ORIGINAL PAPER

Implicit chaining in cotton-top tamarins (*Saguinus oedipus*) with elements equated for probability of reinforcement

Charles Locurto · Laura Dillon · Meaghan Collins · Maura Conway · Kate Cunningham

Abstract Three experiments examined the implicit learning of sequences under conditions in which the elements comprising a sequence were equated in terms of reinforcement probability. In Experiment 1 cotton-top tamarins (*Saguinus oedipus*) experienced a five-element sequence displayed serially on a touch screen in which reinforcement probability was equated across elements at .16 per element. Tamarins demonstrated learning of this sequence with higher latencies during a random test as compared to baseline sequence training. In Experiments 2 and 3, manipulations of the procedure used in the first experiment were undertaken to rule out a confound owing to the fact that the elements in Experiment 1 bore different temporal relations to the intertrial interval (ITI), an inhibitory period. The results of Experiments 2 and 3 indicated that the implicit learning observed in Experiment 1 was not due to temporal proximity between some elements and the inhibitory ITI. The results taken together support two conclusion: First that tamarins engaged in sequence learning whether or not there was contingent reinforcement for learning the sequence, and second that this learning was not due to subtle differences in associative strength between the elements of the sequence.

C. Locurto (✉) · L. Dillon · M. Collins · M. Conway · K. Cunningham
Department of Psychology, College of the Holy Cross, Worcester, MA 01610, USA
e-mail: clocurto@holycross.edu

Introduction

Much human learning occurs in the absence of explicit instructions. There has been considerable interest in this form of learning inasmuch as many human motor skills and abilities ranging from learning to ride a bicycle to language acquisition appear to be learned in this manner. Reber (1967) was the first to study this form of learning. In his procedure, subjects were presented with letter strings that were generated 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 subjects who were presented with letter strings that had been composed randomly with respect to the grammar. Following training, a number of experimental subjects were unable to articulate the rules of that grammar, an observation that led to the name "implicit" for this type of learning (Reber 1996). In subsequent usage, the term implicit learning has come to refer to learning that occurs without explicit instructions to learn and without subjects' awareness of the experimental contingencies (see Locurto et al. 2010).

The artificial grammar design developed by Reber has been used in a large number of subsequent studies. A second form of implicit learning, the serial reaction time task (SRT), was 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 asterisk's location followed a repeating pattern and subjects were required to tap an arbitrarily chosen keyboard key when the asterisk appeared at each stimulus location. The form followed the pattern DBCACBDCBA, where each letter referred to one of four

spatial locations on the screen where the asterisk might occur. The string was presented without breaks between the end of one presentation of the string and the beginning of the next presentation. Results indicated that the latencies of experimental subjects were lower than those of control subjects who experienced the same stimuli presented randomly. Further, experimental subjects often were unable to articulate the nature of the pattern, particularly when the task was made more complex by introducing a second task to be performed concurrently during pattern training (Nissen and Bullemer, Experiment 2).

Subsequent work has confirmed Nissen and Bullemer's basic findings using variations of the initial SRT procedure (reviews: Clegg et al. 1998; Seger 1994). One variation marries the SRT task with Reber's original artificial grammar procedure, so that training strings are developed using an artificial grammar. This variation makes presentation of an individual element probabilistic given that artificial grammars permit some variation in terms of which element may be next in the sequence. 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 position that implicit learning has come to occupy in the study of human cognition makes it valuable to ask whether this form of learning is unique to humans or has comparative origins. It might be argued that there can be no nonhuman analog of human implicit learning inasmuch as the there can be no assessment of subjects' awareness of a pattern, one of the striking features of human implicit learning. In contrast to this line of reasoning, Locurto et al. (2009) argued that an implicit learning procedure for a nonhuman would present patterned information to subjects, but reinforcement would not depend on learning that pattern. The question under consideration would then be whether in this format a nonhuman learns more than is demanded by the experimental contingencies. In support of this idea, there have been a number of adaptations of the basic SRT procedure for nonhumans in which subjects are exposed to a serial pattern of stimuli and they 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. 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 2004; in mice: Christie and Hersch 2004; in rhesus macaques: Procyk et al. 2000; in pigeons: Herbranson and Stanton 2011; in Bengalese finches: Yamazaki et al. 2012). These results, which parallel the basic findings in human

SRT studies, indicate that nonhuman subjects learned something about the pattern.

While confirming that learning took place, these results do not address with precision the issue of what subjects learned, particularly whether they learned something about serial order. To illustrate this issue, consider procedures developed from a parallel literature involving both humans and nonhumans, the study of explicit serial learning. One standard procedure, termed simultaneous chaining, proceeds in the following manner: If the to-be-learned pattern includes, as is typical, five arbitrarily chosen visual elements (A through E), training begins with the presentation of element pair A and B. The subject must choose the two stimuli in the correct order, A then B, to receive reinforcement. When that order is mastered, the sequence is increased to A → B → C. and so on until the entire 5-element sequence must be responded to in correct order for reinforcement (Merritt et al. 2007; Terrace 1993, 2005). After an acquisition criterion is reached, testing proceeds using a number of partial element sequences (e.g., B–D, A–E) to determine subjects' understanding of the ordinal position of the elements. Results from a substantial body of literature indicate that subjects, human and nonhuman, are able to order the test sequences correctly, thereby indicating that they appreciate the ordinal nature of the chain (Conway and Christiansen 2001, D'Amato 1991, provide reviews of other forms of explicit sequence learning in nonhumans).

The importance of ordinal learning of this sort cannot be over-emphasized. Appreciation of serial order has been implicated as being of fundamental importance in a number of domains beyond serial learning itself, including the development of numerical competence, the ontogeny of language, and the development of forms of inferential reasoning, including transitivity (Merritt and Terrace 2011; Terrace and McGonigle 1994. For examples of transitive inference: in rats, Davis 1992; in corvids, Bond et al. 2003; in rhesus macaques, Rapp et al. 1996; see Scarf and Colombo 2008, for comparative analyses).

It would be valuable to adapt this technology to an SRT task to understand what has been learned about serial 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 which aspects of the pattern have been learned. One feature is that the same stimulus location is used repeatedly throughout the pattern. In Nissen and Bullemer's study, for example, some element positions were used two or three times in one complete presentation of the pattern. As a consequence, a particular element was not uniquely associated with a particular ordinal position. 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 learn the sequential nature of the pattern.

To address these limitations, Locurto et al. (2009) developed a procedure that wedded 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 touch screen. 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 dark-screen intertrial interval (ITI). In this procedure, which is a variant of a fixed-ratio schedule, 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. 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.

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 ordered 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.

In a subsequent study, Locurto et al. (2010) extended this basic procedure. In one experiment, the elements were five different visual stimuli that could occur in any one of six positions on the touch screen, thereby separating the visual characteristics of an element from its spatial location. In a second experiment, using the procedure in which the same visual element occurred in different spatial locations, two new tests were introduced following training in addition to pair-wise tests and the random test. One test, termed the wild-card test introduced a new visual element in the same location as an expected element periodically to determine whether subjects were attending to the perceptual nature of the elements or were learning essentially a behavioral pattern (i.e., touch location A then location B, etc.). In a second test termed the running start test, the pattern was interrupted periodically and a choice was given to the subject that consisted of the next two elements in the string (e.g., A → B → C, then the D/E pair-wise test). This test presented subjects with a choice between the next element in the chain (D, in the example above) versus one that was closer to the end of the trial and, hence, closer to

reinforcement (E). 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. If this association were predominant, subjects would be expected to choose the element closer to food (E). If an earlier element were chosen (D), it would reflect the influence of other types of associations, those between stimuli (C and D being contiguous in the chain), or sub-chains of stimuli and responses (a response to C is followed by D, which is followed by a response to D, etc.).

