



Published in final edited form as:

J Exp Psychol Anim Behav Process. 2012 October ; 38(4): 331–345. doi:10.1037/a0030306.

Cognitive mechanisms for transitive inference performance in rhesus monkeys: Measuring the influence of associative strength and inferred order

Regina Paxton Gazes,

Department of Psychology, Emory University and Yerkes National Primate Research Center, Atlanta, GA

Nicholas W. Chee, and

Program in Neuroscience and Behavioral Biology, Emory University. Atlanta, GA

Robert R. Hampton

Department of Psychology, Emory University and Yerkes National Primate Research Center, Atlanta, GA

Abstract

If Ben is taller than Emily and Emily is taller than Dina, one can accurately infer that Ben is taller than Dina. This process of inferring relations between stimuli based on shared relations with other stimuli is called transitive inference (TI). Many species solve TI tasks in which they learn pairs of overlapping stimulus discriminations (A+B-, B+C-, etc.) and are tested with non-adjacent novel test pairings (e.g. BD). When relations between stimuli are determined by reinforcement (i.e., A is reinforced when paired with B, B when paired with C), performance can be controlled by the associative values of individual stimuli or by logical inference. In Experiment 1 rhesus monkeys (*Macaca mulatta*) chose the higher ranked item on non-adjacent test trials after training on a 7-image TI task. In Experiment 2 we measured the associative values of 7 TI images and found that these values did not correlate with choice in TI tests. In Experiment 3 large experimental manipulations of the associative value of images did influence performance in some TI test pairings, but performance on other pairs was consistent with the implied order. In Experiment 4 monkeys linked two previously learned 7-item lists into one 14 item list after training with a single linking pair. Linking cannot be explained by associative values. Associative value can control choice in TI tests in at least some extreme circumstances. Implied order better explains most TI choices in monkeys, and is a more viable mechanism for TI in social dominance hierarchies, which has been observed in birds and fish.

Keywords

associative value; list linking; order; symbolic distance effect; spatial representation

Knowing that Ben is taller than Emily and that Emily is taller than Dina, we readily infer without direct comparison that Ben is taller than Dina. This is known as transitive inference (TI), the process of inferring the relation between two items based on their shared relation with a third item. Transitive inference is a prototypically cognitive process thought to emerge late in development in humans (Bryant & Trabasso, 1971; Piaget, 1960). It can be used to

correctly determine relations among any items along linearly ordered continua, such as height, mass, and linear social dominance (Paz-y-Miño et al., 2004).

In typical laboratory tests of transitive inference subjects are trained on a set of overlapping two-choice conditional discriminations ($A+B-$, $B+C-$, $C+D-$, $D+E-$, $E+F-$, $F+G-$) consistent with an implicit order ($A>B>C>D>E>F>G$). After subjects master individual premise pairs, inference is evaluated in tests with never before seen non-adjacent pairings (e.g. BD). The most strict tests of transitive inference exclude the end anchor items because these items are either always (A), or never (G), reinforced in training. The internal non-adjacent pairs such as BD have more complex reinforcement histories because they consist of items that were both reinforced (B when presented with C ; D when presented with E) and non-reinforced (B when presented with A ; D when presented with C) in premise pair training. Larger image sets are preferable in studies of TI because they provide more of these critical non-adjacent internal pairs. A set of 5 images provides only one critical test pair (BD), while a set of 7 images provides 6 critical test pairs (BD , BE , BF , CE , CF , DF). Four year old children and a diverse group of animal species perform above chance on these critical test trials, consistent with use of transitive inference (corvids: Bond et al., 2003; Bond et al., 2010; children: Bryant & Trabasso, 1971; rats: Davis, 1992; chimpanzees: Gillan, 1981; crows: Lazareva et al., 2004; squirrel monkeys: McGonigle & Chalmers, 1977; Merritt & Terrace 2011; Rapp et al. 1996; mice: Van der Jeugd et al., 2009; pigeons: von Fersen, et al., 1991; geese: Wei et al., 2010; but see Benard & Giurfa, 2004 for an exception in honeybees).

Transitive inference requires that items be processed as a ranked set in which there are no circular relationships (i.e. $A>B$, $B>C$, and $A>C$, not $C>A$). Differences in relative position along a linearly organized continuum result in a Symbolic Distance Effect (SDE) such that widely separated items are easier to rank correctly than are less widely spaced items (e.g. BF tests are easier than BD tests). Humans show the SDE in tests with ranked items, such as size and height (Moyer & Landauer, 1967; Woocher et al., 1978) and humans and non-humans also show the SDE in inference experiments, suggesting that they rank TI items along an ordered continuum (Bond, et al., 2003; D'Amato, 1991; Maclean et al., 2008; Merritt & Terrace, 2011; Woocher, et al., 1978).

Logical inference is not the only cognitive process that could give rise to the patterns of performance observed in many studies of transitive inference. Both successful performance with non-adjacent internal test pairs and the SDE may be accounted for by other processes. For example, choice behavior may be controlled by the associative values of individual stimuli in some cases (Siemann, Delius, Dombrowski, et al., 1996; von Fersen, et al., 1991; Wynne, 1998). In premise pair training with nonhuman animals, one item in a pair is reinforced with food, while the other is not. Therefore inference-like patterns of performance could emerge from the resultant variation in associative values of individual stimuli (Lazareva & Wasserman, 2006; Siemann, Delius, Dombrowski, et al., 1996; Steirn et al., 1995; von Fersen, et al., 1991; Wynne, 1998). The integral role played by reinforcement in these tasks makes it difficult to distinguish between the contributions of associative values and inference. Differences in premise pair performance during training results in differential reinforcement of the items in a TI list that could lead to differences in associative values of those items (Siemann, Delius, Dombrowski, et al., 1996, von Fersen, et al., 1991; Wynne, 1998, Lazareva & Wasserman 2006). Additionally, the associative value of one item can transfer to other items when they are presented together (von Fersen, et al., 1991; Zentall, T. R., & Sherburne, L. M., 1998). Through these mechanisms, the associative values accrued to individual TI stimuli may follow the same order as the implied by TI.

Mathematical modeling studies have shown that it is theoretically possible for associative values to generate transitive inference like patterns of performance (Siemann, Delius, Dombrowski, et al., 1996; Steirn et al., 1995; von Fersen, et al., 1991; Wynne, 1998), but empirical studies are required to directly test predictions based on associative value. One method used to differentiate between mechanisms is to manipulate the associative values of individual stimuli so that associative mechanisms and inference mechanisms would produce different patterns of choice. When the DE premise pair is over-trained prior to TI tests, potentially increasing the associative value accrued to item D, pigeons (*Columbia livia*) continued to correctly select item B on the BD test pair (Lazareva & Wasserman, 2006). Crows (*Corvus cornix L.*) that received similar training fell to chance (Lazareva, et al., 2004) in these tests, suggesting that crows but not pigeons relied on associative values to solve the task. Resistance-to-extinction measures showed no systematic differences in associative value between B and D for pigeons after overtraining, but resistance-to-reinforcement measurements indicated an increase in the value of D compared to B for most birds (Lazareva & Wasserman, 2012). However, discrepancy models of associative learning like the Rescorla-Wagner model are based on the assumption that once performance has reached asymptote, the only way to dramatically increase associative value is to increase the value of the reinforcer (Rescorla & Wagner, 1972).

