

# Characteristics of implicit chaining in cotton-top tamarins (*Saguinus oedipus*)

Charles Locurto · Matthew Gagne · Lauren Nutile

Received: 2 August 2009 / Revised: 16 January 2010 / Accepted: 20 January 2010 / Published online: 7 February 2010  
© Springer-Verlag 2010

**Abstract** In human cognition there has been considerable interest in observing the conditions under which subjects learn material without explicit instructions to learn. In the present experiments, we adapted this issue to nonhumans by asking what subjects learn in the absence of explicit reinforcement for correct responses. Two experiments examined the acquisition of sequence information by cotton-top tamarins (*Saguinus oedipus*) when such learning was not demanded by the experimental contingencies. An implicit chaining procedure was used in which visual stimuli were presented serially on a touchscreen. Subjects were required to touch one stimulus to advance to the next stimulus. Stimulus presentations followed a pattern, but learning the pattern was not necessary for reinforcement. In Experiment 1 the chain consisted of five different visual stimuli that were presented in the same order on each trial. Each stimulus could occur at any one of six touchscreen positions. In Experiment 2 the same visual element was presented serially in the same five locations on each trial, thereby allowing a behavioral pattern to be correlated with the visual pattern. In this experiment two new tests, a Wild-Card test and a Running-Start test, were used to assess what was learned in this procedure. Results from both experiments indicated that tamarins acquired more information from an implicit chain than was required by the contingencies of reinforcement. These results contribute to the developing literature on nonhuman analogs of implicit learning.

**Keywords** Implicit learning · Serial learning · Chaining · Cotton-top tamarins

## Introduction

In the study of human cognition there has been substantial interest in the phenomenon of implicit learning, broadly defined as learning that occurs in the absence of explicit instructions to learn, and without subjects becoming aware of the experimental contingencies. Interest in this form of learning is understandable given that many human abilities, ranging from language to a variety of motor skills, appear to be learned in the absence of explicit instructions (Reber 1996). Implicit learning was first studied by Reber (1967) using a format in which experimental subjects were presented with letter strings that followed a pattern determined by an artificial grammar. Artificial grammars specify an arbitrary set of rules that govern transitions between members of a set of elements. Subjects were asked to memorize the strings, but they were not informed about the presence of a pattern. Subjects nevertheless learned these strings more readily than did control subjects who were presented with letter strings that had been composed randomly with respect to the grammar. Additionally, following training, many experimental subjects were unable to articulate the rules of that grammar, an observation in this and other studies that led to the name “implicit” for this type of learning.<sup>1</sup>

Following Reber’s demonstration, a considerable number of studies explored the nature of implicit learning. Two types of procedures have predominated: one or another form of the original artificial grammar design studied by

<sup>1</sup> In the human cognition literature there are overlapping usages of the terms implicit learning and incidental learning. We suggest that incidental learning refers to all types of learning in human subjects that occur unintentionally, that is, without instructions to learn. Implicit learning is a subset of these incidental learning procedures in which learning occurs typically without subjects’ awareness of the experimental contingencies.

C. Locurto (✉) · M. Gagne · L. Nutile  
Department of Psychology, College of the Holy Cross,  
1 College St., P.O. Box 75A, Worcester, MA 01610, USA  
e-mail: clocurto@holycross.edu

Reber, and a second form of implicit learning, the serial reaction time task (SRT), developed by Nissen and Bullemer (1987). In their procedure, experimental subjects received repeated presentations of a visual stimulus, an asterisk, on a computer screen. The presentations of the asterisk followed a continuously repeating pattern in the form DBCACBDCBA, where each letter referred to one of four spatial locations on the screen where the asterisk might occur. Subjects were required to tap an arbitrarily chosen keyboard key for each stimulus (e.g., if D, press the key for the digit 9 on the keyboard). The string was presented without breaks between the end of one presentation of the string and the beginning of the next presentation, thereby making detection of the beginnings and ends of the string difficult. Results indicated that the latencies of experimental subjects were lower than those of control subjects who experienced the same stimuli presented randomly. Further, as was true of Reber's procedure, experimental subjects often were unable to articulate the nature of the pattern, particularly when the task was made more complex by introducing a second, concurrent task to be performed during pattern training (Nissen and Bullemer, Experiment 2).

Subsequent work has confirmed the basic finding that in the SRT procedure response latencies are lower during presentations of the pattern compared to a random presentation condition, despite the absence of instructions to learn the pattern. Moreover, frequently but not invariably, subjects cannot verbalize the nature of the pattern. (See Clegg et al. 1998, and Seger 1994, for reviews.) These findings are demonstrable using a number of variations of the initial SRT procedure. One variation marries the SRT task with Reber's original artificial grammar procedure, such that training strings are developed in accordance with the rules of an artificial grammar. This variation makes presentation of an individual element probabilistic (i.e., any artificial grammar permits some degrees of freedom in terms of which element may be next in the sequence), not deterministic (i.e., the elements always follow the same pattern), as was true in Nissen and Bullemer's original study. Despite this complication, subjects trained with strings generated from an artificial grammar typically demonstrate lower reaction times during training than do subjects trained with randomly constructed strings (Deroost and Soetens 2006, Experiment 4; Soetens et al. 2004).

The important role that implicit learning has come to play in the study of human cognition makes it valuable to understand its comparative origins. It might be argued that there can be no adequate nonhuman analog of human implicit learning, inasmuch as there can be no assessment of subjects' awareness of a pattern, one of the features of human implicit learning that has conferred so much interest in it. As counterpoint, it may be suggested that an implicit learning procedure for a nonhuman would present

information to subjects, but reinforcement would not depend on learning that information. The question under consideration would then be whether in this format a nonhuman learns more than is demanded by the experimental contingencies (Locurto et al. 2009).

There have been a number of adaptations of the basic SRT procedure for nonhumans. In these adaptations subjects are exposed to a serial pattern that involves a number of stimuli. Subjects are required to respond to each stimulus to advance the pattern to the next stimulus. Reinforcement is presented periodically, or after each response, but it is not contingent on performance that reflects knowledge of the pattern. Subjects can adopt a simple rule in the form "Respond to the stimulus, and periodically food will come" that does not involve pattern information. Following training, a random test is used to observe whether latencies rise compared to their training baseline. These adaptations have yielded evidence of decreases in latencies during pattern training, and increases during a random test (in rats: Christie and Dalrymple-Alford 2004; Domenger and Schwarting 2005; in mice: Christie and Hersch 2004; in rhesus macaques: Procyk et al. 2000). These results, which parallel the basic findings in human SRT studies, indicate that nonhuman subjects learned something about the pattern. Without that learning, there would have been no rise in latencies during the random test.

While confirming that learning of some sort took place, these results do not address the question of precisely what subjects learned in this procedure. To illustrate this problem, consider procedures developed from a parallel literature, the study of explicit serial learning. In these studies, which have included both humans as well as a number of nonhuman species, a variety of training and testing procedures have been used. In one standard procedure, training proceeds in the following manner: If the series includes, as it typically does, five arbitrarily chosen visual elements (A through E), training begins on the first adjacent pair of elements, A and B. The subject must choose between A and B when the pair is presented, with reinforcement contingent upon choosing A. A variant of this procedure requires subjects to choose the two stimuli in the correct order, A → B (e.g., Merritt et al. 2007). This first pair is designated A+/B−, where "+" designates the correct response, or the correct first choice, for each pair. Following mastery of that pair, subsequent pairs follow the pattern, B+/C−, C+/D−, D+/E−. In this procedure, the complete chain of elements is never seen in its entirety. (See Terrace 1993, for a different form of sequence learning in which all elements are presented concurrently.)