The results of the first experiment indicated that implicit learning occurred in approximately the same magnitude as in Locurto et al. (2009) when the visual characteristics of an element were separated from its spatial location. In this second experiment, it was found that subjects were indeed attending to the perceptual nature of the elements in the wild-card test, as latencies increased significantly when new visual elements replaced the expected. In the running start test, subjects chose next elements over elements closer to reinforcement on nearly 70 % of choice tests, thereby indicating that learning in this procedure was not simply mediated by the associative relations between the elements and food.

These results extend the basic finding that sequence learning will occur under conditions in which reinforcement is not contingent on such learning. In one sense, however, these results were generated from a set of limited conditions. The limitation comes from the procedural aspect of these studies that reinforcement was programmed to be delivered at the end of A → E sequences. As a consequence, elements in these studies differed in their temporal relations to reinforcement, that is, their associative strengths. These differences in associative strength may have mediated some if not all of the observed learning. The finding in Locurto et al. (2010, Experiment 2) that other relations, those between stimuli and/or between stimuli and responses, were concurrently learned, does not diminish this fundamental difference between the elements in the A → E chains.

The present experiments were designed to eliminate this difference between elements in terms of associative strength, within the context of an implicit learning task. In Experiment 1, reinforcement was programmed to occur with equal probability following each element in an A → E chain. While this procedure eliminated between-element differences in associative strength, it did not eliminate differences owing to the fact that the elements bore different relations to the ITI which might be considered an inhibitory period in that reinforcement never occurs during this time. To eliminate this source of difference between elements, in Experiment 2, the time between all elements, including the time between elements E and A, the ITI, was made equal at 1 s. In Experiment 3, the

associative strength between elements was made equal by using strings generated by an artificial grammar. In this procedure, each element began and ended the strings with equal probability.

Experiment 1

Methods

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. Prior to the present study, both tamarins served as subjects in other implicit learning experiments (Locurto et al. 2009, 2010). 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. Daily feeding consisted of a morning feeding and a second one in the afternoon. On training days, they received their morning feeding for approximately 2 h. Feedings consisted of monkey chow supplemented with fresh fruit. Food was removed 3–4 h prior to training, which took place in the early afternoon. 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.

Apparatus

The stimulus consisted of an image of a banana approximately 200 × 200 pixels displayed against a white background on a touch screen (ELO CarrollTouch infrared; Tyco Electronics, Berwyn, PA). The screen was 38.1 cm on the diagonal. The front of the touch screen was fitted with a clear Plexiglas bezel that allowed access to the touch screen at six locations (10 cm × 10 cm holes) arranged in a 2 × 3 array across the top of the touch screen. The image was displayed in approximately the center of each location. 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 touch screen via a rotary pellet feeder (Coulbourn Instruments, Allentown, PA).

Procedure

Subjects were tested in their home cage. The touch screen 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 touch screen and food dish.

Baseline training

Since both tamarins were experienced, we presented them on the first day of training with a full set of training stimuli. Sessions terminated after 40 trials or after 50 min, whichever criterion was met first. Each trial consisted of the presentation of the five-element spatial chain. Each response produced an auditory stimulus approximately 100 ms in duration. The first touch to an element advanced the chain to the next element without delay. 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 an element with a probability of .16 for each element. If reinforcement were programmed for an element, a screen touch of the element changed the image to one with a blue border for approximately one second concurrent with the delivery of a reinforcer. The screen was blank for 1 s during the presentation of a reinforcer. Following reinforcement, the sequence continued. A 20 s ITI, a black screen, followed element E.

Training proceeded until each tamarin completed at least 15 sessions of training, and in four, the last five sessions median latency was within ±10 % of the median latency across those sessions. Latency was used as the training criterion because there are no errors in the sense of incorrect choices in this procedure. Hence, criteria based on correct performance were not possible. Windsor met this criterion in 23 sessions. Winston met this criterion in 26 sessions. Windsor was continued for three additional sessions to equate training between subjects.

Pair-wise tests

Following acquisition, pair-wise testing began. Each pair-wise test session contained 40 trials. The first 10 trials in each session were composed of the same chains that were presented during acquisition. During all chain trials, reinforcement was delivered according to the same schedule that was in effect during baseline training. The remaining 30 trials consisted of 20 chain trials and 10 pair-wise tests. Each of the 10 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. Reinforcement was scheduled randomly following choices with the same probabilities used during chains to

maintain the implicit nature of the task. A subject's particular choice did not determine whether reinforcement would be delivered. There were 12 pair-wise test sessions.

Random test

For all experiments, 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 with a probability of .16 following an element, 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 the random test serves as global assessment of whether learning in any form occurred, we present the results of this test before the results of the pair-wise tests, which are more fine-grained assays of what was learned.

Data analyses

For each experiment, the principal dependent measure was latency. Choices made during the pair-wise tests constituted a second measure. Analysis of latencies focused on differences between baseline and random trials in the random test session. As is detailed in the results of the individual experiments, latencies failed to meet homogeneity of variance criteria for correlated samples using the Pitman–Morgan test (Kenny 1953) across all experiments combined and for each experiment individually. Therefore, all latency analyses for each experiment were accomplished using logged data. The logged latencies also failed Pitman–Morgan test for all experiments combined and for Experiments 1 and 3 individually. As a result, tests of between-condition differences in each experiment used the nonparametric Wilcoxon signed ranks test and results are reported in terms of the Z distribution. All pair-wise test results were evaluated using the sign test. For all analyses, an alpha level of .05 was used.

The analysis of all baseline trials combined and all random trials combined across the three experiments indicated strong positive skewness in both types of trials (baseline trials skewness = 5.25, SE = .12; random trials skewness = 14.4, SE = .12, using the Skew function in Microsoft Excel 2007). As a result, medians were used throughout to represent average performance. The use of medians as the measure of central tendency encouraged the use of median-based measures of dispersion as well. We report two such measures that are each more robust, in the sense of being less influenced by extreme values, than their mean-based counterparts. We adopted a median-based

measure to replace the standard deviation, the Median Absolute Deviation (MAD), calculated as the median of the absolute value of each latency, x_i, minus the median[1]:

$$\mathrm{Median}(|x_i - \mathrm{median}(x_i)|) \tag{1}$$

In place of the 95 % confidence interval measure that is mean based, we adopted a similar measure that is median based, the 95 % confidence interval for the median (CIM). CIM is calculated by rank-ordering all latencies for a given treatment and then applying the following formula to determine the rank of the relevant latency[2]:

$$n/2 + / - 1.96\sqrt{n/2} \tag{2}$$

Results and discussion

Baseline training

Both subjects completed 40 trials per session over the last three baseline sessions and during all test sessions. There were significant differences between the first three sessions and the last three sessions in latency for both subjects combined, measured as the time to complete a sequence (first three sessions: Mdn = 7.60 s, MAD = 3.22 s; last three sessions: Mdn = 3.30 s, MAD = 1.04 s, Z = 3.75, p = .00), and for each subject individually (Windsor: first three sessions–Mdn = 11.27 s, MAD = 5.26 s; last three sessions–Mdn = 3.28 s, MAD = 1.38 s, Z = 2.83, p = .00; Winston: first three sessions–Mdn = 6.44 s, MAD = 1.70 s; last three sessions–Mdn = 3.45 s, MAD = 0.74 s, Z = 2.27, p = .02).

Random test

The results of Pitman–Morgan test indicated that the combined data for both subjects failed to meet homogeneity of variance either for raw or logged data (MAD baseline = .45 s, MAD random = 1.10 s; Pitman–Morgan test (raw data): t(166) = 40.74, p = .00; Pitman–Morgan test(logged data): t(166) = 3.65, p = .00).

Figure 1 presents a histogram of latencies for each subject during the baseline and random portions of the random test session. For all figures, the median is included as the measure of central tendency. MAD and CIM are given as measures of dispersion. The last bin in each figure is a catch-all bin for latencies that exceeded 8.0 s. All latencies were included in the analysis except for latencies that immediately followed reinforcement.