Another method for dissociating the contributions of associative value and inference to TI task performance is to create a modified task that can only be solved using one of these methods. One such task is to require linking of two separate ordered lists (e.g. A>B>C; X>Y>Z) into a single larger ordered list by training only a single linking pair (C>X). Test trials consist of never before seen pairs of items, one from each of the original lists. If subjects used inference to create a single ordered representation of the two previously separate lists, then they will correctly choose the item from the higher ranked list in these probe tests. By contrast, if choice is driven by associative values, subjects would be unable to correctly select the higher ranked item on between list pairs. This is because items occupying the same position in the two lists would have acquired similar associative values during initial training. None of the current associative models predict above chance performance on linked lists (Lazareva, 2012). Monkeys taught five item lists can link two and three of these lists together after training on single linking pairs (Treichler & Raghanti, 2010; Treichler et al., 2003; Treichler & Van Tilburg, 1996). There are no published reports of list linking by other species. However fish, jays, and chickens can solve social TI tasks, indicating knowledge of inferred dominance information (Grosenick et al., 2007; Hogue et al., 1996; Paz-y-Miño, et al., 2004). As these tasks do not use explicit reinforcement, this performance is difficult to account for using associative values alone, suggesting that these animals use the implied order of stimuli to make choices.

In the present series of experiments we determined the contributions of inference and associative values to performance on TI tasks in monkeys. In Experiment 1 we documented performance patterns on a 7 item TI task. In Experiment 2 we measured the associative values of individual stimuli to evaluate whether these values predicted performance on test trials. In Experiment 3 we manipulated the associative values of multiple stimuli in a TI task. In Experiment 4 we presented monkeys with a list linking task that cannot be solved by associative value.

Experiment 1- Transitive Inference

In Experiment 1 we established baseline patterns of performance of rhesus monkeys in computerized 7-item transitive inference tests.

Method

Subjects—Subjects were twelve four to six-year-old male rhesus monkeys (*Macaca mulatta*) that had been raised by their biological mothers in a large social group until the age of approximately 2.5 years before moving to the laboratory. Monkeys were pair-housed whenever possible and kept on a 12:12 light:dark cycle with light onset at 7:00 am. They received a full ration of food daily and water was available ad libitum. Monkeys had previous experience with cognitive testing, but no experience with TI tasks.

Apparatus and procedure—Monkeys were tested in their home cages. Computerized touch-screen test systems, each consisting of a 15-inch LCD color monitor (3M, St. Paul, MN) running at a resolution of 1024 X 768 pixels, generic stereo speakers, two automated food dispensers (Med Associates Inc., St. Albans, VT), and two food cups below the screen, were attached to the front of each monkey's cage. Sessions were conducted daily between 10 am and 5 pm, six days per week.

During testing, each pair of monkeys was separated by an opaque plastic divider with holes in it that allowed visual, auditory, and tactile contact, but prevented the monkeys from touching the computer screen in the adjacent cage. Computer screens were locked to the front of each monkey's cage and the door was raised, giving subjects full access to the screen during testing. To prevent counting spurious touches as responses, all choices required two consecutive touches within the image border of a single stimulus (FR2). After a 3-second inter-trial interval (ITI), a green square appeared at the bottom of the screen and remained until the monkey touched it (FR2) to start a trial. Two images from the training set appeared on the right and left sides of the screen (position was counterbalanced over trials), and remained until one was touched (FR2). Selection of the correct item always resulted in a 1 second long positive auditory reinforcer that had a history of being associated with food reinforcement; a food reward was also delivered on 75% of trials (85% of food rewards were nutritionally balanced banana flavored pellets; Bio-Serv, Frenchtown, NJ and a random 15% of food rewards were miniature chocolate candies). Selection of the incorrect item in the pair resulted in a .5 second long negative auditory stimulus and a five second time out during which the screen was black.

Premise pair training—Stimuli consisted of two sets of seven 350 X 350 pixel color clip art images shown in Figure 1, presented in overlapping adjacent pairs (A+B-, B+C-, etc) that could be organized into an implied linear hierarchy (A>B>C>D>E>F>G). Images appeared semi-randomly counterbalanced on the left and right sides of the screen.

Training proceeded one premise pair at a time, with pairs at the bottom of the implied hierarchy (FG) trained first (Treichler & Van Tilburg, 1996). Each premise pair was introduced individually in 25 trial sessions until subjects reached at least 80% correct in a single session. Except in the case of the first pair learned, twenty-five trials of that pair were then pseudo-randomly intermixed in a session including 25 trials of each of the previously learned pairs until subjects performed above 80% on each pair present in one session simultaneously. This pattern continued as indicated in Table 1 until monkeys met criterion with all 6 training pairs in a single session (phase 11, Table 1).

Transitive inference test trials—One trial with each of the 15 non-adjacent test pairings (e.g. AF, BD, CE, etc.) was pseudo-randomly intermixed with 25 trials of each of the 6 premise pairs to generate a session of 165 trials (15 TI test pairs and 150 premise pairs). All TI test trials were reinforced with the positive auditory reinforcer only, regardless of whether responses were correct with respect to implied order or not. Auditory only reinforcement was consistent with the pattern of reinforcement monkeys had come to expect

for correct responses, which included 25% auditory only reinforcement trials. Reinforcing probe trials non-differentially this way encouraged monkeys to continue executing whatever rules they had adopted in training. Subjects received 4 testing sessions.

The entire training and testing procedure was completed twice for each subject with two distinct sets of stimuli (Run1 and Run2). The order of training with the two image sets was counterbalanced across subjects.

Data analysis—All response latency analyses in this paper used median latencies from correct trials only (Montgomery, 1953). All accuracy data were arcsin transformed before analyses (Aron & Aron, 1999). Accuracy refers to percent correct with respect to the implied TI order. All analyses were conducted using a two-tailed alpha level of .05, except where otherwise indicated.

Results and Discussion

Performance with the two image sets did not differ either in the total number of errors made before reaching criterion (independent samples t-tests: Run 1: $M_{set1}=673$ errors, $SEM=144$, $M_{set2}=750$ errors, $SEM=63$; $t_{10}=-0.55$, $p=.60$; Run 2: $M_{set1}=530$ errors, $SEM=13$, $M_{set2}=626$ errors, $SEM=50$; $t_{10}=-0.77$, $p=.46$) or in accuracy on internal test pairs (independent samples t-tests; Run 1: $M_{set1}=69.9\%$, $SEM=0.4$, $M_{set2}=69.5\%$, $SEM=0.2$; $t_{10}=0.05$, $p=.96$; Run 2: $M_{set1}=80.8\%$, $SEM=0.6$, $M_{set2}=70.0\%$, $SEM=0.5$; $t_{10}=1.18$, $p=.27$). The two image sets were therefore combined for further analyses.

Monkeys made more total errors before reaching criterion in the first run of the experiment than in the second run of the experiment ($M_{run1}=718$, $SEM=67.2$, $M_{run2}=586$, $SEM=60.6$; paired samples t-test: $t_{11}=4.07$, $p=.002$). There was no difference between the two runs of the experiment in performance on the six adjacent training pairs during the criterion session (RMANOVA: $F_{5, 55}=0.34$, $p=.88$), or performance on the critical non-adjacent internal test pairs ($M_{run1}=69.6$, $SEM=0.1$, $M_{run2}=74.7$, $SEM=0.3$; paired samples t-test: $t_{11}=-0.82$, $p=.43$), therefore all test data were combined across the two runs for further analysis.