After an acquisition criterion is reached for each training pair, testing proceeds using pairs that were never presented together during training, e.g., A/C, B/E, etc., to determine subjects' ability to organize these arbitrary

stimuli into an ordinal sequence. Results from a substantial body of literature indicate that subjects are able to order previously untrained test pairs correctly, thereby indicating that they appreciate the ordinal nature of the chain (Conway and Christiansen 2001, and D'Amato 1991, provide reviews of other forms of explicit sequence learning in nonhumans. For examples of transitive inference, a similar type of serial learning: in rats, Davis 1992; in corvids, Bond et al. 2003 in rhesus macaques, Rapp et al. 1996.)

It would be valuable to adapt this technology to the SRT task to understand what has been learned about stimulus order under conditions in which such learning is not demanded. Unfortunately, there are procedural features in the standard SRT task that render it impossible to determine how individual elements are appreciated. One feature is that the same stimulus location is used repeatedly throughout the pattern. In Nissen and Bullemer's study, for example, particular elements are used twice (locations A and D) or three times (locations B and C) in one complete presentation of the pattern. As a consequence, a particular element was not uniquely associated with a particular ordinal position. In explicit serial learning procedures elements are unique, spatially or visually. Additionally, in SRT procedures the pattern is repeated continuously, without demarcation of a beginning or end point. This feature contributes to the difficulty of assessing to what extent subjects appreciate the sequential nature of the pattern.

To address these limitations, Locurto et al. (2009) developed a procedure that married aspects of the standard SRT task with features taken from explicit serial learning procedures. The task, termed implicit chaining, consisted of a five-element chain in which the same visual element, an orange slice, was displayed at one of five locations on a touchscreen. Location presentations followed a pattern, and subjects, two cotton-top tamarins (*Saguinus oedipus*), were required to touch the element each time it appeared to advance the chain to the next element. Reinforcement was presented following the fifth element, and was followed by a 20 s intertrial interval (ITI). In this procedure, there was a pattern that possessed discernible beginning and end points, and each element in the chain occupied a unique ordinal and spatial position in the chain, but reinforcement was not dependent on pattern information. A subject need only adopt a simple rule, "Touch the orange slice and food comes periodically," to receive reinforcement. Following training on the pattern, pair-wise tests were conducted, as is typical in explicit serial learning procedures, as was a random test, which is characteristic in SRT studies.

The differences between these tests mirror their different origins. The pair-wise tests derive from the serial learning

literature. While in our adaptation of this test there are no correct answers, and reinforcement is delivered independently of a subject's choice, they are, in a sense, explicit tests in that subjects are asked which stimulus of two they prefer. The random test, derived from the implicit learning literature, measures latencies and is more of an embedded test in that latencies are recorded in the course of a subject's completion of the task, but they do not interrupt the task.

Results indicated that the tamarins learned something about the pattern. For each tamarin, latencies were higher during the random test than during training. The results of the pair-wise tests indicated that each tamarin chose later elements in the chain compared to earlier elements (e.g., D over A), indicating that they appreciated the elements of the chain in terms of their proximity to reinforcement. The pattern of choosing later over earlier elements was true for nonadjacent elements, and was stronger for pairs that contained an end element.

While these results indicate that the tamarins learned more than was required by the experimental contingencies, they leave open a number of issues. Addressing those issues forms the rationale for the present experiments. The first issue comes from the recognition that in Locurto et al.'s procedure, the perceptual pattern of presenting the visual icon in different locations was accompanied by a behavioral pattern. The issue arising from this covariation is whether learning was principally perceptual, in the form "The icon occurs here, then there," etc., principally behavioral in nature, "Touch this location, then this one," etc., or some combination of each type of learning. This issue has been of interest in the human SRT literature, and to date there is little consensus as to whether SRT learning is principally behavioral, or includes perceptual components apart from the learning of a behavioral sequence (Eimer et al. 1996; Deroost and Soetens 2006, Experiment 1; Remillard 2003; Rowland and Shanks 2006; Willingham et al. 1989, Experiment 3; see Goschke 1998, for a review). Experiment 1 asked whether implicit learning is evident in a procedure in which the perceptual and behavioral components of an implicit chain were dissociated by creating a purely perceptual pattern. Five different visual elements were presented in the same order on each trial. Each element could occur at any one of six locations on the touchscreen on a given trial. Reinforcement was delivered at the end of each chain, as was the case in Locurto et al.

In Experiment 2, a behavioral pattern was correlated with the perceptual pattern, and two new tests, a Wild-Card test and a Running-Start test, were introduced to assess what was learned in an implicit chaining procedure. Additionally, reinforcement came at the end of chains with a probability of 0.33, rather than a probability of 1.0 that had been programmed in Experiment 1 and in Locurto et al.

## Experiment 1

### Method

#### *Subjects*

Subjects were two male cotton-top tamarins. One subject, Windsor, was 19 years old at the start of training; the second subject, Winston, was 14 years old. Both tamarins were among the 19% of their colony of approximately 150 tamarins that were more than one standard deviation above the average colony age of 7.5 years. Prior to the present study both tamarins served as subjects in Locurto et al., and in a pilot study involving an implicit chaining procedure similar to that reported in Locurto et al. The tamarins had been born and were housed individually at the New England Primate Research Center, Southborough, Massachusetts. Subjects were maintained in accordance with the guidelines of the Committee on Animals of the Harvard Medical School. Their daily feeding regimen consisted of a morning feeding and a second one in the afternoon. On training days they received their morning feeding for approximately 2 h. Food was removed 3 h prior to training, which took place in the early afternoon. Feedings consisted of monkey chow supplemented with fresh fruit. The tamarins received their afternoon feeding immediately following a session. With this schedule, each tamarin was maintained at about 95% of its free-feeding weight. Training was conducted on average four times each week.

Cotton-top tamarins were chosen for this work because they have been used to study a wide range of learning mechanisms. Of relevance to the present work, Hauser and colleagues have used cotton-top tamarins to study forms of implicit learning and related phenomena including statistical learning (e.g., Fitch and Hauser 2004; Hauser et al. 2001; Ramus, et al. 2000). There are presently more than 300 references in the PsycINFO database for tamarins as subjects in psychological research. More than 100 of those references are for cotton-tops.

#### *Apparatus*

The stimuli consisted of five different icons, each approximately  $200 \times 200$  pixels. The icons were as follows: A (sunglasses), B (airplane), C (Leaning Tower of Pisa), D (apple), and E (flower). Each stimulus was displayed against a white background on a touchscreen (ELO Carroll-Touch infrared; Tyco Electronics, Berwyn, PA). The screen was 38.1 cm on the diagonal. The front of the touchscreen was fitted with a clear Plexiglas bezel that allowed access to the touchscreen at six locations ( $10 \text{ cm} \times 10 \text{ cm}$  holes) arranged in a  $2 \times 3$  array across the touchscreen. The images were displayed in the middle of one of the six loca-

tions. Reinforcement consisted of two 45 mg chocolate sucrose pellets (Test Diet, Purina Mills, LLC) delivered into a dish located directly in front of the bottom left side of the touchscreen via a rotary pellet feeder (Coulbourn Instruments, Allentown, PA).

#### *Procedure*

Two changes were made in the procedures for this experiment that shortened the length of acquisition and testing compared to Locurto et al. These changes reflected the fact that each tamarin had, by the start of this experiment, participated in two implicit chains procedure (Locurto et al. and a pilot study). The first change was that acquisition was reduced from a minimum of 15 sessions to 10 sessions. The second change was that pair-wise testing was reduced to 5 sessions from the 12 sessions that were used in Locurto et al., because in that earlier study we observed that pair-wise choices during the first 5 test sessions did not differ from choices during the last 5 sessions. With this testing format, each element was presented on a total of 20 trials.