---

[1] Formula was taken from: http://en.wikipedia.org/wiki/Median_absolute_deviation, downloaded 1.2.13.

[2] Formula was obtained from the Statistics and Research Methodology website: https://epilab.ich.ucl.ac.uk/coursematerial/statistics/non_parametric/confidence_interval.html, downloaded 12.28.12.
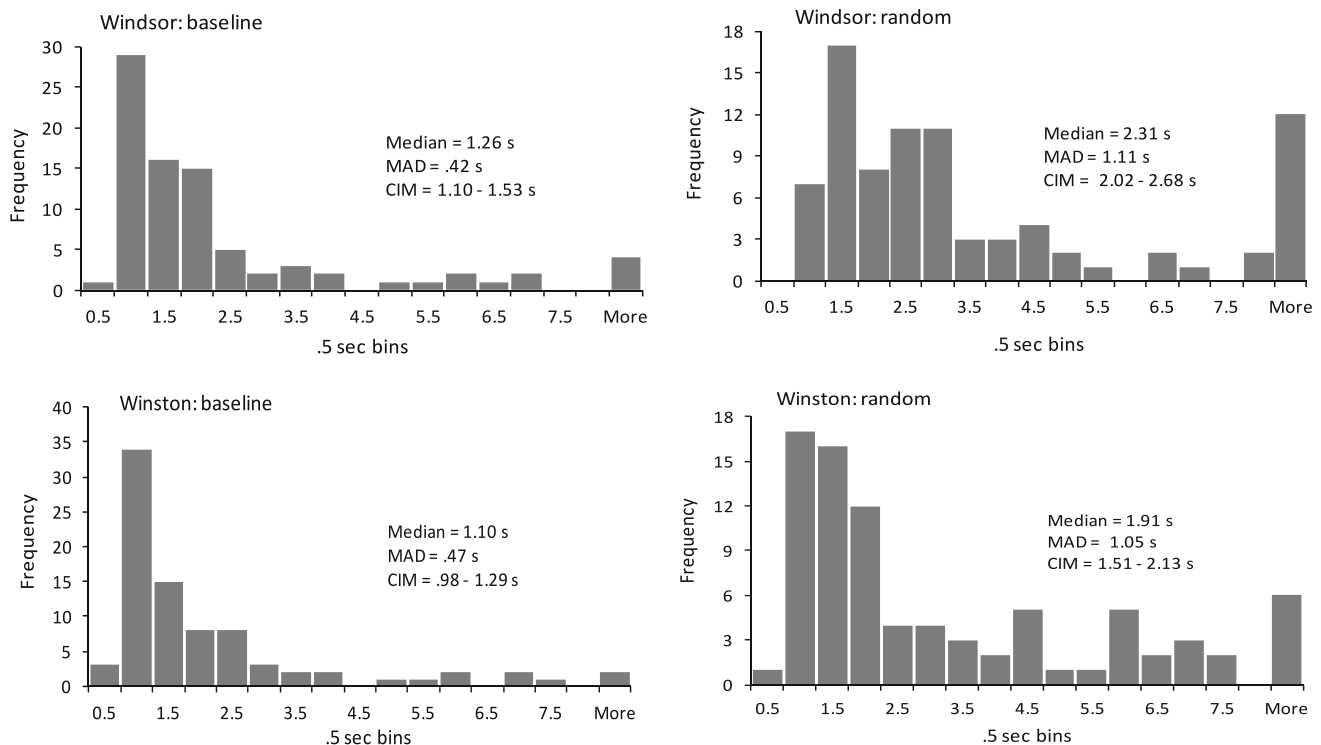
Fig. 1 Histograms of latencies for subjects in Experiment 1 during baseline and the random test. Given the positive skewness of latency distributions, descriptive indices used the median as the measure of central tendency and used median-based measures of dispersion, the Median Absolute Deviation (MAD), and the 95 % confidence interval for the median (CIM). The last bin, termed "More," contains all latencies greater than 8.0 s. Baseline latencies were taken during the 20 trials preceding the 20 random test trials

There were significant differences in latencies between baseline and random conditions averaged across the two subjects and for each subject individually (average: $Z = 5.41$, $p = .00$; Windsor: $Z = 4.28$, $p = .00$; Winston: $Z = 3.45$, $p = .01$). Overall, there was approximately a 78 % increase in latencies between baseline and random conditions. The pattern of latency changes between the baseline and random test portions of the session for each subject included an increase in median latency, an increase in MAD, an increase in the range of CIM, and the presence of more frequent extreme latencies during the random test, that is, latencies in the catch-all bin of latencies that exceeded 8.0 s. For Windsor, there were 4 such latencies during baseline and 12 such latencies during the random test. For Winston, there were 2 such latencies during baseline and 6 during the random test.

Pair-wise tests

As earlier noted the pair-wise tests were used to determine whether the tamarins learned something about the ordinal position of the elements. We evaluated this issue in two ways. First, we asked whether subjects demonstrated a tendency to choose either the earlier or later element in a pair, across all pairs. This analysis indicated that there was no tendency to do so, for either tamarin alone or for the two subjects combined ($p > .05$ for all comparisons using the sign test). For Windsor, 62 out of 120 pair-wise choices were for the earlier element in a pair. For Winston, 70 out of 120 choices were for the earlier element in a pair. Second, we examined the choices that involved only element E. We conducted this analysis given the fact that element E was contiguous with the ITI, and the ITI might be considered an inhibitory stimulus, as might be the 5-s ready signal that followed the ITI. This relationship poses a type of artifact in this procedure in that while the procedure was designed to equate the probability of reinforcement across elements, there is a distinct inhibitory period which is reliably preceded by element E. This second analysis using the sign test revealed that earlier elements were chosen significantly more than was element E across all pair-wise tests that involved element E. This result was significant for both tamarins combined ($p = .04$), and individually for Winston ($p = .00$) but not for Windsor ($p = .44$). For Windsor, 23 out of 48 choices were for an element other than E, while for Winston 34 out of 48 choices were for an element other than E.

The results of the pair-wise tests involving element E suggest that the implicit learning observed in this experiment may have been influenced to some extent by

differences in the temporal relations that each element bore to the ITI. It might be argued that this influence was relatively weak, given that all pair-wise test results combined did not indicate a systematic choice of elements earlier in the chain over those elements that were closer temporally to the ITI (e.g., B was not chosen more than D, although D bears a closer forward-paired temporal relationship to the ITI). These observations are admittedly speculative. The issue is best addressed experimentally, by designing an implicit sequence task that does not contain the potential confound of the elements in the sequence bearing different forward temporal relations to an inhibitory interval.

Experiment 2

The strategy used in Experiment 2 was to follow the termination of each stimulus with a 1-s interval (interstimulus interval or ISI). The ITI between the termination of element E and element A, the beginning of the next sequence, was also set at 1 s. Reinforcement was not programmed to occur following either the ISIs or the ITI. This strategy equated the temporal intervals that followed the end of each element in the sequence. In this procedure, then, the A → E sequence repeated continuously, without a discernable temporal break between sequences, much in the manner that the sequence repeated in the procedure used by Nissen and Bullemer.

Methods

Subjects

Windsor served in this experiment along with two other males, Fergus and Homer. Windsor was 20 years old at the start of the study while Fergus was 8 years old and Homer was 7 years old. Both Fergus and Homer had served previously in implicit learning experiments. Fergus had participated in one of the experiments reported in Locurto et al. (2010). Homer had participated in an unpublished pilot study in which he learned a series of three-item lists implicitly. The subjects were treated the same as the subjects in Experiment 1 with respect to daily care and their food deprivation schedule.