Subjects performed significantly above chance on critical non-adjacent internal test pairs ($M = 71.6$, $SEM=0.1$; $t_{11}=7.95$, $p<.001$). These results reinforce previous findings that rhesus macaques and other species solve 5 and 7 item transitive inference tasks (corvids: Bond et al., 2010; chimpanzee: Gillan, 1981; crows: Lazareva, et al., 2004; rhesus monkeys: Buckmaster et al., 2004; Merritt & Terrace, 2011; Rapp, et al., 1996; Treichler & Raghanti, 2010; Treichler et al., 2007; Treichler & Van Tilburg, 1996; pigeons: Von Fersen et al., 1991; greylag geese: Wei, et al., 2010).

To test for the SDE all test pairs were grouped according to the number of list images intervening between members of the pair. For example, a symbolic distance of 1 included test pairs AC, BD, CE, DF, and EG, while a distance of 2 included AD, BE, CF, and DG. Response latency decreased and accuracy increased as the symbolic distance between tested images increased, as shown in Figure 2 (RMANOVA Response Latency: $F_{4, 44}=14.51$, $p<.01$; RMANOVA Accuracy: $F_{4, 44}=52.75$, $p<.01$). When pairs containing end items were excluded from the SDE analysis, leaving symbolic distance 1 represented by pairs BD, CE, DF; distance 2 by pairs BE, CF; and distance 3 by pair BF, accuracy increased but response latency did not change as the distance between tested images increased (RMANOVA Accuracy: $F_{2, 22}=4.35$, $p=.03$; RMANOVA Response Latency: $F_{2, 22}=0.23$, $p=.98$).

At larger symbolic distances the proportion of contributing pairs that contain end anchor images increases (i.e. distance 5 has 100% of pairs containing images A and G while distance 1 has 40% of pairs including an end anchor). Because end anchors are either always

or never reinforced, these pairings may be easier to solve than others and may lead to overestimates of the SDE (Vasconcelos, 2008). We examined the SDE by the first item in the pair so that only one pair contributes to performance at each distance. As shown graphically in Figure 3, and statistically in Table 2, the SDE patterns still hold across different first images indicating that the overall effect is not driven entirely by the disproportionate contributions of end anchor pairs to the larger symbolic distances. This suggests that monkeys represent the implied order of all images.

Monkeys performed above chance on internal test pairs in two seven item lists and showed a robust SDE for both accuracy and latency. At least two mechanisms can potentially explain these patterns. Monkeys might use logical inference to form an ordered representation of stimuli or choice behavior might be controlled by associative value differences that mirror the implied rank order of the images. In Experiment 2 we empirically measured the associative values of individual stimuli after premise pair training to determine whether performance on standard TI tasks can be accounted for by these values.

Experiment 2- Measurement of Associative Values

Associative value accounts of TI performance posit that items used in TI tests gain associative values consistent with the implied order through association with primary reinforcers during premise pair training. In TI tests, differences in associative value manifest as preference for particular stimuli on non-adjacent test trials. Modeling studies have illustrated how TI stimuli could accrue associative values that produce above chance performance on test trials (Siemann & Delius, 1998; Siemann, Delius, & Wright, 1996; von Fersen et al., 1991). However the only study to empirically measure the values accrued to stimuli used in transitive inference tests found that associative values indexed by resistance to extinction and resistance-to-reinforcement do not predict performance on test pairs in pigeons (Lazareva & Wasserman, 2012). In order to determine the extent to which associative values could produce the same patterns of performance expected from inference on standard TI tasks in monkeys, we explicitly measured the associative values accrued to the stimuli used in Experiment 1.

We presented monkeys with two identical concurrent schedules; one schedule was associated with neutral images and the other with the seven trained TI images. This allowed us index the relative associative value of each of the seven TI stimuli. Secondary reinforcers superimposed on stable schedules of reinforcement facilitate responding, permitting comparisons of the associative values of secondary reinforcers (Armus & Garlich, 1961; D'Amato & Lachman, 1958). Because we used two identical RI schedules, differences in rates of responding between the schedules can be attributed to differences in the associative values of the stimuli (Miller, 1976). The relative response rate to each of the seven transitive inference stimuli was used to index associative values.

To the extent that monkeys' choices in the TI tests were determined by the associative values of individual stimuli, they should select the stimulus with the higher associative value in each test pair regardless of whether that stimulus was higher or lower in the implied hierarchy. Accordingly, larger differences in the associative values of the images in a test pair should result in greater preference for the higher valued image in TI tests. We determined the extent to which choice was controlled by associative value by assessing the extent to which choice behavior on TI tests correlated with associative value.

Method

Subjects and apparatus—Subjects and apparatus were the same as in Experiment 1.

Stimuli and procedure—Monkeys were trained on two concurrent random interval schedules for at least one hundred 25 minute sessions by which point responding was expected to be stable. The two schedules were represented by 300×300 pixel images presented simultaneously on the left and right sides of the screen. Images for the two schedules were drawn from a pool of 20 familiar clip art images and changed in synchrony every 30 seconds, independent of the state of either schedule. The schedules operated independent of one another, and the first touch after each random interval was reinforced with one food pellet and an auditory reinforcer. All other touches were recorded, but did not result in reinforcement. There was no changeover delay. Each session lasted 25 minutes.

The TI images for which associative values were to be measured were presented as probe trials pseudorandomly intermixed with the 20 familiar images used during concurrent RI training. A probe trial began when the images at both schedule locations changed. The right hand location displayed a different one of the 20 familiar clip art images, while the left location changed to be one of the to-be-measured probe stimuli. Probe images appeared only on the left side of the screen to control for side bias. During the 30 seconds that the probe images were on the screen, neither schedule terminated and the monkeys were not rewarded for any touches. When the 30 sec probe period was over, the images each changed to one of the 20 standard images, and a shortened RI 10s schedule was initiated on both keys. Consecutive probe trials were separated by at least 90 seconds. To increase the likelihood that subjects were attending to the screen when the probe images appeared, probes were only presented after subjects had touched each key at least once since the presentation of the last probe.

Object discrimination reversal—To evaluate the efficacy of our methods for measuring associative value, we first measured the accrual of associative value during training on a series of object discriminations and reversals. Monkeys learned 4 two-image object discriminations concurrently. One of the images in each pair was randomly designated the S+ and was reinforced, while the other was the S-. Monkeys received three 8 trial sessions per day (two trials of each of the 4 object discriminations) until they chose correctly at least 23 of the 24 trials in a single day. The contingencies of each discrimination were then reversed so that choice of the item that had been rewarded was now non-rewarded (S-) and choice of the item that had been non-rewarded was now rewarded (S+). Monkeys were trained to the same criterion on this reversal, followed by a final reversal in which the reward contingencies were the same as in the original training. Immediately after each eight trial object discrimination training session one 25 minute concurrent RI30s test session was conducted in which each stimulus from the object discrimination pairs was presented once as a probe.

Transitive inference—Following completion of the object discrimination training and measurement, the associative values of the transitive inference stimuli used in the first iteration of Experiment 1 were measured. Monkeys received one transitive inference test session with reinforcement contingencies as described in Experiment 1, followed by one measurement session in which the probe stimuli were the seven images from the TI set. The associative value index for each probe stimulus was calculated using data from this measurement session. Three more TI test sessions were conducted, one per day over the next three days. Accuracy for each of the 15 TI test pairs was calculated by averaging performance over these four test sessions.

