychology*, *44*(1), 453-495.

Tu, H.-W., & Hampton, R. R. (2012, October 29). One-Trial Memory and Habit Contribute Independently to Matching-to-Sample Performance in Rhesus Monkeys (Macaca mulatta). *Journal of Comparative Psychology*. Advance online publication. doi:10.1037/a0030496

Williams, B. A. (1976). Short-term retention of response outcome as a determinant of serial reversal learning. *Learning and Motivation*, *7*(3), 418-430.

Yokoyama, C., Tsukada, H., Watanabe, Y., & Onoe, H. (2005). A dynamic shift of neural network activity before and after learning-set formation. *Cerebral Cortex*, *15*(6), 796 801.