Apparatus

The apparatus was the same as the apparatus used in Experiment 1 with the exception that due to difficulties with the rotary reinforcement delivery system, we changed the reinforcer to one 97-mg pellet as compared to the two 45-mg pellets that were used in Experiment 1. The stimulus consisted of the image of the Chinese symbol for justice. The image was approximately 200×200 pixels, and it consisted of black lettering displayed against a white background. As was true in Experiment 1, the image could be displayed in one of six spatial locations on the touch screen.

Procedure

The general procedures including the baseline training criterion were the same as used in Experiment 1. The one change in procedure consisted of arranging for equal temporal intervals between all elements of the chain, including the interval between elements E and A, the ITI. This change was accomplished by inserting a 1-s delay following a response to an element and also using a 1-s ITI. The 1-s delays between elements consisted of a white touch screen, whereas the ITI consisted of a dark touch screen. Following the 1-s ITI, element A was presented and a new sequence was presented, without the 5-s ready signal that had been used in Experiment 1. The ready signal was only used prior to the first presentation of the sequence in a session. This procedure had the effect of presenting the subject with a continuous loop of the sequence, interrupted only by reinforcement. When reinforcement followed an element, the 1-s interval followed reinforcement. Each completion of the sequence was considered a trial. A comparison between the procedures used in Experiments 1 and 2 is given in Table 1.

As detailed in the Results and discussion section, with this procedure, none of the three tamarins regularly completed all 40 trials within the 50 min time limit for a session during baseline training. Extensions of session time to 75 min during the last 10 baseline sessions of training prior to the random test did not lead to the completion of 40 trials for any subject. As a consequence of the fact that the tamarins did not complete all trials programmed for a session during training, the pair-wise tests and the random test were altered from the procedure used in Experiment 1 to accommodate the possibility that subjects would not complete 40 trials per session. For the pair-wise tests, we included only two tests, B-D and A-E, each of which was presented four times a session. These two tests were chosen because they are each typically used in assessing ordinal information in explicit serial learning (e.g., Davis 1992; von Fersen et al. 1991). For Fergus and Windsor, the pair-wise tests were conducted over four sessions for a total of 16 tests of each pair. For Homer, due to scheduling constraints on our access to him, pair-wise testing was conducted over three sessions for a total of 12 tests of each pair. As was true for all pair-wise tests in these experiments, reinforcement was presented randomly following choices.

Table 1 Sequences in Experiments 1 and 2

Experiment	Sequence within trial
1	(ready signal, 5 s. white) → A → B → C → D → E → (ITI, 20 s dark) →
2	A → (1 s, white) → B → (1 s) → C → (1 s) → D → (1 s) → E → (ITI–dark; 1 s) →

We altered the random test by reducing the test session to 30 trials, 10 baseline trials, and 20 random trials. These 10 baseline trials were supplemented for analyses with the last 10 trials during the session prior to the random test, so that baseline and random blocks each consisted of 20 trials. Fergus and Homer completed the 30 trials of the random test. Windsor completed only 19 trials. We did not run Windsor for 4 weeks, then gave him six baseline sessions followed by a second random test. During those additional baseline sessions, Windsor completed all 40 trials in each session, and he completed 29 trials in the second random test. The random test results for Windsor are taken from this second random test.

Results and discussion

Baseline training

Fergus averaged 22 completed trials per session across all baseline sessions, Homer completed on average 24 trials, and Windsor completed an average of 34 trials. Fergus required 38 sessions to meet the acquisition criterion, Homer required 35 sessions, and Windsor required 34 sessions. The amount of baseline training for the three subjects combined averaged 944 trials, compared to 1,040 trials for the two subjects in Experiment 1.

There was no difference in latency between the first three sessions and the last three sessions for the three subjects combined (first three sessions: Mdn = 6.29 s, MAD = 4.22 s; last three sessions: Mdn = 5.28 s, MAD = 2.88 s, $Z = 1.11$, $p = .27$). There were significant differences for Fergus (first three sessions: Mdn = 9.36 s, MAD = 6.00 s; last three sessions: Mdn = 4.98 s, MAD = 2.86 s, $Z = 2.16$, $p = .03$), but not for Homer (first three sessions: Mdn = 10.13 s, MAD = 2.56 s; last three sessions: Mdn = 8.71 s, MAD = 2.44 s, $Z = 1.92$, $p = .06$) or for Windsor (first three sessions: Mdn = 2.26 s, MAD = .63 s; last three sessions: Mdn = 3.37 s, MAD = 1.21 s, $Z = -2.70$, $p = .02$).

Random test

The results of Pitman–Morgan test indicated that logged data for all subjects combined met homogeneity of variance, whereas raw data did not meet homogeneity of variance(MAD baseline = 1.05 s, MAD random = 1.39 s;

Pitman–Morgan test(raw data): $t(235) = 3.79$, $p = .00$; Pitman–Morgan test(logged data): $t(235) = .33$, $p = .74$).

Figure 2 gives the histograms of latencies during the random test session for each subject. Averaged across the three subjects, there were significant increases in latencies between baseline and random conditions (latencies: $Z = 2.03$, $p = .04$). These differences between baseline and random trials were significant for Fergus and Windsor individually, but not for Homer (Fergus: $Z = 2.18$, $p = .03$; Homer: $Z = .56$, $p = .57$; Windsor: $Z = 2.13$, $p = .03$). There was approximately a 20 % median increase in latencies between baseline and random conditions across the three subjects. The pattern of latency changes between baseline and random treatments that was noted in Experiment 1, that is, increases in median, increase in MAD and CIM, and a greater frequency of extreme latencies (i.e., >8.0 s) was observed in this experiment for Fergus and Windsor, with the exception that the range of Windsor's CIM was the same in both segments of the random test session. For Fergus, the frequency of extreme latencies went from 7 to 12 across baseline and random test segments of the random test session. For Windsor, the frequency of extreme latencies went from 1 to 8 during these segments of the random test session. For Homer, there was no evidence of systematic changes between baseline and random test. He evidenced small increases in MAD and CIM between baseline and random test, but there was no evidence of increases in median or increases in the frequency of extreme latencies.

Pair-wise tests

Analysis of all choices for the three subjects combined did not reveal a preference for earlier or later elements in the two pair-wise comparisons, B/D and A/E ($p = .34$). This same result was observed for each subject individually ($p > .05$ for each pair-wise comparison for each subject).