Data analysis—The associative value index of each probe stimulus was assessed by calculating the proportion of total touches allocated to the probe stimulus during the 30 seconds it appeared on screen during the one measurement session using the formula:

$$\text{Value Index} = \frac{\text{Touches to the Probe Stimulus}}{\text{Touches to the Probe Stimulus} + \text{Touches to the NonProbe Stimulus}}$$

Associative value scores therefore range from 0 (aversion to the probe stimulus) to 1 (exclusive preference for the probe stimulus).

For each of the 4 object discrimination pairs and each of the 15 possible non-adjacent TI test pairs we calculated an associative value difference score using the formula:

$$\text{Associative Value Difference Score} = \text{Value of Correct Item} - \text{Value of Incorrect Item}$$

Positive difference scores indicate that the correct item (the S+ in the object discrimination task and the higher ranked item in the TI task) had a larger value index than the incorrect item. Negative difference scores indicate the incorrect item had a larger value index than the correct item. If associative values of individual stimuli control choices, then the valence and magnitude of these differences should predict choice in object discrimination and TI test trials.

We used a Generalized Linear Mixed Model with subject as a random factor to compare average daily associative value difference scores with average daily choice on the object discrimination task over the initial training, first reversal, and second reversal. A Generalized Linear Mixed Model with subject as a random factor compared associative value difference scores with choice on the 15 transitive inference test pairs. Data from one subject were excluded from the transitive inference measurement analysis due to insufficient touching (this monkey touched both control and probe stimuli less than 0.1 times per second, which was our cutoff for inclusion).

Results and Discussion

Measured associative value differences correlated with object discrimination performance. As shown by the summary data in Figure 4, there was a significant relationship between associative value difference scores and percent choice of the S+ across testing days (Generalized Linear Mixed Model, daily associative value difference X daily discrimination accuracy: $F_{1,170} = 7.90$, $p = .01$). This positive relationship demonstrates that the method we used to measure associative values detected differences that are relevant to choice behavior in tasks in which associative values are expected to control choice.

Associative value models predict a significant linear relationship between the implied order of the TI images and their individual value index scores. However the value indices of the individual TI images were not significantly correlated with the implied order of those images, Figure 5 (Spearman Rank Order correlation: $r_7 = -.23$, $p = .61$), indicating that associative values accrued to individual images in standard TI tasks do not necessarily follow the implied order.

However, monkeys' choices on TI test trials may still be driven by associative value even when associative values do not correlate with the implied TI order of images. If TI test pair performance was controlled by associative values, then associative value difference scores should correlate with TI accuracy. By contrast if performance was controlled by the implied order of the TI images, accuracy should be unrelated to value differences and instead would vary with symbolic distance between the images in the test pairs (D'Amato & Colombo, 1990). Symbolic distance was not correlated with associative value difference scores

(Spearman Rank Order correlation: $r_{13}=.18$, $p=.53$). Therefore any effect of symbolic distance in these results cannot be explained by associative value differences alone.

There was no relationship between associative value difference scores and accuracy across the fifteen test pairs from the TI image set, as shown in Figure 6 (*left*) (Generalized Linear Mixed Model, associative value difference X pair accuracy: $F_{1,12.68}=.68$, $p=.42$). In contrast the correlation between accuracy and symbolic distance of the test pairs was nearly significant (Spearman Rank Order correlation: $r_{13}=.51$, $p=.05$), as shown in Figure 6 (*right*). These results show that performance on test trials in the TI task was not controlled by associative value as was the case in the object discrimination task, and suggest that choice instead was controlled by the implied order of the stimuli.

Associative value models predict that animals will select the item with the higher associative value on TI test trials. In contrast, our results support those recently reported in pigeons (Lazareva & Wasserman, 2012) that choices in the TI tests were not correlated with associative value. In fact, in TI tests, monkeys often chose the stimulus with the lower associative value (all those points on the left side of Figure 6 where associative value difference is negative and accuracy is above 50%). This indicates that when associative value and implied order conflict, choice behavior is controlled by the implied order. In Experiment 3, we further evaluated the relative influence of associative values and inference by explicitly manipulating the magnitude of the reinforcement associated with different images in a TI task.

Experiment 3- Manipulation of Associative Values

In standard transitive inference training paradigms, images high in the implied hierarchy may accrue larger associative values than images lower in the hierarchy due to differences in reinforcement and non-reinforcement during training. To the extent that this is the case, associative value is congruent with the implied order of the hierarchy, making it difficult to determine the extent to which the implied order and associative value control choice behavior. But in Experiment 2, we found that associative value and implied order were often incongruent: in six of the TI test pairs, the associative value of the image lower in the implied order had accrued a larger associative value than item at the top of the implied order. Despite this incongruency in Experiment 2, monkeys chose images consistent with the implied order more often than expected by chance, indicating supremacy of implied order over associative value under these conditions. However, it is likely that stronger differences in associative value can affect choice behavior in TI tests. In Experiment 3 we directly and dramatically manipulated the associative values of individual stimuli to more clearly test for contributions of associative values to choice in TI tests.

Other investigators have overtrained selected premise pairs in an effort to increase the associative value of a lower ranked item and create incongruencies between implied order and associative strength. For example, overtraining the DF pair could increase the value of D relative to other stimuli because it is reinforced on many additional trials. Results from such overtraining in pigeons and crows are mixed; pigeons continued to select B over D in transitivity tests, suggesting use of inference, while crows chose D and B equally often, suggesting that the associative value manipulation influenced their choices (Lazareva, et al., 2004; Lazareva & Wasserman, 2006). The fact that in neither case was D chosen over B in transitivity tests may either show either that choice is not controlled by associative value or that overtraining has only a modest effect on associative value. Because all premise pairs are trained to a high accuracy criterion using the same reinforcer, all images in the TI set may already have values close to maximum supported by the reinforcer (Rescorla & Wagner,

1972). Additional reinforced trials administered in overtraining may have only a small effect.

To produce large systematic differences in associative value among stimuli in an implied hierarchy we manipulated reward magnitude. Selection of some images during training was rewarded with a single food pellet while other images were reinforced with two pellets. When this manipulation produces associative values that are congruent with the implied order, monkeys should perform above chance whether or not their behavior is controlled by associative value. When the resultant associative values are incongruent with the implied order, accuracy should decrease to the extent that choice is controlled by associative values.

Method

Subjects and apparatus—Subjects and apparatus were the same as in Experiments 1 and 2.

Data analysis—In order to give equal weight to the associative value and implied order hypotheses, we assessed statistical tests using both corrected and Bonferonni corrected alpha levels where applicable. Corrected alpha levels are presented when used.

Procedure

Premise pair training—Stimuli were two new sets of seven 300×300 pixel photographs. Premise pair training was conducted as in Experiment 1, with the following exceptions. On the 80% of correct choices that were rewarded with a food reinforcer in the Congruent condition, correct choices of images A, B, and C were rewarded with 2 pellets, while correct choices of images D, E, and F were rewarded with one pellet (G was never correct and therefore never rewarded as before). On the 80% of correct choices that were rewarded with a food reinforcer in the Incongruent condition, correct choice of images A, B, and C were rewarded with 1 pellet, while correct choices of images D, E, and F were rewarded with 2 pellets. Incorrect choices in both conditions resulted in no food reward, a negative auditory stimulus, and a five second time out during which the screen was black. All monkeys received both conditions with order of the conditions and image sets counterbalanced across subjects.

Transitive inference test trials—Test trials were presented as in Experiment 1. Correct choices were defined as those consistent with the implied order regardless of how the images had been rewarded during training.