Subjects were tested in their home cage. This procedure has precedent in other work that has identified advantages in allowing primates to remain in their home cage during training (e.g., Crofts et al. 1999). The touchscreen was placed in front of the subject's home cage by using a load lifter (Genie Industries, Redmond, WA). The front door to the animal's cage was then opened, allowing the subject access to the touchscreen and food dish.

*Acquisition* Since both tamarins were experienced we presented them on the first day of training with a full set of training stimuli. Sessions consisted of 40 trials. Each trial consisted of the presentation of the five-element chain. During element presentations the screen background was white. Each response produced an auditory stimulus approximately 100 ms in duration. The first touch to an element advanced the chain to the next element. Touches to blank areas of the screen were recorded but had no consequences. A 5 s blank white screen preceded the first element in a trial as a ready signal. Reinforcement was delivered at the end of each chain. A 20 s ITI, a dark screen, followed reinforcement.

Training proceeded until each tamarin completed at least 10 40-trial sessions of training, and sessions were completed with  $\pm 10\%$  of the average session time in four the last five sessions. Session time was used as the training criterion, inasmuch as there are no errors in the sense of incorrect choices in this procedure. Hence, criteria based on correct performance were not possible. Given that the ready signal and the ITI were constant across trials, variations in session time consisted of differences in a subject's completion of the chains. Windsor and Winston met this criterion

in ten sessions. Sessions were scheduled for 40 trials, or were terminated after 50 min.

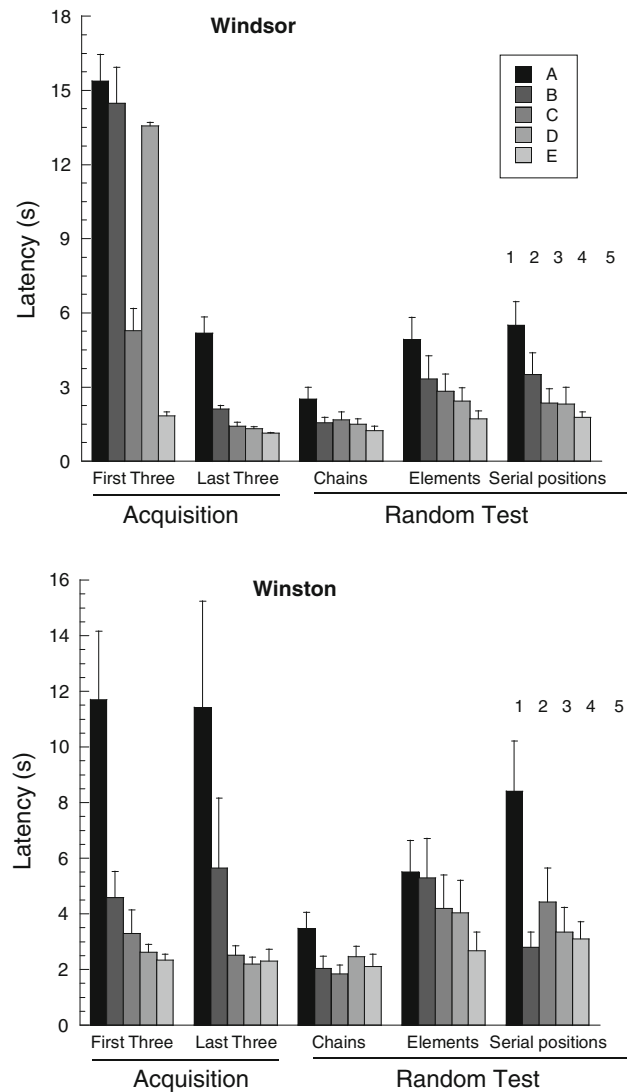
**Testing: pair-wise tests** Following acquisition, pair-wise testing began. Each pair-wise test session contained 40 trials. The first ten trials in each session were composed of the same chains that were presented during acquisition, with reinforcement delivered at the end of each chain as was true during acquisition. The remaining 30 trials consisted of 20 chain trials and ten pair-wise tests. Each of the ten pair-wise test trials contained one of the possible comparisons between elements (e.g., A/B, C/E). These 30 trials were programmed such that at least one chain trial was interspersed between pair-wise tests. During a pair-wise test trial, the two elements were presented simultaneously. The first response to either element ended the trial and produced reinforcement. Reinforcement was scheduled randomly following choices during this test. A subject's particular choice did not determine whether reinforcement would be delivered. There were five pair-wise test sessions.

**Testing: random sequences** The random test session followed the last pair-wise test session. In this session, the first 20 trials were chain trials, identical to those used in acquisition. For the next 20 trials, the elements were presented in randomized order, with reinforcement delivered at the end of each chain, as was done during acquisition. Randomization of elements was accomplished with two constraints: on any trial an element could occur only once; and, each element occurred four times in each serial position during the test.

For all latency analyses, both during acquisition and the random test session, latencies that fell more than two standard deviations from the mean were eliminated from analyses. By this criterion, during the random test session, which consisted of 200 element presentations (40 trials (20 chains and 20 random trials)  $\times$  5 serial positions), 4 latencies were eliminated for Windsor, two each during chain trials and the random test. For Winston, 3 latencies were eliminated, one during chain trials and two during the random test. All eliminated latencies were associated with the first element position of a chain. For assessing within-subject changes during acquisition, one-tailed paired-*t* statistics or repeated measure analyses of variance were used. Preferences between members of the chain were assessed using the binomial test.

## Results and discussion

Figure 1 provides mean latencies and standard errors for the first three and last three acquisition sessions and for the random test. During acquisition and during the chain trials of the random test session, elements and serial positions were con-



**Fig. 1** Mean latencies and standard errors (*vertical lines* represent 1 standard error) across all elements for the first and last three acquisition sessions, and the random test in Experiment 1. The random test consisted of 20 chain trials followed by 20 trials in which element order was randomized. Data for the randomized trials is organized by elements (A–E) and by serial positions (1–5). During acquisition, elements and serial positions were confounded (e.g., element B always occurred in the second serial position). During the random test they were dissociated (e.g., element B could occur in any serial position)

founded (e.g., element B always occurred in the second serial position). In the random test, they were dissociated (e.g., element B occurred in each serial position four times). As a consequence, latencies for the 20 random trials during the random test session are arranged by elements (e.g., the average latency for element B across all presentations) and by serial positions (e.g., the average latency for all elements occurring in the second serial position). For Windsor there were significant reductions in the time taken to complete a chain during acquisition (first three sessions:  $M = 50.52$  s  $SD = 12.57$  s; last three sessions:  $M = 11.15$  s  $SD = 1.61$  s;

$t_{14} = 2.49$ ,  $P = 0.012$ ). During the first three sessions, Windsor's latencies did not evidence a pattern as a function of elements/serial position, but during the last three sessions his latencies showed a systematic decrease across elements/serial positions. Naturally, total time taken to complete a chain includes time to begin responding to the first element. To determine whether latencies were reduced during acquisition apart from start-up latencies, differences in running time were also calculated. Running time latencies were measured beginning with a response to the first element in a chain and terminating with a response to the last element. The changes in the pattern of Windsor's latencies during acquisition also resulted in a significant difference in running time latencies between the first three and last three sessions (first three sessions:  $M = 43.95$  s,  $SD = 13.28$ ; last three sessions:  $M = 7.46$  s,  $SD = 0.41$  s,  $t_{11} = 1.93$ ,  $P = 0.039$ ).