These results indicate that when elements in a sequence are equated for associative strength both in terms of the probability of reinforcement and the relations between individual elements and the nonreinforced temporal intervals that follow them, implicit learning of the sequence occurs. While implicit learning was observed, several aspects of these results suggest that the magnitude of learning was weaker in this experiment compared to Experiment 1. First, whereas both subjects in Experiment 1 produced significant reductions in latency during baseline
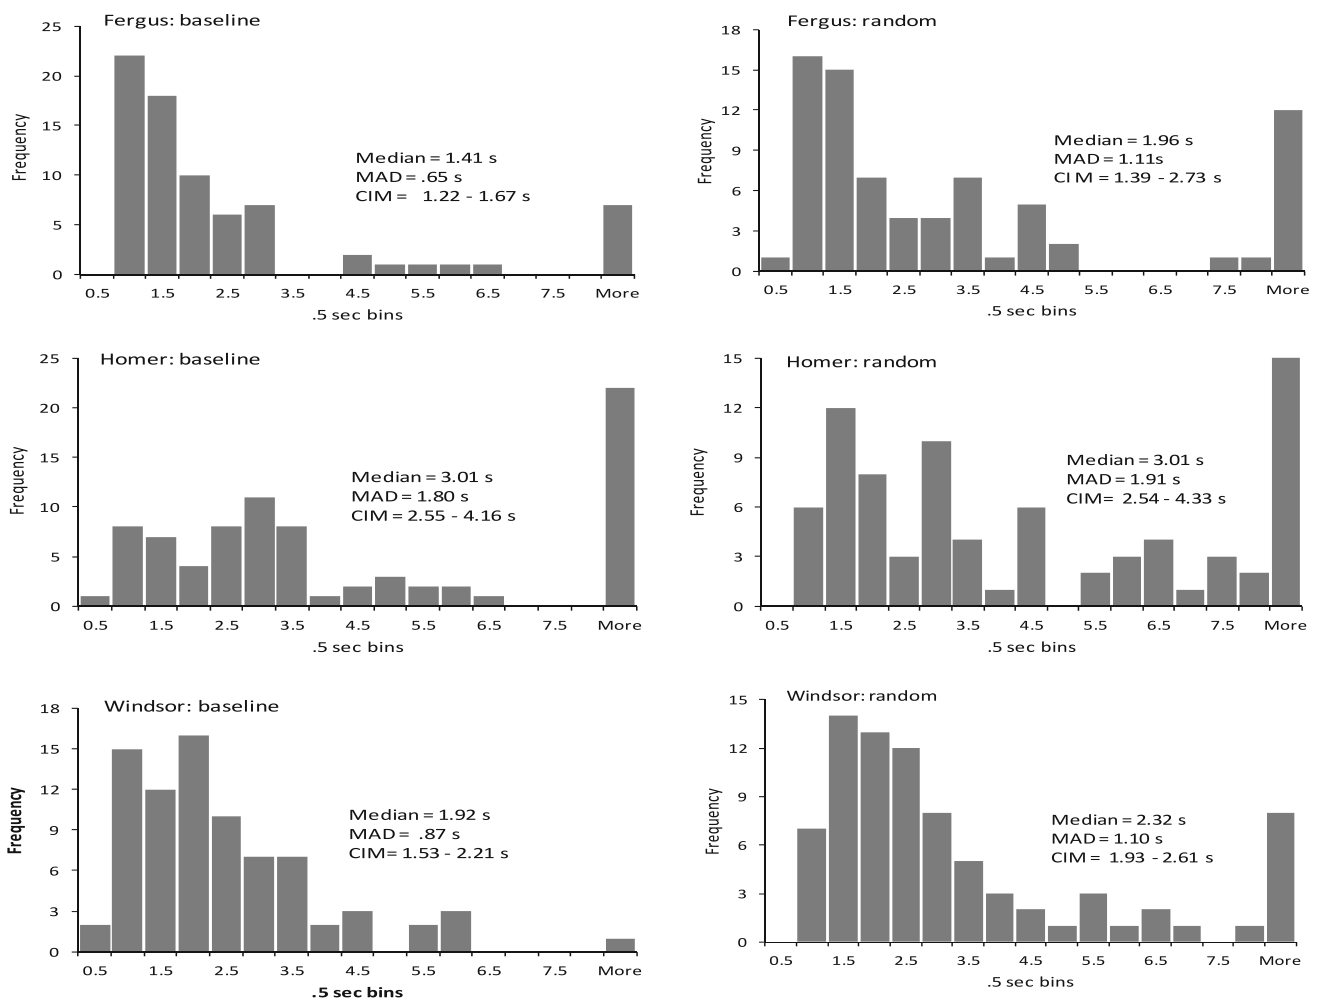
Fig. 2 Histograms of latencies for subjects in Experiment 2 during baseline and the random test. Given the positive skewness of latency distributions, descriptive indices used the median as the measure of central tendency and used median-based measures of dispersion, the Median Absolute Deviation (MAD), and the 95 % confidence interval for the median (CIM). The last bin, termed "More," contains all latencies greater than 8.0 s. Baseline latencies were taken during the 20 trials preceding the 20 random test trials. The random test session consisted of 10 baseline trials and 20 random trials. To equate the number of elements contributing to the baseline and random analysis, element latencies were added from the end of the prior baseline session until the frequency of each element in the baseline and random conditions was equal

training, only Fergus in Experiment 2 produced significant reductions in latency. Second, increases in latencies in this experiment during the random test were more modest than the increases observed in Experiment 1. In the first experiment, the rise in latencies during the random test was 78 %, whereas in the present experiment, the rise was 20 %. Third, while subjects in this experiment completed approximately the same number of trials prior to the random test session as had subjects in Experiment 1, subjects in this experiment evidenced difficulty in completing 40 trial sessions, necessitating some alterations in the structure of training and testing sessions. Last, unlike Experiment 1, one subject, Homer, did not show evidence of implicit learning.

Regarding these differences, it should be noted that the conditions of learning with respect to temporal variables were more favorable in Experiment 1, with its longer ITI (20 s), than was the case in Experiment 2 (ITI = 1 s). A number of findings have demonstrated that over a broad range of temporal values, longer ITIs produce more rapid acquisition and higher levels of maintained responding than do shorter ITIs (e.g., Gallistel and Gibbon 2000; Gibbon 1977; Locurto 2005; Staddon et al. 1999). If so, this pattern suggests that factors that affect explicit learning, in this case temporal spacing between trials, operate similarly in implicit learning. The difference in the magnitude of latency increases during the random test between Experiments 1 and 2 may also indicate that the differences in associative strength between elements that were possible in Experiment 1 mediated some of the differences in latency increases observed.

The finding that there was no tendency to choose either earlier or later elements of a pair in the second experiment

stands in contrast to the results of Experiment 1 in which element E was chosen less frequently than would be expected by chance. In one sense, it may stand to reason that the pair-wise tests would reveal no preference between element pairs, given that this experiment was designed to equate elements in terms of associative strength, even with respect to the ITI. In another sense, this result raises the question of what the nature of learning was in this procedure. The rise in latencies during the random test indicates that implicit learning occurred, but from the results of the pair-wise tests, it appears that this learning did not include ordinal information, at least as revealed in the form of choosing one element over another. In the usual course of explicit sequence learning, subjects are differentially reinforced for choosing elements in a particular sequence during training. That differential training results in appropriate choices during pair-wise testing. In the present experiment, with its continually looping chain of elements, there was no contingency that encouraged subjects to choose one item over another in the pair-wise tests.

Experiment 3

Given the lower percent increase in latencies during the random test in Experiment 2 compared to Experiment 1, we developed another type of implicit learning procedure in which elements within a sequence were equated both with respect to associative strength and to the relationships between elements and temporal intervals, as was done in Experiment 2, but one that would retain the feature of discrete trials, as was true in Experiment 1. In Experiment 3, a finite state grammar was used to achieve this purpose. The finite state grammar used in this experiment is shown in Fig. 3. This grammar is a relatively simple system compared to other types of finite grammars (e.g., Bailey and Pothos 2008). Each of the four letters corresponds to the presentation of an icon at a specified location on a touch screen. The arrows represent grammatically correct transitions between elements. While finite state grammars do not capture the flexibility of true language grammars, they are clearly a step removed from standard sequence learning (Chomsky 1957 noted the shortcomings of finite state grammars compared to true language grammars). Using an artificial grammar to specify transitions creates a sequence that has probabilistic, not certain transitions between elements. In the grammar shown in Fig. 3, each element can be followed by one of two grammatical possibilities: element A, for example, can be followed by itself or by element D and so on for each remaining element.

Note that in this procedure, individual elements can occupy more than one serial position within a sequence and
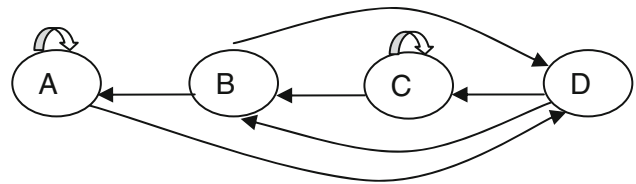

Fig. 3 The finite state grammar used in Experiment 3. In this grammar, there are two grammatical transitions and two nongrammatical transitions for each element (e.g., for element B, A and D represent grammatical transitions, whereas B and C represent nongrammatical transitions)

can occupy different serial positions across sequences. The result of this feature is that while this procedure can be used to equate associative strength across elements, it cannot be used to determine what subjects may learn about serial order, as was the case in Experiments 1 and 2.