We assessed the influence of associative value in two ways. First, we compared the number of errors required to reach criterion in premise pair training in the Congruent and Incongruent conditions. Second, we examined performance on the critical internal non-adjacent test pairs BD, BE, BF, CE, and CF in the Congruent and Incongruent conditions. For these pairs, one image in the pair had been reinforced with a single pellet and the other image had been reinforced with two pellets. Because the implied order (A>B>C>D>E>F>G) did not differ between the two conditions, choices based on the implied order would result in above chance performance in both the Congruent and Incongruent conditions. In contrast, to the extent that choice behavior is controlled by associative value monkeys should select the item that was reinforced with two food rewards during training over the item that was reinforced with only one, even when these choices conflict with the implied order. In the Congruent condition choices based on associative value would still result in above chance performance, as the higher ranked images also had higher associative values, but in the Incongruent condition choices based on associative

value would result in below chance performance, because the lower ranked images had higher associative values.

Results and Discussion

Order of the congruent and incongruent condition presentation did not affect learning rates or accuracy within either condition, as shown by comparison of conditions as a function of whether they were trained first or second (errors to criterion, Congruent: $t_{10}= .12$, $p=.90$; Incongruent: $t_{10}=-1.20$, $p=.26$; internal test pair performance, Congruent: $t_{10}=1.40$, $p=.19$; Incongruent: $t_{10}=-1.58$, $p=.14$). Therefore data were collapsed within each condition with respect to order of testing.

Premise pair training—In both the Congruent and Incongruent conditions, reinforcement of premise pairs FG, EF, and DE is congruent with implied order and should not be difficult to learn. In the incongruent condition, reinforcement of images C, B, and A is in conflict with the implied order. To the extent that choices are controlled by associative value, acquisition of these pairs, but not others, should be retarded relative to the congruent condition. As shown in Figure 7, monkeys learned pairs FG, EF, and DE in the two conditions with similar numbers of errors (paired samples t-tests: FG: $t_{11}= 0.22$, $p=.83$; EF: $t_{11}= 0.20$, $p=.85$; DE: $t_{11}= -1.44$, $p=.18$), but made significantly more errors reaching criterion on pairs CD, BC, and AB in the Incongruent condition than the Congruent condition (2×6 RMANOVA: Condition: $F_{1,11}=11.39$, $p=.01$; Premise Pair: $F_{5,55}=9.99$, $p<.01$; Condition X Premise Pair: $F_{5,55}=8.77$, $p<.01$; paired samples t-tests: CD: $t_{11}= -5.15$, $p<.00$; BC: $t_{11}= -2.35$, $p=.04$; AB: $t_{11}= -2.82$, $p=.02$). Using a corrected alpha of .008, monkeys still made significantly more errors reaching criterion on the critical CD pair, suggesting that the incongruity of reinforcement value and implied order retarded learning.

Despite differences in how well the pairs were learned, by the end of training monkeys performed above 85% correct on all six test pairs in both conditions.

Transitive inference testing—On critical internal non-adjacent TI test trials on which the two images in a test pair had been trained with different reinforcement values (BD, BE, BF, CE, CF), monkeys performed above chance in the Congruent condition, but below chance in the Incongruent condition (Congruent: $M= 89.6\%$, $SEM=5.8$; one sample t-test: $t_{11}= 6.50$, $p<.00$; Incongruent: $M= 37.1\%$, $SEM=4.6$; one sample t-test: $t_{11}=-2.79$, $p=.02$). This indicates that when differences in associative value are sufficiently large, the influence of associative value on choice can overwhelm control by implied order. Because premise pairs in both conditions were trained to the same criterion, deficits in internal test pair performance in the Incongruent condition cannot be explained by differences in how well the pairs were learned.

Monkeys correctly chose the higher ranked item in both the Congruent and Incongruent conditions if images in the tested pair were associated with the same reward (i.e. pairs AC and DF; Congruent: $M= 89.6\%$, $SEM=4.6$, one sample t-test: $t_{11}=7.01$, $p<.00$; Incongruent: $M= 76.0\%$, $SEM=3.6$, one sample t-test: $t_{11}=5.42$, $p<.00$). Despite both images being rewarded with the same number of reinforcers, it is possible that the higher ranked item in these pairs accrued a higher associative value during training. However, the results of Experiment 2 suggest that TI choices are not driven by associative value differences of this size. Therefore, this result suggests that monkeys had latent knowledge of the inferred order of all the images, but the two-fold reinforcement difference between test images in pairs BD, BE, BF, CE, and CF resulted in control of choice by associative values, or knowledge of reward magnitude, on these trials.

This is the first study we know of in which associative value was manipulated using different reward magnitudes. The results show that when associative value differences are large, they do influence both premise pair learning and performance in TI tests. These results differ from previous studies that found little or no effect of efforts to manipulate associative value in crows and pigeons using overtraining (Lazareva, et al., 2004; Lazareva & Wasserman, 2006). Differences in reward magnitude likely have considerably stronger effects on associative values than does overtraining. Discrepancy models of associative learning assume that the associative value of a stimulus will reach an asymptotic value based on the value of the reinforcer (Rescorla & Wagner, 1972), therefore the larger reinforcer values in this experiment likely resulted in larger associative values than those hypothesized to exist between TI images based on computational modeling of standard TI training. Thus, while this study clearly shows that associative values can control choice in TI tests, the differences in associative value created in this study may be well outside the magnitude of differences that would be relevant in standard tests of TI.

We found that choice in TI tasks can be controlled by differences in associative value when associative value and inference are incongruent with one another. But even in the Incongruent condition, monkeys selected images consistent with the implied order if the target and the distracter had been reinforced with the same number of pellets (pairs AC and DF). This suggests that monkeys may have encoded both implied order and associative value during premise pair training and associative value may mask knowledge of implied order when differences in associative value are sufficiently large. To further evaluate whether monkeys infer order in TI tests, we presented a list-linking task that can only be solved using implied order. Choices based on associative value and inferred order would produce very different patterns of performance in this test.

Experiment 4- List Linking

To further evaluate the extent to which choices on TI tasks are controlled by the implied order of stimuli, we tested the ability of monkeys to link two previously learned lists into a single list after experience with just one linking pair of images. This linking could not be accomplished by associative strength. Because the two original lists (A, B, C, D, E, F, G and T, U, V, W, X, Y, Z) were trained separately and performance did not differ between them, images occupying the same relative location in the lists (e.g. B and U, D and W) should have accrued similar associative values during training. Associative values therefore predict interleaving of the lists rather than linking of the lists end to end. We trained monkeys with a single linking image pair consisting of the lowest item (G) in one of the lists learned in Experiment 1 (A>B>C>D>E>F>G) with the highest image (T) in the other list (T>U>V>W>X>Y>Z).

If choices in TI tests are controlled by associative values, training on the linking pair G+T- will not result in systematic selection of all images from the higher ranked list over all images in the lower ranked list but instead would result in inter-mixing the images in the two lists. Presentation of Same Location pairs, which consist of images occupying the same location in their respective lists, would result in chance performance. In tests with pairs in which the image from the higher ranked list occupied a relatively lower position in the training list (Lower Location pairs, e.g. F and U) choice by associative value would result in below chance accuracy with respect to implied order. Additionally, because images in the two lists have similar associative values, choice by associative value does not predict an SDE spanning the two linked lists. In contrast, if choices on TI tasks are controlled by inferred order, then monkeys should link the two independently learned lists into a single 14 item list (A>B>C>D>E>F>G>T>U>V>W>X>Y>Z) and correctly select any item from the

higher list over any item from the lower list. Additionally, because they have linked the two lists into one large list, monkeys should show a SDE that spans the entire 14 item list.