For Winston, there was no difference in chain latencies for the first three sessions vs. the last three sessions (first three sessions:  $M = 24.52$  s,  $SD = 3.90$  s; last three sessions:  $M = 24.08$  s,  $SD = 4.42$  s,  $t_{14} = 0.15$ ,  $P = 0.437$ ). Running time latencies also did not show a decrease as a function of acquisition (first three sessions:  $M = 16.03$  s,  $SD = 1.20$  s; last three sessions:  $M = 13.16$  s,  $SD = 2.16$  s,  $t_{14} = 0.063$ ,  $P = 0.475$ ). For Winston, latencies decreased as a function of element/serial position, and this trend is apparent within the first three sessions. For Windsor, this pattern was evident in the last three acquisition sessions, but not in the first three sessions. It can also be seen in Fig. 1 that both tamarins produced shorter latencies during the chains portion of the random test session than they produced at the end of acquisition, particularly with respect to element A. This reduction in latencies indicates that the pair-wise tests that intervened between acquisition and the random test session did not disrupt performance. Rather, performance continued to improve during the pair-wise tests.

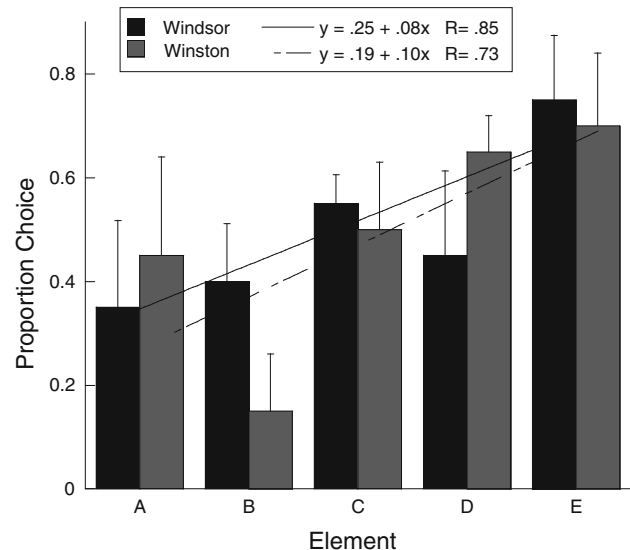
Each tamarin also demonstrated significant increases in latencies during the random test as compared to the chains baseline (Windsor: chains,  $M = 8.24$  s,  $SD = 2.77$  s; random,  $M = 14.90$  s,  $SD = 7.00$  s,  $t_{19} = 3.98$ ,  $P = 0.000$ ; Winston: chains,  $M = 11.74$  s,  $SD = 5.02$ , random:  $M = 21.72$  s,  $SD = 11.22$  s,  $t_{19} = 3.30$ ,  $P = 0.002$ ). For each tamarin there were also significant increases in running time between chains and random trials (Windsor: chains,  $M = 5.97$  s,  $SD = 2.15$  s; random,  $M = 9.95$  s,  $SD = 5.71$  s;  $t_{19} = 2.75$ ,  $P = 0.006$ ; Winston: chains,  $M = 8.45$  s,  $SD = 3.92$  s; random:  $M = 13.67$  s,  $SD = 7.78$  s;  $t_{19} = 2.48$ ,  $P = 0.011$ ). Running time latencies during the random test were calculated using the data arranged by serial position.

For each tamarin during the random test, latencies were generally higher for earlier elements and serial positions in the chain than for later elements and serial positions, as they were during acquisition. The regression of latency on elements was significant for each tamarin (Windsor:

$F_{1,3} = 35.47$ ,  $P = 0.009$ ; Winston,  $F_{1,3} = 38.35$ ,  $P = 0.008$ ). The regression of latency on serial position was significant for Windsor ( $F_{1,3} = 16.21$ ,  $P = 0.027$ ), and for the combined data from both tamarins ( $F_{1,8} = 8.23$ ,  $P = 0.020$ .) This pattern of results suggests that both elements and serial positions retained something of their relative values during the random test.

Figure 2 shows the element choices for the pair-wise tests. Both tamarins were more likely to choose later elements compared to earlier elements. The regression of proportion choice on serial position was not significant for either tamarin alone, but was significant for the combined data from both tamarins ( $F_{1,8} = 9.14$ ,  $P = 0.016$ ). It was also the case that both tamarins consistently chose element E over A. For Windsor, 5 of 5 choices were for E over A; for Winston 4 of 5 choices were for E over A. This pattern of choices was significant for the combined data from both tamarins using the binomial test,  $P = 0.010$ .

The results for training latencies, for latencies during the random test, and for the pair-wise tests, reveal evidence of perceptual learning in an implicit chains procedure. For Windsor, latencies were elevated during the random test by 81%; for Winston the elevation in latencies was 85%. These latency increases were significant, although they were slightly lower than those observed in Locurto et al. in which a behavioral pattern was correlated with the perceptual pattern (in that study, latency increases: Windsor, 81%; Winston, 99%). As was also observed in Locurto et al., the tamarins in this experiment demonstrated a significant



**Fig. 2** Proportion choice of each element during pair-wise testing in Experiment 1. Data are averaged across all pair-wise tests. Pair-wise testing consisted of presenting all possible paired elements (e.g. B and D) once during each of five sessions of testing. In each pair-wise test, reinforcement was presented for the first response to either element. In each test session pair-wise tests were intermixed with chain trials

preference in the pair-wise tests for the later element in a pair, particularly for element E over element A.

## Experiment 2

In Experiment 2, tamarins were exposed to a five-element implicit chain in which the same visual element was presented serially in the same five locations on each trial, thereby allowing a behavioral pattern to be correlated with the visual pattern, as had been done in Locurto et al. Several changes were made in this experiment compared to Experiment 1 and Locurto et al. First, the probability of reinforcement at the end of a chain was set at 0.33, instead of 1.0 as had been done in Experiment 1 and in Locurto et al. This change was instituted in light of an extensive literature that has identified differences in responding under continuous versus partial reinforcement (Gibbon et al. 1980).

Additionally, two new tests were introduced to explore further the nature of learning in an implicit chains procedure. In the Wild-Card test, the visual stimulus was changed at one serial position on selected trials to determine whether subjects were learning to perform essentially a behavioral sequence, in the form, “Touch here, then here,” etc., or whether they were also engaged in perceptual learning, that is, learning to expect a particular image to appear at particular spatial locations. In the Running-Start test, subjects received the beginning portion of a chain (A, A → B, or A → B → C) and then were given a pair-wise choice between the next two elements in the chain. For example, following A → B, the choice would be between C and D. This test presented subjects with a choice between the next element in the chain (C, in the example above) versus one that was closer to the end of the trial, and, hence, to reinforcement on 0.33 of the trials (D, in this example). This test was developed to determine the extent to which subjects’ choices were controlled principally by the temporal relation of a given element to food, that is by what are referred to as stimulus-reinforcer (S–S\*) relations. If this association were predominant, subjects would be expected to choose the element closer to food (D). If an earlier element were chosen (C), it may reflect the influence of other types of associations, those between stimuli (S–S; B and C being contiguous in the chain), or stimuli and responses (S–R; a response to C follows stimulus B).

## Method

### Subjects

Three tamarins served as subjects. Fergus and Marcel were, respectively, 7 and 6 year-old males. Spencer was a 5 year-old female. They were housed in the same colony room as

Windsor and Winston, and were maintained under the same feeding schedule as described for Experiment 1.

### Apparatus

The apparatus consisted of two touchscreens configured identically to the system used in Experiment 1. The icon was a strawberry, 200 × 200 pixels.