Methods

Subjects

Two male tamarins, Fergus and Homer, served in this experiment. They were housed and treated according to the same conditions outlined for Experiments 1 and 2.

Apparatus

The apparatus was the same as that used in Experiments 1 and 2.

Procedures

We used the finite state grammar depicted in Fig. 3 to create 300 grammatically correct strings of elements. A randomly chosen sample of 40 strings was used in each session. The strings ranged from 4 to 12 elements in length. Reinforcement was set at .16/element, as was the case in Experiments 1 and 2. The ITI was set at 15 s, and the starting stimulus was set at 5 s. This procedure, therefore, presented subjects with the discrete-trial format of Experiment 1, with an ITI separating trials, while at the same time equating elements in terms of their associative strength. The construction of the strings was accomplished by choosing one element as the start element by using a random number generator and following that element for the number of grammatical transitions that equaled a number that was generated using a second random number generator set at a range of from 4 to 12 elements. Each element occupied the starting position and the ending position of the string an equal number of times (using chi Square, $p > .05$ for both starting and ending positions as judged against chance). With this procedure, therefore,

each element bore equal temporal relations to the ITI and to the 5-s starting stimulus. An issue attendant to the construction of these strings is that two of the elements in this grammar, elements A and C, can be followed by themselves, with no logical constraints on how many such repetitions can occur sequentially. To implement this grammar experimentally, we adopted the rule that elements A and C could occur up to 4 times consecutively.

Sessions were scheduled to consist of 25 trials. This amount of training nominally equated the number of elements experienced per session (average = 200 elements/session) with the nominally scheduled training in Experiments 1 and 2 in which there were 40 5-element trials scheduled (200 elements/session if all trials completed).

The pair-wise test was designed to determine what the subjects learned about the artificial grammar. For these tests, we used a procedure in which a string was interrupted at pre-determined points, and the subject was presented with a choice between a grammatical and a nongrammatical element. A grammatical element was one that constituted a valid grammatical transition from the last element in the string. So, for example, if D were the last element in the string prior to a test, the pair-wise test might consist of A and B, where A would be a nongrammatical transition and B would be a grammatical transition. Reinforcement was presented randomly following choices. Fergus completed 17 pair-wise tests spaced over six sessions. Homer completed 20 pair-wise tests spaced over five sessions.

The random test was designed to begin with five baseline trials followed by 15 random trials. To equate the number of elements contributing to the baseline and random analysis, we added element latencies from the end of the prior baseline session until the frequency of each element in the baseline and random conditions was equal. Fergus completed six random trials during the random test (50 elements). Homer completed 15 trials (122 elements).

Results and discussion

Baseline training

Fergus averaged completing 17 trials per session, while Homer averaged 23 trials per session. Both subjects met the acquisition criterion in within 20 sessions. They were both continued until they had completed 23 sessions to allow the pair-wise tests and the random test to be administered without the interruption of an academic holiday. At the end of baseline training, Fergus had completed 388 trials whereas Homer had experienced 427 trials. There were significant differences between the first three sessions and the last three sessions in latency for both subjects combined, measured as the median element latency (first three sessions: Mdn = 3.06 s, MAD = 1.12 s; last three sessions:

Mdn = 1.84 s, MAD = .69 s, $Z = 4.43$, $p = .00$) and for each subject individually (Fergus: first three sessions–Mdn = 2.27 s, MAD = .74 s; last three sessions–Mdn = 1.40 s, MAD = .35 s, $Z = 3.59$, $p = .00$; Homer: first three sessions–Mdn = 4.04 s, MAD = 1.24 s, last three sessions–Mdn = 2.31 s, MAD = .93 s, $Z = 2.67$, $p = .00$).

Random test

The results of Pitman–Morgan test indicated that the combined data for both subjects failed to meet homogeneity of variance for raw or logged data (MAD baseline = 1.02 s, MAD random = 1.72 s; Pitman–Morgan test(raw data): $t(142) = 123.01$, $p = .00$; Pitman–Morgan test(logged data): $t(142) = 3.79$, $p = .00$).

Figure 4 presents latency histograms for the baseline and random test portions of the random test session for each subject. The pattern of results that was noted for each subject in Experiment 1, and for Fergus and Windsor in Experiment 2, was also evident for each subject in this experiment. That pattern included higher medians, larger MADs, broader ranges in CIMs, and a greater frequency in extreme latencies during the random test compared to the baseline treatment. There were significant differences in latencies between baseline and random conditions averaged across the two subjects (latencies: $Z = 2.84$, $p = .00$), and there was a 43 % increase in latencies between baseline and random conditions averaged across the two subjects. These differences between baseline and random latencies were significant for Homer ($Z = 2.48$, $p = .01$), but not for Fergus ($Z = 1.43$, $p = .15$). Fergus did evidence the pattern of increased medians, MAD and CIM, along with a greater frequency of extreme latencies during the random test compared to baseline (baseline: 7, random test: 12). Homer also emitted more extreme latencies during the random test (baseline: 17; random test: 21).

These results for the random test in this experiment mirrored the results of Experiments 1 and 2. As a way of summarizing the random test results across experiments, data for all subjects in the three experiments were combined into one summary analysis. That analysis was significant (Wilcoxon, $Z = 4.75$, $p = .00$). The combined data failed Pitman–Morgan test for both raw and logged data ($p > .05$ for each measure).

Pair-wise tests

As earlier noted, pair-wise tests consisted of a choice between grammatical and nongrammatical elements. These tests immediately followed an element presentation. Fergus chose grammatical elements in 13 of 17 pair-wise tests ($p = .02$), where Homer chose grammatical elements on 7 of 20 tests ($p > .50$). The combined choices of grammatical
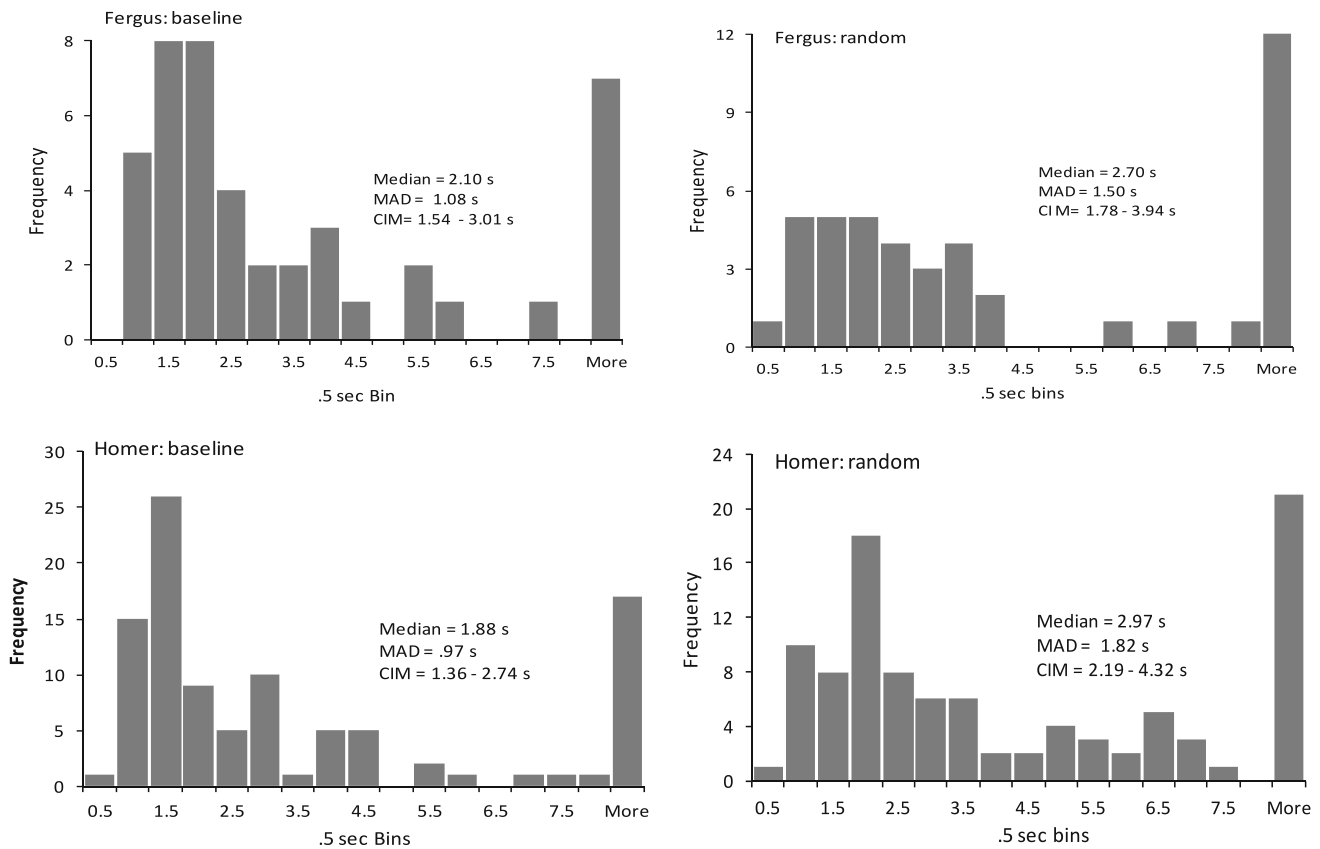
Fig. 4 Histograms of latencies for subjects in Experiment 3 during baseline and the random test. Given the positive skewness of latency distributions, descriptive indices used the median as the measure of central tendency and used median-based measures of dispersion, the Median Absolute Deviation (MAD), and the 95 % confidence interval for the median (CIM). The last bin, termed "More," contains all latencies greater than 8.0 s. The random test session consisted of five baseline trials followed by 15 random test trials. To equate the number of elements contributing to the baseline and random analysis, element latencies were added from the end of the prior baseline session until the frequency of each element in the baseline and random conditions was equal