Method

Subjects and apparatus—Subjects and apparatus were the same as used in Experiments 1, 2, and 3.

Stimuli and procedure—Stimuli consisted of the two sets of seven color clip art images that were used in the two repetitions of Experiment 1. Training and testing reward contingencies were the same as in Experiment 1.

Re-familiarization—In order to ensure that subjects remembered the premise pairs learned in Experiment 1, they were presented with re-familiarization sessions consisting of 25 trials of each of the 6 previously trained adjacent premise pairs from one of the two lists (AB, BC, CD, DE, EF, FG). Once they reached >80% on all six premise pairs simultaneously in one session, they were presented with sessions containing the 6 premise pairs from the second list (TU, UV, VW, WX, XY, YZ) until they reached this same criterion. Finally, they were presented with sessions in which all 12 of the premise pairs from the two lists were intermixed. During this re-familiarization phase none of the pairs spanned the two lists, thus monkeys were familiarized with test sessions containing 12 test pairs intermixed, but could not link the two previously learned lists at this stage.

Linking—List linking training sessions presented 25 trials of the linking pair in which the lowest item (G) from the to-be- higher ranked list was rewarded when paired with the highest item (T) from the to-be-lower ranked list until subjects performed above 80%. For half of the subjects the higher ranked list was the first one learned in Experiment 1, for the other half of the subjects it was the second one learned in Experiment 1. Next subjects received training sessions in which all 13 training pairs were intermixed (the 12 premise pairs from the two previously learned lists and the one linking pair) until they performed above 80% on all 13 pairs in a session.

TI Testing—Test sessions consisted of all possible non-adjacent test pairings pseudo-randomly intermixed with the 13 training pairs in a session containing 403 trials. The 13 premise pairs and linking pair made up 325 of these trials (25 of each trial type), within list test pairs (non-adjacent pairs of stimuli from within the same list, e.g. AC, BD, TW) made up 30 of these trials, and between list test pairs (never before seen non-adjacent test pairs spanning the two lists, e.g. AZ, BW) made up 48 of the trials. Monkeys received four test sessions.

Results and Discussion

Re-familiarization and training—The number of errors made before reaching criterion in the re-familiarization phase did not differ between the two previously trained lists ($M_{List1} = 362$, $SEM = 107$; $M_{List2} = 360$, $SEM = 76$; paired samples t-test $t_{11} = 0.04$, $p = .97$). Monkeys made an average of 94 ($SEM = 9.7$) errors on the linking pair (GT) before reaching criterion.

In order for associative values to result in above chance performance on both within and between list test trials, images in the lower list would have to accrue values below those accrued to images in the higher list, but images within each list would have to maintain their relative values. This change would occur only after linking, when all pairs were intermixed, and would be expected to result in changes in premise pair performance compared to pre-linking performance levels. In contrast to associative predictions, in the first session in

which the 13 premise pairs (6 from the higher list, 6 from the lower list, and 1 linking pair) were intermixed, subjects performed at pre-linking rates on all within-list pairs except for the pairs containing the linking images (FG, GT, and TU). As shown in Figure 8, above chance performance was maintained on all these pairs throughout intermixed training sessions. Performance on pairs containing the linking images (pairs FG, GT, and TU) did decrease compared to performance during re-familiarization both with and without corrections for family wise error (uncorrected $\alpha = .05$, Bonferroni corrected $\alpha = .017$; paired samples t-tests: FG: $t_{11} = 9.48, p < .01$; GT: $t_{11} = 4.47, p < .01$; TU: $t_{11} = 6.23, p < .01$). This pattern of decreased performance after addition of a new adjacent pair is typical in sequential transitive inference training (Treichler & Van Tilburg, 1996) and can be explained by either inclusion of new images into an ordered representation or changes to the associative values of the individual images.

Test—After linking, subjects maintained test trial performance within the previously learned list lists, performing significantly above chance on the critical internal non-adjacent within list test pairs (Higher list: $M = 70.3\%$, $SEM = 4.0$; $t_{11} = 5.17, p < .01$; Lower list: $M = 66.8\%$, $SEM = 4.6$; $t_{11} = 3.56, p < .02$). They also showed the SDE for accuracy within both lists, although there was no difference in response latency across the symbolic distances (Higher list: accuracy: $F_{4, 44} = 15.69, p < .01$, latency: $F_{4, 44} = 0.87, p = .49$; Lower list: accuracy: $F_{4, 44} = 6.07, p < .01$, latency: $F_{4, 44} = 1.87, p = .13$).

Accuracy and response latency on between list test pairs did not differ as a function of whether the list learned first or second in Experiment 1 took the higher or lower position when the lists were linked (independent samples t-tests: accuracy $t_{10} = 0.18, p = .86$; response latency $t_{10} = 3.45, p = .09$). Therefore data were collapsed for further analyses.

Monkeys were more accurate than expected by chance on critical internal non-adjacent between list test trials ($M = 66.5\%$, $SEM = 3.7$; $t_{11} = 7.15, p < .001$). Associative value accounts predict below chance performance on Lower pairs, in which the image from the higher ranked list occupies a relatively lower position in its initially trained list, and chance performance on Same pairs, in which the images occupy the same location in their respective lists. In contrast, monkeys performed above chance regardless of the relative locations of the test images in their originally trained lists (Lower: $M = 59.4\%$, $SEM = 2.7$; $t_{11} = 3.43, p = .01$; Same: $M = 77.3\%$, $SEM = 2.5$; $t_{11} = 9.82, p < .01$), suggesting that their choices were driven by the inferred order of the stimuli. Accuracy was significantly lower on Lower pairs than on Same pairs (paired-samples t-test: $t_{12} = -5.31, p < .01$). This difference may reflect the influence of associative values on choices in cross list pairs. However, these pairs differed in symbolic distance, with Lower pairs having small symbolic distances ranging from 1 to 5, while Same pairs had the larger symbolic distance of 6. Thus either associative values or implied order could account for the difference between Lower and Same test pair performance.

To determine if subjects integrated the two separate lists into one large 14 item list after linking training, we examined the SDE for novel between list test pairs. As shown in Figure 9, the SDE for accuracy was evident across the 12 between list symbolic distances (RMANOVA: $F_{11, 121} = 38.27, p < .01$). This pattern suggests that monkeys formed a unified linear representation of one large 14 item list. However monkeys did not show a significant SDE for response latency, Figure 9 (RMANOVA: $F_{11, 121} = .56, p = .86$). Latencies on between list test pairs were highly variable, as indicated by the large standard errors. Previous studies of TI in animals have found similar results, with an SDE for accuracy but none for response latency (Vasconcelos, 2008). The novelty of the task may have led to this large variability and may account for the lack of a systematic pattern in response latency.

Together, success with between list test pairs and the SDE for accuracy spanning the two lists suggest that monkeys linked the two separately learned seven item lists into one fourteen item list. This linking was done after exposure to only one linking pair (GH) and on the monkeys' first experience with this type of task. These results support and expand upon previous findings that experienced monkeys link five item lists and select the higher ranking item on between list pairs regardless of their relative rankings in their initially trained lists (Treichler & Raghanti, 2010; Treichler, et al., 2003; Treichler et al., 2007; Treichler & Van Tilburg, 1999; Treichler & VanTilburg, 1996). List linking results like those shown here cannot be explained by any of the current associative models of performance (Lazareva, 2012), and suggest that performance on TI tasks is controlled by the implied order of stimuli.