### Procedure

*Pre-acquisition* All subjects had served in one or more pilot studies on implicit chaining before the present study. One subject, Fergus, required three sessions of shaping before beginning acquisition training. Shaping employed an autoshaping procedure in which the strawberry icon was presented for 6 s in the same location on each trial, upper center of the touchscreen, and was followed by food. If the subject touched the image during this time, the image was removed from the screen and food was immediately presented. If no response occurred, food was presented at the end of 6 s. Sessions consisted of 40 trials. This procedure was continued until a touch response occurred to the image on 80% of the trials in a session for two consecutive sessions, and all reinforcers had been consumed in each session. For Fergus, this procedure lasted for three sessions.

In the second phase of pre-acquisition training for all subjects, the image was moved to different locations on the touchscreen, and remained there until a response terminated the trial and produced food. This phase lasted until responding occurred on each trial of the 40 trial sessions for two consecutive sessions. Sessions terminated after 1 h if the subject did not complete the 40 trials. For Fergus, that criterion was reached in 13 sessions; for Marcel and Spencer the criterion was reached in 5 sessions.

*Acquisition* Following completion of this phase, training began on the five-item chain. As was the case in Experiment 1, the first response to an element advanced the chain to the next element. Sessions consisted of 40 trials. Touches to blank areas of the screen were recorded but had no consequences. A 5 s blank white screen preceded the first element in a trial as a ready signal. A 20 s ITI, a dark screen, followed reinforcement.

The probability of reinforcement at the end of a trial was progressively decreased from 1.0 at the beginning of acquisition to 0.67, then to 0.50, and then to 0.33 at the end of acquisition. At each reinforcement level, subjects were required to meet the criterion that sessions had to be completed with ±10% of the average session time in four the last five sessions. Fergus met these successive criteria in a total of 61 sessions across the four reinforcement levels, Marcel in 35 sessions, and Spencer in 72 sessions.

**Testing: pair-wise tests** Pair-wise tests were conducted for 5 sessions following acquisition using the same method used in Experiment 1. Reinforcement followed the first choice during a pair-wise trial with a probability of 0.33.

**Testing: wild-card test** This test was conducted during one 40-trial session following the completion of the pair-wise tests. Following the first ten chain trials in this session, a wild card, a novel image, was periodically substituted for the standard strawberry image. The novel images covered a wide range of types, from objects (e.g., a camera, a can of soda, an electric drill) to natural scenes (e.g., a snowflake, a galaxy, planet Earth). The substitution was made twice at each serial position during these trials, resulting in a total of 10 wild-card trials in this session. A chain could contain only one wild card, and at least one standard chain was interspersed between wild-card trials. Each wild card was used only once per subject.

**Testing: running-start test** This test was conducted during one session following the wild-card test session. Following the first ten chain trials, subjects were periodically presented with the start of a chain, and were then given a pair-wise choice between the next two elements in the chain. The starting points of the chain were either an  $A \rightarrow$  start followed by a B/C choice, an  $A \rightarrow B$  start followed by a C/D choice, or an  $A \rightarrow B \rightarrow C$  start followed by a D/E choice. Each type of start was presented 4 times during the session, resulting in a total of 12 running start trials in the test session. The first response to either element during the pair-wise choice terminated the trial. Subjects' choices did not determine whether reinforcement was delivered. Reinforcement was presented with a probability of 0.33 following a response to either element. A trial could contain only one Running-Start test.

**Testing: random sequences** The random test session followed the Running-Start test. In this session, as was the case in Experiment 1, the first 20 trials were chain trials, identical to those in acquisition. For the next 20 trials, the elements were presented in randomized order, with reinforcement delivered at the end of each chain with a probability of 0.33, as was done during acquisition. Randomization of elements was accomplished with two constraints: on any trial an element could occur only once; and, each element occurred four times in each serial position during the test. As in Experiment 1, latencies that fell more than two standard deviations from the mean were eliminated from analyses. By this criterion, during the random test session, of the 200 latencies recorded for each subject during the random test session, 4 latencies were eliminated for Fergus, 5 for Marcel, and 6 for Spencer.

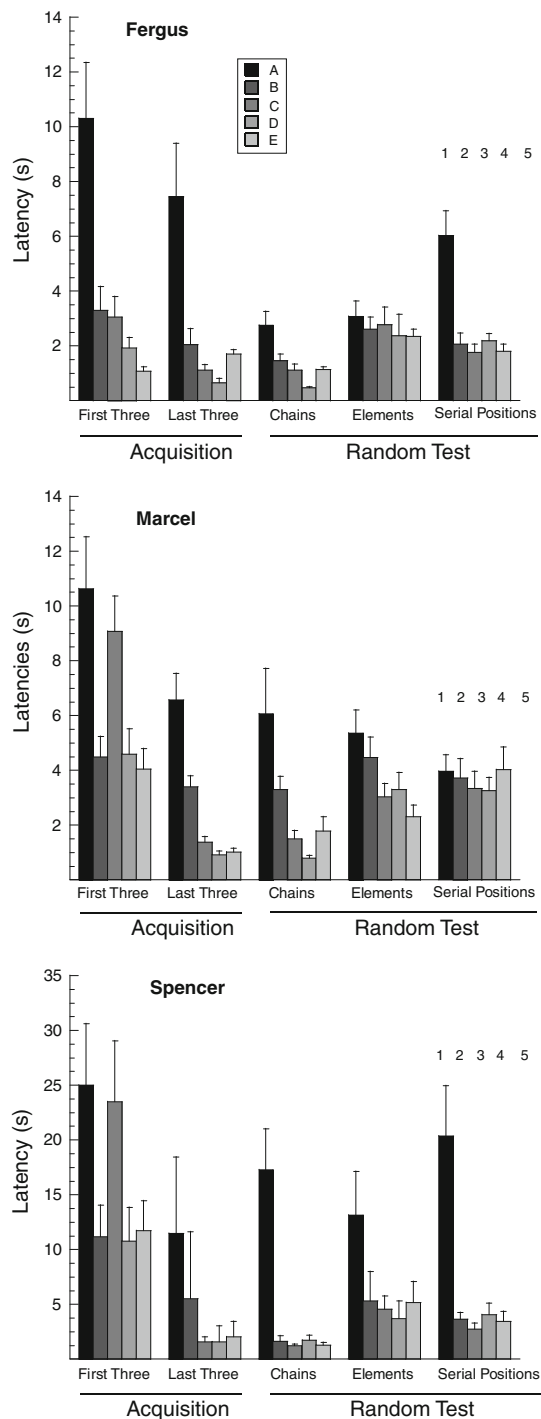
## Results and discussion

Figure 3 provides the mean latencies and standard errors across elements and serial positions during the first and last three acquisition sessions and during the random test session in Experiment 2. There were significant reductions in the time taken to complete a chain between the first and last three acquisition sessions for each subjects (Fergus: first three,  $M = 19.7$  s,  $SD = 3.52$  s; last three,  $M = 13.01$  s,  $SD = 2.95$  s,  $t_{14} = 2.25$ ,  $P = 0.020$ ; Marcel: first three sessions,  $32.83$  s,  $SD = 4.33$  s; last three sessions,  $M = 13.21$  s,  $SD = 2.56$  s,  $t_{14} = 3.69$ ,  $P = 0.002$ ; Spencer: first three sessions,  $M = 22.23$  s,  $SD = 12.24$  s; last three sessions,  $M = 22.34$  s,  $SD = 4.26$  s,  $t_{14} = 3.66$ ,  $P = 0.001$ ). The first three acquisition sessions were taken from sessions in which the probability of reinforcement following a trial was set at 1.0. The last three sessions were taken from sessions in which the probability of reinforcement was 0.33. There were also significant reductions in running time latencies for each subject (Fergus: first three,  $M = 11.71$  s,  $SD = 1.30$  s; last three,  $M = 6.91$  s,  $SD = 0.62$  s,  $t_{14} = 2.51$ ,  $P = 0.014$ ; Marcel: first three,  $M = 27.75$  s,  $SD = 4.19$  s; last three,  $M = 8.38$ ,  $SD = 1.09$  s,  $t_{14} = 3.09$ ,  $P = 0.005$ ; Spencer: first three,  $M = 71.45$  s,  $SD = 11.34$  s; last three,  $M = 13.45$  s,  $SD = 1.80$  s,  $t_{14} = 3.37$ ,  $P = 0.003$ ).