elements for the two tamarins were not greater than chance ($p = .37$).

The results of this experiment indicate that tamarins can engage in sequence learning when transitions between elements consist of probabilistic not certain relation and when elements bear equal temporal relations to reinforcement (see Herbranson and Stanton 2011, for a similar conclusion). The pair-wise tests did not indicate that, averaged across both tamarins, they learned the difference between grammatical and nongrammatical transitions, at least as expressed by choices in these tests. These choices were significant for Fergus but not for Homer, even though Homer's latencies increased significantly during the random test.

General discussion

The results of these three experiments taken together support two conclusions: First, tamarins engaged in sequence learning whether or not there was contingent reinforcement

for learning the sequence. This finding was true for determined sequences as are typically studied in explicit serial learning (i.e., A → E) and for probabilistic sequences as are generated by a finite state grammar. Second, this learning was not due to differences in associative strength between the elements of the sequence.

One aspect of these results that requires comment is the somewhat different outcomes that were observed in the latency measures compared to those observed in the pair-wise tests. In each experiment, there was a significant increase in latency during the random test. The random test data for all three experiments combined yielded the same result. Additionally, in Experiments 1 and 3, and for one subject in Experiment 2, there were significant reductions in latency during baseline training, another indication that implicit learning took place. Yet, as judged by the pair-wise tests, there was not always clear evidence of precisely what had been learned. In Experiment 1, there was no evidence for consistently choosing elements that were either earlier or later in the chain, although there was

evidence for choosing elements other than element E in pair-wise tests involving element E. In Experiment 2, there were no reliable choices in the pair-wise tests B/D and A/E, tests that are often used to indicate the learning of serial order in explicit sequence learning(e.g., Rapp et al. 1996). In Experiment 3, there were no consistent choices of grammatical over nongrammatical elements.

In seeking to understand the differences in these two sets of results, it is useful to acknowledge that the two tests derive from different literatures, both of which are concerned with sequence learning. The pair-wise tests derive from the explicit serial learning literature. In our adaptation of this test there are no correct answers and reinforcement is delivered independently of a subject's choice. They are, nonetheless, explicit tests in that the flow of the element presentations was interrupted and subjects were asked which stimulus of two they preferred. The random test, derived from the implicit learning literature, measures latencies and is more of an embedded test in that latencies were recorded in the course of a subject's engagement in the task; the test did not interrupt the task. In explicit sequence learning, the subject is well adapted to choosing between presented stimuli during the course of training, whereas in the present set of implicit experiments, the pair-wise tests represented the first time in training that the subject was asked to choose between stimuli. Additionally, as noted above, in attempting to retain the implicit nature of the task reinforcement was delivered randomly following pair-wise test choices, not for correct responses. The combination of these factors may result in the pair-wise tests being relatively insensitive indicators of precisely what was learned in these procedures. The random test, in comparison, indicates simply whether or not learning occurred. It should be added that these experiments are not the first in which we noted that the pair-wise test results were not fully consistent with the results obtained in studies of explicit sequence learning. Locurto et al. (2009) noted that there appeared to be less precise control over internal pair choices (e.g., B/D) in an implicit chaining procedure than is typically evident in studies of explicit serial learning (compare Fig. 2, Locurto et al. 2009, with Terrace and McGonigle 1994).

The finding that tamarins engaged in sequence learning without contingent reinforcement for learning the sequence is the most salient feature of these results. This finding may have implications not only for the establishment of animal models of implicit learning, but it also may serve as a more general challenge to the nature of associative learning. Historically, the fundamental prototypes that represent associative learning have emphasized some combination of Pavlovian and/or operant relationships. In each of these prototypes, it is the presence of biologically relevant stimuli in the form of reinforcement/unconditioned stimuli that drives acquisition. Most if not all major models of learning embrace this fundamental idea, whether or not they deal directly in changes in associative strength or are approaches that rely on reinforcement to drive representational learning and timing processes (e.g., Hull-Spence, Spence 1960; Rescorla and Wagner 1972; Pearce and Hall 1980; Scalar Expectancy Theory, Gibbon 1977).

In these experiments, while there was no contingent reinforcement for correct responses, it should be pointed out that there was a consequence for responding: Each response moved the sequence forward to the next element, and in doing so reduced the temporal delay to the next reinforcement. A wide variety of evidence, summarized by Fantino's delay reduction theory, supports the idea that conditioned responding is highly sensitive to temporal delays to reinforcement (Fantino and Abarca 1985). As a corollary, the effectiveness of a stimulus as a conditioned reinforcer is best predicted by the reduction in time to primary reinforcement correlated with its onset (Fantino et al. 1993). In this sense, there was a consequence to responding, although it can also be said that in each experiment, reinforcement was delivered randomly following elements, more as an incentive to maintain performance rather than a contingently programmed stimulus to drive the learning of particular associative relationships. In this manner, a subject could adopt a simple rule: "Touch the element and reinforcement comes periodically," and could receive reinforcement at the same rate as a subject who learned the sequence (Locurto et al. 2009). It may be that this option was adopted by Homer in Experiment 2. He did not evidence a significant increase in latencies during the random test in that experiment, although he demonstrated robust implicit learning in Experiment 3. That pattern suggests that his performance in Experiment 2 was a function of the specific structure of the implicit task in that experiment rather than an inability to learn without contingent reinforcement. The adoption of this simple rule is *de facto* possible given that we have defined implicit learning in a nonhuman procedure as learning that occurs when reinforcement is not contingent on correct responding.

Despite the option of adopting this rule, subjects were clearly learning more than was strictly demanded by the experimental contingencies. One may argue that this learning was a function of a long prior experiential history of reinforcement in which subjects were contingently rewarded for the learning of sequences in other situations. Even in laboratory-reared animals, there are likely numerous opportunities to engage in sequence learning and to be differentially rewarded by that learning. An animal that learns to judge the sequential signs of an impending-scheduled feeding may be at an advantage compared to other animals.