General Discussion

Together the findings from three of our four experiments indicate that choice by monkeys in TI tests is typically controlled primarily by the implied order of stimuli. While under the extreme conditions of reward magnitude manipulation used in Experiment 3 associative values played a clear role in choice, the results from Experiment 2 showed that choice on a standard TI task did not correlate with measured associative values, but instead correlated with the distance between images in a pair. Additionally, in Experiment 4, monkeys linked two lists in a way that would not be possible if choice was primarily controlled by associative value. Like most behavior, TI performance appears to be subject to influence by multiple cognitive systems, but our results make clear that the capacity to represent the implied order of TI stimuli is present and strong in rhesus monkeys.

Contributions of Associative Values and Implied Order

It is well established that choice behavior in various reinforcement schedules and object discrimination tasks can be controlled by rates of reinforcement and non-reinforcement (Domjan, 2004; Tarpay, 1997). In Experiment 3 monkeys selected the item associated with the larger reinforcer over the item that was correct according to the implied order. This result indicates that under certain conditions, TI choices, like choices on many cognitive tasks, can be controlled by associative values. Associative value models successfully predict transitive inference performance by pigeons (Siemann, et al., 1996; von Fersen, et al., 1991; Wynne, 1995, 1998). While these models have not yet been applied to monkey data, our empirical measurements of associative values in Experiment 2 suggest that performance cannot be well accounted for by associative values alone. Therefore, to the extent that associative models correctly model associative values, they would not predict monkey performance.

Choices by monkeys were primarily controlled by the implied order of stimuli in three of the four TI experiments presented here. Seamless linking of two seven item lists into one 14 item list in Experiment 4 cannot be explained by any of the current associative models (Lazareva, 2012). Even when choices were controlled by associative values in Experiment 3, above chance performance on the equally reinforced pairs suggested latent knowledge of the implied order was present but was masked by the large differences in reinforcement associated with some TI images. These results suggest that monkeys extract the implied order from overlapping pairs of stimuli independent of the associative values of individual stimuli. It would be of interest to test this idea using a post training manipulation of reward values to unmask latent knowledge, such as by selective satiation or reward devaluation.

The results of the present series of experiments implicate contributions of both associative values and implied order to TI performance in monkeys and highlight the dual nature of the mechanisms underlying choice. While this means that choices may sometimes be driven by

associative values, monkeys may simultaneously have an underlying representation of the implied order of the stimuli that is not expressed. The extent to which associative value and implied order control performance may depend on species, task parameters, or subject expertise (Bond et al., 2010; Lazareva, et al., 2004; Lazareva & Wasserman, 2006; Maclean, et al., 2008).

Comparative Implications

The results of our experiments indicate that even for an individual subject, multiple cognitive mechanisms are involved in TI choice. Within humans the mechanisms behind TI task performance vary between individuals based on age, experience, and awareness of the task (Bryant & Trabasso, 1971; Goswami, 1995; Pears & Bryant, 1990; Russell et al., 1996). TI choices by adult humans who are highly aware of the implied order of stimuli appear to be primarily driven by an explicit representation of the relations between images, while correct choices by participants unaware of the implied order appear to be driven by implicit, perhaps associative, knowledge of which item was correct (Siemann & Delius, 1996; Smith & Squire, 2005). This dissociation between cognitive mechanisms used by aware and unaware participants is further supported by findings from neuroimaging studies indicating that aware and unaware participants show differential brain activity patterns during TI test trials (Greene et al., 2006; Moses, Brown, et al., 2010).

Differences in the relative contribution of cognitive mechanisms to TI performance between subjects of the same species are likely to extend to differences in the relative contribution of cognitive mechanisms between species. Even across closely related species, learning rates and performance patterns on transitive inference tasks vary enough to suggest differences in the relative contributions or functioning of the relevant cognitive mechanisms (Bond, et al., 2010; Bond, et al., 2003; Lazareva, et al., 2004; Maclean, et al., 2008). While associative models can account for TI performance in pigeons, they have been less successful at predicting performance in other avian species (Bond, et al., 2010). Within the corvid and lemur families, species that live in complex social environments show performance patterns that are more consistent with use of inferred order than species with natural histories that do not include these complex cognitive demands (Bond, et al., 2003; Maclean, et al., 2008). A species' natural history may be predictive of differences in the relative contributions of different cognitive mechanisms to task performance. It is therefore invalid to assume that choice in all species is controlled by the same mechanisms to the same extent, and studies like Experiment 3 that manipulate the salience of associative values may be useful in elucidating species differences.

Mental Representations

To fully understand how animals and humans solve TI tasks, the mechanisms responsible for representing implied order and how they differ from associative mechanisms must be better elucidated. Whereas the mechanisms of associative learning are comparatively well understood (Domjan, 2004; Rescorla & Wagner, 1972) and the mechanisms by which associative values could account for TI performance have been extensively modeled (Siemann, et al., 1996; Vasconcelos, 2008; von Fersen, et al., 1991; Wynne, 1997), the mechanisms underlying representation of implied order are poorly characterized. Logical inference is invoked largely because of evidence against associative accounts, rather than because of positive evidence for specific alternative cognitive representational systems. Neurobiological evidence implicates hippocampal processing and explicit memory, but still falls short of providing a clear cognitive account (Greene, et al., 2006; Greene et al., 2001; Smith & Squire, 2005).

Informal descriptions of the process involved in TI performance are almost surely misleading. For example, online inferences of the type generally referred to in TI examples (if Ben is taller than Emily and Emily is taller than Dina, then Ben is taller than Dina), would not produce the patterns of performance seen in TI tasks like those reported here. If a subject who is presented with test pair BD actively inferred “if $B > C$ and $C > D$, then $B > D$,” this online inference would lead to longer response latencies and decreased accuracy with increasing symbolic distance between items, as more inferences need to be made for more disparate items (Banks, 1977; McGonigle & Chalmers, 1992; Vasconcelos, 2008). However the prevalence of the SDE in TI tasks, which shows increased accuracy and often shorter response latencies with larger symbolic distances, suggests that a representation of the relations between items is formed during training, and is later referenced to solve test trials. This hypothesis that a mental representation of the ordered list is created during training is further supported by limited evidence that animals trained on TI premise pairs before hippocampal system disruption perform well on post-lesion TI test trials (Van der Jeugd, et al., 2009), whereas animals trained after lesion perform poorly (Buckmaster et al., 2004; Dusek & Eichenbaum, 1997). However disruptions to the prefrontal cortex produce retarded premise pair learning and deficits in performance on TI test trials (DeVito, et al. 2010). The hippocampus may be necessary for forming a representation of TI stimuli during training, but not for accessing this pre-existing representation at test (Van der Jeugd, et al., 2009).