It can also be seen that latencies during the chains portion of the random test session were similar to the latencies recorded at the end of acquisition. For Fergus and Marcel, the latencies during the chains portion of the random test session were lower than those produced at the end of acquisition. For Spencer latencies appeared to be comparable to those recorded at the end of acquisition. For Spencer, latencies to element A were higher during the chains portion of the random test session, but latencies to the remaining elements were lower compared to their acquisition baseline. The similarities between these two sets of latencies for each subject indicate that the tests that intervened between the end of acquisition and the chains portion of the random test session (Pair-Wise test, the Wild-Card test, and the Running-Start test) did not disrupt performance during chain trials.

Each tamarin produced significantly higher latencies during the random test than during the chains portion of the random test session. These elevations in latency averaged 69% across the three tamarins. Increases in latency were noted for total chain time (Fergus: chains,  $M = 6.96$  s,  $SD = 4.79$  s; random,  $M = 13.88$  s,  $SD = 9.27$  s,  $t_{19} = 5.22$ ,  $P = 0.002$ ; Marcel: chains,  $13.45$  s,  $SD = 2.68$  s; random,  $M = 21.81$  s,  $SD = 4.07$  s,  $t_{19} = 2.52$ ,  $P = 0.010$ ; Spencer: chains,  $M = 21.69$  s,  $SD = 2.16$  s; random,  $M = 32.66$  s,  $SD = 3.26$  s,  $t_{19} = 2.10$ ,  $P = 0.024$ ). Elevations in latency were also noted for running times (Fergus: chains,  $M = 5.26$  s,  $SD = 0.65$  s; random,  $M = 9.81$  s,  $SD = 1.32$  s,  $t_{19} = 4.82$ ,  $P = 0.000$ ; Marcel: chains,  $M = 9.23$  s,  $SD = 2.03$  s, random,  $M = 20.12$  s,






**Fig. 3** Mean latencies and standard errors (*vertical lines* represent 1 standard error) across all elements for the first and last three acquisition sessions, and the random test in Experiment 2. During the first three acquisition sessions, the probability of reinforcement at the end of a trial was 1.00. That value was set at 0.33 during the last three acquisition sessions. The random test consisted of 20 chain trials followed by 20 trials in which element order was randomized. Data for the randomized trials is organized by elements (A–E) and by serial positions (1–5). During acquisition elements and serial positions were confounded (e.g., element B always occurred in the second serial position). During the random test they were dissociated (e.g., element B could occur in any serial position)

SD = 1.83 s,  $t_{19} = 4.72$ ,  $P = 0.007$ ; Spencer: chains,  $M = 7.39$  s, SD = 0.739 s, random,  $M = 17.46$  s, SD = 1.81,  $t_{19} = 4.63$ ,  $P = 0.008$ ). The elevations in running time indicate, as they did in Experiment 1, that the disruptions owing to the random test were not simply of function of increased latencies for the first element in a trial.


As was true in Experiment 1, inspection of the latencies across elements and serial positions during the random test suggested that both elements and serial positions retained something of their relative values during the random test. Fergus and Spencer recorded higher latencies to the first serial position than to any other position during the random test. Similarly, for each subject higher latencies were associated with element A than for any other element. The regression of latency on elements was significant for each tamarin ( $P < 0.05$ ) and for the combined sample ( $F_{1,12} = 8.45$ ,  $P = 0.013$ ). The regression of latency on serial positions was not significant for any tamarin, nor for the combined sample (all values:  $0.05 < P < 0.10$ ).

Figure 4 shows element choices during the pair-wise tests. There was a tendency to choose later elements over earlier elements, although for Marcel and Spencer, element B was chosen more often than would be expected by this generalization. The preference for element B may represent a pre-existent preference that was not eliminated by acquisition training, although there was no evidence in the acquisition latencies of a preference for this element. The regression of proportion choice on serial position was not significant for individual tamarins ( $P > 0.05$ ), but was significant for the combined sample ( $F_{1,12} = 6.67$ ,  $P = 0.022$ ). As was true in Experiment 1, subjects predominantly chose element E over element A. Each tamarin chose element E in 4 out of 5 A/E choices ( $P = 0.017$  for the combined sample).

Figure 5 provides latency summaries for the Wild-Card test for each subject across elements. The comparison of chain and wild-card latencies was accomplished by comparing latencies at each serial position during the chain trials of this test session with the wild-card latencies. It can be seen that the presentation of wild cards resulted in appreciably higher latencies for each subject (Fergus: total chain time,  $M = 9.67$  s, SD = 1.31 s; wild card,  $M = 153.63$  s, SD = 25.63 s,  $t_5 = 2.64$ ,  $P = 0.028$ ; Marcel: total chain time, 38.53 s, SD = 8.24 s; wild card,  $M = 264.40$  s, SD = 35.07 s,  $t_5 = 3.14$ ,  $P = 0.017$ ; Spencer: total chain time,  $M = 16.03$  s, SD = 3.34 s; wild card,  $M = 36.60$  s, SD = 4.53 s,  $t_5 = 2.08$ ,  $P = 0.522$ ). It was also the case that latency increases for wild cards were observed at each serial position for each subject, with the exceptions of element C for Spencer and element D for Marcel. While the increases were on average higher for elements A and E and lower for internal elements C and D, an analysis of variance found no significant differences in wild-card latencies across wild-card elements

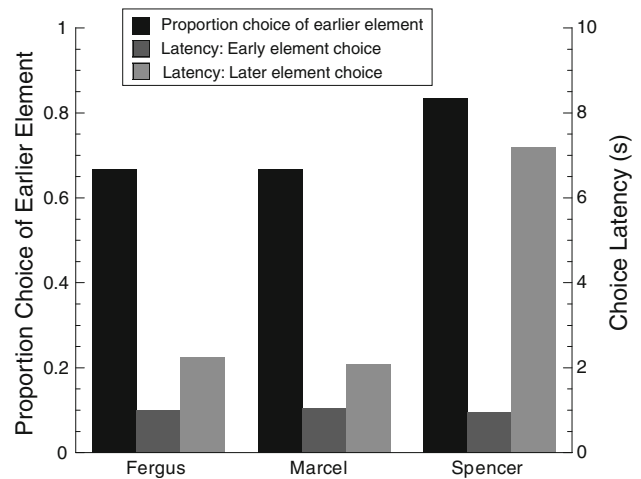


**Fig. 4** Proportion choice of each element during pair-wise testing in Experiment 2. Data are averaged across all pair-wise tests. Pair-wise testing consisted of presenting all possible paired elements (e.g. B and D) once during each of five sessions of testing. In each pair-wise test, reinforcement was presented for the first response to either element with a probability of 0.33. In each test session pair-wise tests were intermixed with chain trials



**Fig. 5** Latencies for chain trials and wild-card trials during the Wild-Card test session. Wild cards were novel images that occurred twice at each serial position during the test session. The session began with ten chain trials. Thereafter, chain trials were intermixed with wild-card trials in which one element in the chain was replaced by a novel image

( $P > 0.10$ ), and the regression of difference scores (wild card–chain) on elements was not significant for linear, quadratic or cubic fits ( $P > 0.10$  for each analysis). These



**Fig. 6** Element choices and latencies during the Running-Start test. Subjects received part of a chain, either A, or  $A \rightarrow B$ , or  $A \rightarrow B \rightarrow C$ , before being presented with a choice between two elements: the first element was the next element in the chain, and the second element was the following element. The following element was closer to food than the next element. Results are provided for the proportion of choices of the earlier element in a pair (read from the Y1 axis) and for latencies for choices of earlier and later elements choices (read from the Y2 axis)

results reveal that subjects were responding not only in terms of spatial location, in the form “Touch here, then touch here,” but were also observant of the particular element that occurred at each spatial location.