That prior history may indeed predispose an animal to sequence learning in other situations, but it is nonetheless true that this learning was not directly reinforced in the present experiments and was accomplished under conditions that were dissimilar in many respects to the type of sequence learning engaged in by an animal in its colony room. It is likely that this type of learning also has distal roots in an evolutionary propensity to learn sequences before such learning results in positive outcomes. Consider an animal foraging in its habitat. The route to food may be uncertain. It is only after finding food that the prior sequence of turns and pathways becomes relevant. It would be extraordinarily disadvantageous if it were only at the moment of finding food that the animal began to wonder: "How did I get here?" The storing of, in this case, a spatial sequence had to occur before it could be determined that the sequence was useful. Something similar must be true of predator avoidance.

In other words, sequence learning, at least in some situations, must proceed without contingent reinforcement or unconditioned stimuli attendant to such learning. If so, perhaps the reliance on reinforcement and unconditioned stimuli in standard prototypes of learning should be questioned. A more fitting prototype may be sensory preconditioning, where a sequence of contiguous elements is experienced first (A → B), before the later element is subsequently paired with an unconditioned stimulus (B + food), thereby allowing the learning of the pre-reinforced association to be assessed (test A; e.g., Blaisdell et al. 2009; Sawa et al. 2005; see Espinet et al. 2004, for an example of inhibitory sensory preconditioning).

The domain of this type of learning is not confined to the kind of implicit procedures studied in these experiments. There are a number of nonhuman analogs of implicit learning, including not only SRT and finite state grammar procedures as used in these experiments (also, Herbranson and Shimp 2003, 2008), but in addition, latent learning (Tolman 1948), statistical learning in cotton-top tamarins (Fitch and Hauser 2004; Hauser et al. 2001) and pigeons (Froehlich et al. 2004), song bird learning, particularly recursion in song bird learning (Gentner et al. 2006; Petkov and Erich 2012), and repetition priming in pigeons (Blough 1993). A particularly interesting example of complex learning that appears to occur in the absence of reinforcement is the acquisition by birds that migrate at night of the north–south directional axis in the night sky. This learning occurs as a result of early exposure to stellar rotation, well before the birds have use for that information (Emlen 1975).

The range of these types of learning is impressively broad, encompassing as it does rather straightforward forms of conditioning such as sensory preconditioning and latent learning, as well as apparently more intricate

learning forms such as song bird learning and celestial navigation. It should be added that while the first exemplars of implicit learning emerged within the procedural motor learning literature, the examples cited above are not confined to motor learning per se. Instead, a number of them appear to require a broader interpretative framework (See, Goschke 1998; Remillard 2003; van Tilborg and Hulstijn 2010 for similar concerns in the human implicit learning literature). Perhaps most importantly, to the extent that these types of learning do not appear to depend fundamentally on the presence of contingent reinforcement or unconditioned stimuli, they represent a challenge to models of learning that do depend on those contingencies.

References

Bailey TM, Pothos EM (2008) AGL StimSelect: software for automated selection of stimuli for artificial grammar learning. Behav Res Methods 40:164–176

Blaisdell AP, Leising KJ, Stahlman W, Waldmann MR (2009) Rats distinguish between absence of events and lack of information in sensory preconditioning. Int J Comp Psychol 22:1–18

Blough D (1993) Effects on search speed of the probability of target-distracter combinations. J Exp Psychol Anim B 19:231–243

Bond AB, Kamil AC, Balda RP (2003) Social complexity and transitive inference in corvids. Anim Behav 65:479–487

Chomsky N (1957) Syntactic structures. Mouton, The Hague/Paris

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 Memory 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

D'Amato MR (1991) Comparative cognition: processing of serial order and serial pattern. In: Dachowski L, Flarherty CF (eds) Current topics in animal learning: brain, emotion and cognition. Erlbaum, Hillsdale, 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 (2004) Sequential behavior in the rat: a new model using food-reinforced instrumental behavior. Behav Brain Res 160:197–207

Emlen ST (1975) The stellar-orientation system of a migratory bird. Sci Am 233:102–111

Espinet A, González F, Balleine BW (2004) Inhibitory sensory preconditioning. Q J Exp Psychol B 57:261–272

Fantino E, Abarca N (1985) Choice, optimal foraging, and the delay-reduction hypothesis. Behav Brain Sci 8:315–330

Fantino E, Preston RA, Dunn R (1993) Delay reduction: current status. J Exp Anal Behav 60:159–169

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

Gallistel CR, Gibbon J (2000) Time, rate, and conditioning. Psychol Rev 107:289–344

Gentner TQ, Fenn KM, Margoliash D, Nusbaum HC (2006) Recursive syntactic pattern learning by songbirds. Nature 440:1204–1207

Gibbon J (1977) Scalar expectancy theory and Weber's law in animal timing. Psychol Rev 84:279–325

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, pp 401–444

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

Herbranson WT, Shimp CP (2008) Artificial grammar learning in pigeons. Learn Behav 36:116–137

Herbranson WT, Stanton GL (2011) Flexible serial response learning by pigeons (*Columba livia*) and humans (*Homo sapiens*). J Comp Psychol 125:328–340

Kenny DT (1953) Testing of differences between variances based on correlated variates. Can J Psychol 7:25–28

Locurto C (2005) Further evidence that mice learn a win-shift but not a win-stay contingency under water-escape motivation. J Comp Psychol 119:387–393

Locurto C, Gagne M, Levesque K (2009) Implicit chaining in cotton top tamarins (*Saguinus oedipus*). J Exp Psychol Anim B 35:116–122

Locurto C, Gagne M, Nutile L (2010) Characteristics of implicit chaining in cotton-top tamarins (*Saguinus oedipus*). Anim Cogn 13:617–629

Merritt DJ, Terrace HS (2011) Mechanisms of inferential order judgments in humans (Homo sapiens) and rhesus monkeys (*Macaca mulatta*). J Comp Psychol 125:227–238

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

Pearce JM, Hall G (1980) A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. Psychol Rev 87:532–552

Petkov CI, Erich DJ (2012) Birds, primates, and spoken language origins: behavioral phenotypes and neurobiological substrates. Front Evol Neurosci 4:12

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

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 University Press, New York

Remillard G (2003) Pure perceptual-based sequence learning. J Exp Psychol Learn Memory Cogn 29:581–597

Rescorla RA, Wagner AR (1972) A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black AH, Prokasy WF (eds) Classical conditioning II: current research and theory. Appleton Century Crofts, New York, pp 64–99

Sawa K, Leising KJ, Blaisdell AP (2005) Sensory preconditioning in spatial learning using a touch screen task in pigeons. J Exper Psychol Anim B 31:368–375

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

Spence KW (1960) Behavior theory and learning. Prentice-Hall, Englewood Cliffs

Staddon J, Higa J, Chelaru I (1999) Time, trace, memory. J Exp Anal Behav 71:293–330

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

Terrace HS, McGonigle B (1994) Memory and representation of serial order by children, monkeys, and pigeons. Curr DirI Psychol Sci 3:180–189

Tolman EC (1948) Cognitive maps in rats and men. Psychol Rev 55:189–208

van Tilborg IA, Hulstijn W (2010) Implicit motor learning in patients with Parkinson's and Alzheimer's disease: differences in learning abilities? Mot Control 14:344–361

von Fersen L, Wynne CD, Delius JD, Staddon JE (1991) Transitive inference formation in pigeons. J Exper Psychol Anim B 17:334–341

Yamazaki Y, Suzuki K, Inada M, Iriki A, Okanoya K (2012) Sequential learning and rule abstraction in Bengalese finches. Anim Cogn 15:369–377