Figure 6 summarizes the results of the Running-Start test. Results are provided for the proportion of choices for the earlier element in a pair (read from the Y1 axis), as well as for latencies associated with choices of early and later elements (read from the Y2 axis). Each type of running-start trial was presented four times during this test session, that is, four trials each for:  $A \rightarrow$ ,  $A \rightarrow B$ ,  $A \rightarrow B \rightarrow C$  running starts, resulting in 12 choices for each subject. Given the relatively small number of choices for each type of running start trial, data were summarized across the three types of running start trials. As can be seen in Fig. 6, each subject chose the earlier element over a later element on the majority of trials. Moreover, their latencies when choosing the earlier element in a pair were shorter than when the choice was for the later element in a pair. The choice of the earlier element in a pair was significant for the combined sample (binomial test,  $P < 0.006$ ). Additionally, the latencies for choices of earlier elements were significantly shorter than the latencies associated with later choices (Wilcoxin signed-rank test,  $Z = 3.99$ ,  $P < 0.01$ ). Recall that this test pitted the next element in the chain against the following element, that is, one that was closer to the end of the trial, and, hence, closer to food on 0.33 of the trials. The data indicate that subjects chose next elements over elements closer to food, and they chose those earlier elements more quickly than they chose elements closer to food.

Overall, the pattern of latencies observed in this experiment during acquisition and the random test appears highly similar to the pattern observed in Experiment 1 and in Locurto et al. In all three studies, there was an orderly decrease in latencies across serial positions by the end of acquisition, and increases in latencies during the random test that were evident across elements and serial positions. As noted earlier, the lack of difference between Experiment 1 and Locurto et al. indicates that the magnitude of implicit learning was approximately equal whether or not elements were associated with unique spatial locations. The similarities between Experiment 2 and these other experiments suggest that the effects of the partial reinforcement schedule used in this experiment did not produce an appreciably different pattern of responding compared to procedures in which reinforcement probability was set at 1.0. The similarities between Experiment 2 and these other experiments may be seen as atypical compared to the large number of studies that have observed different patterns of performance between partial and continuous reinforcement (Gottlieb 2006). It should be noted that the frequently reported differences between partial and continuous reinforcement are centered on acquisition and extinction. These aspects of performance were not studied in this experiment. Rather, maintained performance was the focus of this experiment. In this light, there is a smaller body of evidence that has documented the finding that under some conditions, patterns of maintained performance under partial schedules may be similar to the patterns observed under continuous reinforcement (e.g., Gibbon et al. 1980).

## General discussion

The intent of these experiments was to broaden understanding of the nature of learning in an implicit chains procedure. In doing so, this work touches on a several issues that have been important both in the human implicit learning literature and in the nascent comparative literature that attempts to extend implicit learning to nonhumans. Perhaps the overarching issue attendant to this work is whether it is at all reasonable to speak of studying implicit learning in a nonhuman organism, given that the very name “implicit” refers to human subjects’ inability to articulate the rules of a sequence learning task, while their performance gives evidence of some appreciation of those rules (e.g., Reber 1996). As was suggested earlier, the study of implicit learning may be extended to nonhumans with the following provisions: (1) reinforcement does not depend on the acquisition of specified information, and (2) following acquisition, subjects’ knowledge of that information can be ascertained. The experiments reported here accommodated these requirements. Considered together with Locurto et al.

and a number of other studies, there is growing evidence that nonhumans in this procedure learn more about the sequential nature of the chain than is required by the contingencies of reinforcement.

From a methodological standpoint these results may be characterized as illustrating the advantages to using a procedure in which, unlike the standard SRT task, performance associated with individual elements and serial positions can be teased apart within the framework of implicit learning (See Reed and Johnson 1994, for strategies allowing analysis what is learned in conventional SRT procedures; see Hunt and Aslin 2001, for a similar approach in a format designed to study statistical learning). The analytical tests used in these experiments beyond the random test which is standard in most SRT studies, namely the pair-wise tests, wild-cards, and the running starts, are additional strategies that facilitate analysis of the learning that occurs under implicit conditions.

It is clear that the learning in this implicit procedure includes a form of perceptual learning under conditions where there was no association between an element and a unique spatial location. In Experiment 1, the disruptions during the random test and the pair-wise choices were of similar form and magnitude compared to the results observed by Locurto et al. in which perceptual and behavioral elements were combined. There was further evidence of perceptual learning in Experiment 2 using the Wild-Card test. The large-scale increases in latency to wild cards indicates that within a procedure that associated serial positions with particular spatial locations, subjects were nonetheless attendant to the nature of the element that appeared at each location. This type of learning was not mandated by the experimental contingencies, nor was it confined only to those elements that were temporally close to food. Wild-card disruptions were evident throughout the five-element chain. It should be added that this interpretation leaves open the more precise reason why latencies were elevated during the wild-card test. It may be that the increased latencies reflect the violation of an expectation for an element to occur at a particular location, or a more generalized violation that a new stimulus has been introduced into the perceptual field of the subject.

Perhaps the most interesting question arising from these experiments concerns the nature of the learning that occurs under implicit conditions. Given that reinforcement was delivered only at the end of trials, the possibility cannot be dismissed that these results are due to the differing temporal relations that elements enjoyed with reinforcement (S–S\*). By this account, elements closer to reinforcement would be of higher associative value than elements further removed from reinforcement. From this perspective, it would be expected that later elements would be chosen over earlier

elements in the pair-wise tests. As earlier noted, in the literature on explicit sequence learning there is a substantial body of evidence supporting the conclusion that subjects come to understand the ordinal nature of the chain apart from, or independent of, the associative strengths of the elements (Bond et al. 2003; Conway and Christiansen 2001; D'Amato 1991; Rapp, et al. 1996; Scarf and Colombo 2008; Terrace 2005).

In the present implicit procedure, ruling out this type of associative strength explanation, one based on S–S\* relations, is not possible. There is, however, evidence that points to aspects of what is learned in this procedure that are not explicable in terms of stimulus–food relations. The strongest evidence comes from the Running–Start test in Experiment 2 in which subjects tended to choose next elements over elements temporally closer to food. The choice of next elements cannot be attributed to an element's value based on S–S\* relations, since the value of next elements was *per force* weaker than that of the following element based on S–S\* relations. The choice of next elements in this test suggests that the tamarins' choices were influenced by S–S and/or S–R relations, that is, relations between stimuli in the chain, or connections between stimuli and responses. Either or both of these relationships might lead to next-element choices in the Running Start test.

Interestingly, although S–R associations are the most prominent type of association cited in connection with behavioral theories of learning, their presence may also be the most difficult to document. One source of evidence for the presence of S–R associations has come through devaluation experiments. In these procedures, acquisition is followed by some form of devaluation of the reinforcer, usually through a taste-aversion procedure. It is often observed that following devaluation responding continues at some non-zero level. This persistent responding has been taken as evidence for the existence of S–R associations that were formed prior to the devaluation manipulation and, therefore, did not depend on the present value of the reinforcer (Colwill and Rescorla 1986). It may be that the results of the Running–Start test add to the evidence that S–R associations may be functionally independent of the value of a stimulus based on its current relation to a reinforcer.

Considered broadly, these results should be placed within the context of other nonhuman analogs of implicit learning. These analogs include not only SRT procedures, but also statistical learning in cotton-top tamarins (Fitch and Hauser 2004) and pigeons (Froehlich et al. 2004), recursion in song bird learning (Gentner et al. 2006), repetition priming in pigeons (Blough 1993), and artificial grammar learning in pigeons (Herbranson and Shimp 2003). These analogs hold the possibility of becoming important contributions to the study of comparative cognition. They may enhance our perspective on precisely which aspects of

human cognition find no parallels in nonhuman cognition, and which aspects may share common characteristics with nonhuman cognition. The present work adds to this effort by indicating that without explicit reinforcement for correct responding, nonhuman subjects nevertheless learn several features of patterned information, including the ordinal position of elements in a series and the perceptual characteristics of those elements.

**Acknowledgments** This work was funded by grant BCS-0116089 from the National Science Foundation to the College of the Holy Cross, and by grant RR00168 from the National Institutes of Health to the New England Primate Research Center (NEPRC), Southborough, MA. We thank the staff at NEPRC for its support of this work.

## References

- Blough D (1993) Effects on search speed of the probability of target-distracter combinations. *J Exp Psychol Anim Behav Pro* 19:231–243
- Bond AB, Kamil AC, Balda RP (2003) Social complexity and transitive inference in corvids. *Anim Behav* 65:479–487
- Christie MA, Dalrymple-Alford JC (2004) A new rat model of the human serial reaction time task: contrasting effects of caudate and hippocampal lesions. *J Neurosci* 24:1034–1039
- Christie MA, Hersch SM (2004) Demonstration of nondeclarative sequence learning in mice: development of an animal analog of the human serial reaction time task. *Learn Mem* 11:720–723
- Clegg BA, DiGirolamo GJ, Keele SW (1998) Sequence learning. *Trends Cogn Sci* 2:275–281
- Conway CM, Christiansen MH (2001) Sequential learning in non-human primates. *Trends Cogn Sci* 5:539–546
- Cowill EM, Rescorla RA (1986) Associative structures in instrumental learning. In: Bower GH (ed) *Psychol learn mot*. Academic Press, San Diego, pp 55–104
- Crofts HS, Muggleton N, Bowditch AP, Peach PC, Nutt DJ, Scott EAM (1999) Home cage presentation of complex discrimination tasks to marmosets and rhesus monkeys. *Lab Anim* 33:207–214
- D'Amato MR (1991) Comparative cognition: processing of serial order and serial pattern. In: Dachowski L, Flaherty CF (eds) *Current topics in animal learning: Brain, emotion and cognition*. Erlbaum, Hillsdale, NJ, pp 165–185
- Davis H (1992) Transitive inference in rats (*Rattus norvegicus*). *J Comp Psychol* 106:342–349
- Deroost N, Soetens E (2006) Perceptual or motor learning in SRT tasks with complex sequence structures. *Psychol Res* 70:88–102
- Domenger D, Schwarting RKW (2005) Sequential behavior in the rat: a new model using food-reinforced instrumental behavior. *Behav Brain Res* 160:197–207
- Eimer M, Goschke T, Schlaghecken F, Stürmer B (1996) Explicit and implicit learning of event sequences: Evidence from event-related potentials. *J Exp Psychol Learn Mem Cogn* 22:970–987
- Fitch WT, Hauser MD (2004) Computational constraints on syntactic processing in a nonhuman primate. *Science* 303:377–380
- Froehlich AL, Herbranson WT, Loper JD, Wood DM, Shimp CP (2004) Anticipating by pigeons depends on local statistical information in a serial response time task. *J Exp Psychol Gen* 133:31–45
- Gentner TQ, Fenn KM, Margoliash D, Nusbaum HC (2006) Recursive syntactic pattern learning by songbirds. *Nature* 440:1204–1207
- Gibbon J, Farrell L, Locurto CM, Duncan HJ, Terrace HS (1980) Partial reinforcement in autoshaping with pigeons. *Anim Learn Behav* 8:45–59

- Goschke T (1998) Implicit learning of perceptual and motor sequences: evidence for independent learning systems. In: Stadler MA, Frensch PA (eds) *Handbook of implicit learning*. Sage, Thousand Oaks, CA, pp 401–444
- Gottlieb DA (2006) Effects of partial reinforcement and time between reinforced trials on terminal response rate in pigeon autoshaping. *Behav Process* 72:6–13
- Hauser M, Newport E, Aslin R (2001) Segmentation of the speech stream in a non-human primate: statistical learning in cotton-top tamarins. *Cognition* 78:B53–B64
- Herbranson WT, Shimp CP (2003) Artificial grammar learning in pigeons: a preliminary analysis. *Learn Behav* 31:98–106
- Hunt RH, Aslin RN (2001) Statistical learning in a serial reaction time task: access to separable statistical cues by individual learners. *J Exp Psychol Gen* 130:658–680
- Locurto C, Gagne M, Levesque K (2009) Implicit chaining in cotton-top tamarins (*Saguinus oedipus*). *J Exp Psychol Anim Behav Process* 35:116–122
- Merritt D, Maclean EL, Jaffe S, Brannon EM (2007) A comparative analysis of serial ordering in Ring-Tailed Lemurs (*Lemur catta*). *J Comp Psychol* 121:363–371
- Nissen MJ, Bullemer P (1987) Attentional requirements of learning: evidence from performance measures. *Cogn Psychol* 19:1–32
- Procyk E, Dominey PF, Amiez C, Joseph JP (2000) The effects of sequence structure and reward schedule on serial reaction time learning in the monkey. *Cogn Brain Res* 9:239–248
- Ramus F, Hauser MD, Miller CT, Morris D, Mehler J (2000) Language discrimination by human newborns and cotton-top tamarin monkeys. *Science* 288:349–351
- Rapp PR, Kansky MT, Eichenbaum H (1996) Learning and memory for hierarchical relationships in the monkey: effects of aging. *Behav Neurosci* 110:887–897
- Reber AS (1967) Implicit learning of artificial grammars. *J Verb Learn Verb Behav* 6:855–863
- Reber AS (1996) *Implicit learning and tacit knowledge: an essay on the cognitive unconscious*. Oxford, New York
- Reed J, Johnson P (1994) Assessing implicit learning with indirect tests: determining what is learnt about sequence structure. *J Exp Psychol Learn Mem Cogn* 20:585–594
- Remillard G (2003) Pure perceptual-based sequence learning. *J Exp Psychol Learn Mem Cogn* 29:581–597
- Rowland LA, Shanks DR (2006) Sequence learning and selection difficulty. *J Exp Psychol Human Percept Perform* 32:287–299
- Scarf D, Colombo M (2008) Representation of serial order: a comparative analysis of humans, monkeys, and pigeons. *Brain Res Bull* 76:307–312
- Seger CA (1994) Implicit learning. *Psychol Bull* 115:163–196
- Soetens E, Melis A, Notebaert W (2004) Sequential effects and sequence learning. *Psychol Res* 10:124–137
- Terrace HS (1993) The phylogeny and ontogeny of serial memory: list learning by pigeons and monkeys. *Psychol Sci* 4:162–169
- Terrace H (2005) The simultaneous chain: a new approach to serial learning. *Trends Cogn Sci* 9:202–210
- Willingham DB, Nissen MJ, Bullemer P (1989) On the development of procedural and declarative knowledge. *J Exp Psychol Learn Mem Cogn* 15:1047–1060