While the hippocampus is involved in a wide variety of cognitive processes, and the best description of its function is a matter of debate, there is no debate that it is critical for some types of spatial processing (Hampton et al., 2004; Hampton & Shettleworth, 1996; Lavenex & Lavenex, 2009; Spiers & Maguire, 2007). The implication of the hippocampus in TI performance (Buckmaster, et al., 2004; Dusek & Eichenbaum, 1997; Fortin et al., 2002; Greene, et al., 2006; Heckers et al., 2004; Moses, et al., 2010; Nagode & Pardom, 2002; Van der Jeugd, et al., 2009; Zalesak & Heckers, 2009) suggests that the representation created during TI training may be spatially organized (Moses, et al., 2010). In humans, both spatial and non-spatial ordered information is often cognitively represented as a spatially organized “mental line” (Holmes & Lourenco, 2011; Prado et al., 2008; Previtali et al., 2010; Schwarz & Keus, 2004; Shaki & Fischer, 2008). In a TI task this type of representation could result in a mental line with item A on one end (e.g. left), item G on the other end (e.g. right), and items B, C, D, E, and F located linearly between (Brunamonti et al., 2011; Chen et al., 1997; D’Amato & Colombo, 1990; Roberts & Phelps, 1994). At test this mental line would be referenced and the leftmost item in the pair would be correctly selected, producing above chance performance on TI test trials. Items further apart on this mental line would be easier to distinguish, resulting in the SDE. This spatial representation hypothesis accounts for the contributions of the implied order found in the present experiments, and provides a basis for “logical inference” without the need to invoke more abstract logical cognitive processes such as those described by Piaget (Piaget & Inhelder, 1967).

In humans, TI tasks produce performance patterns consistent with a spatial representation of the relations between stimuli (Brunamonti, et al., 2011; Moses, Ostreicher, et al., 2010; Previtali, et al., 2010). Monkeys may represent ordered information such as time spatially (Merritt et al., 2010), and the limited evidence available in animals supports use of a spatial representation in TI tasks by the species that have been tested. Rats learned a TI task faster when the stimuli were trained in a physical linear order (Roberts & Phelps, 1994). Mongoose lemurs and crows only performed above chance on TI tasks when the training conditions highlighted the linear order of the stimuli (Lazareva, et al., 2004; Maclean, et al., 2008). Together, these results support the possibility that humans and some animals may represent TI stimuli in a linear spatial representation.

Natural Function of TI

Many species live in social groups organized around a linear social dominance hierarchy (Appleby, 1983) that could be learned and represented using implied order, like the stimuli in laboratory TI tasks. An animal capable of organizing dominance information into an ordered representation using TI could place itself within the hierarchy after only a small number of first party interactions (Paz-y-Miño, et al., 2004), and could more easily maintain information about a large number of relationships by representing them in the ordered fashion. Elegant studies using controlled, live social stimuli have found that cichlid fish (*A. burtoni*), pinyon jays (*G. cyanocephalus*), and chickens (*G. domesticus*) use TI to learn new dominance relations, reacting appropriately on the first encounter with an individual they had only observed engaging in dominance interactions (Grosenick, et al., 2007; Hogue, et al., 1996; Paz-y-Miño, et al., 2004). Chickens and jays appropriately display subordinate behavior during their first interaction with an individual they observed dominating an individual known to be dominant to themselves (Hogue, et al., 1996; Paz-y-Miño, et al., 2004). Cichlid fish that have a natural preference for lower ranking individuals choose to be near an individual inferred to be of lower rank based on a series of 4 overlapping dominance interactions (Grosenick, et al., 2007). These studies test TI in non-human animals without food reinforcement. Associative value mechanisms would be unlikely to support such dominance hierarchy learning, as observed third-party dominance interactions are not followed by explicit reinforcement. These results suggest that, like our rhesus monkeys, other species may solve TI tasks without relying on associative values alone, and that the cognitive mechanisms used to solve abstract laboratory TI tasks may be useful in natural social contexts.

Acknowledgments

The Yerkes National Primate Research Center is fully accredited by the American Association for Accreditation of Laboratory Animal Care. This project was funded by the National Center for Research Resources P51RR165, currently supported by the Office of Research Infrastructure Programs/OD P51OD11132, the Center for Behavioral Neuroscience under the STC Program of the NSF under Agreement IBN-9876754, a grant from the James S. McDonnell Foundation, NIH grant 1R01MH082819, NSF grants 0745573 and 1146316. Regina Paxton Gazes was supported in part by NSF grant DGE- 0231900.

We thank Dina P. Chou, Steven R.L. Sherrin, and Emily K. Brown for help with testing, Jack McDowell for consulting on the design of Experiment 2, Nancy Bliwise for statistical advice, and Olga Lazareva for discussions about transitive inference.

References

- Appleby MC. The probability of linearity in hierarchies. *Anim Behav.* 1983; 31:600–608.
- Armus HL, Garlich MM. Secondary reinforcement strength as a function of schedule of primary reinforcement. *Journal of Comparative and Physiological Psychology.* 1961; 54:56–58. [PubMed: 13684319]
- Aron, A.; Aron, E. *Statistics for psychology.* Upper Saddle River, NJ: Prentice Hall; 1999.
- Banks, WP. Encoding and processing of symbolic information in comparative judgments. In: Bower, GH., editor. *The psychology of learning and motivation.* Vol. 11. New York: Academic Press; 1977. p. 101-159.
- Benard J, Giurfa M. A test of transitive inferences in free-flying honeybees: Unsuccessful performance due to memory constraints. *Learning & Memory.* 2004; 11(3):328–336. [PubMed: 15169863]
- Bond A, Wei CA, Kamil AC. Cognitive representation in transitive inference: A comparison of four corvid species. *Behavioral Processes.* 2010
- Bond AB, Kamil AC, Balda RP. Social complexity and transitive inference in corvids. *Animal Behaviour.* 2003; 65:479–487.

- Brunamonti E, Genovesio A, Carbe K, Ferraina S. Gaze modulated non-propositional reasoning: Further evidence for spatial representation of reasoning premises. *Neuroscience*. 2011; 173:110–115. [PubMed: 21075173]
- Bryant PE, Trabasso T. Transitive inferences and memory in young children. *Nature*. 1971; 232(5311): 456–458. [PubMed: 4937205]
- Buckmaster CA, Eichenbaum H, Amaral DG, Suzuki WA, Rapp PR. Entorhinal cortex lesions disrupt the relational organization of memory in monkeys. *Journal of Neuroscience*. 2004; 24(44):9811–9825. [PubMed: 15525766]
- Chen S, Swartz KB, Terrace HS. Knowledge of the ordinal position of list items in rhesus monkeys. *Psychological Science*. 1997; 8(2):80–86.
- D'Amato M, Lachman R. Secondary reinforcement as affected by reward schedule and the testing situation. *Journal of Comparative Physiological Psychology*. 1958; 51:737–741. [PubMed: 13620815]
- D'Amato, MR. Comparative cognition: Processing of serial order and serial pattern. In: Dachowski, L.; Flaherty, CF., editors. *Current topics in animal learning: Brain, emotion and cognition*. Hillsdale, New Jersey: L Erlbaum; 1991. p. 165-185.
- D'Amato MR, Colombo M. The symbolic distance effect in monkeys (*Cebus apella*). *Animal Learning & Behavior*. 1990; 18(2):133–140.
- Davis H. Transitive inferences in rats (*Rattus norvegicus*). *Journal of Comparative Psychology*. 1992; 106:342–349. [PubMed: 1451416]
- DeVito LM, Lykken C, Kanter BR, Eichenbaum H. Prefrontal cortex: Role in acquisition of overlapping associations and transitive inference. *Learning & Memory*. 2010; 17(3):161–167. [PubMed: 20189961]
- Domjan, M. *The Essentials of Conditioning and Learning*. Belmont, CA: Wadsworth Publishing; 2004.
- Dusek JA, Eichenbaum H. The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences of the United States of America*. 1997; 94(13):7109–7114. [PubMed: 9192700]
- Fortin NJ, Agster KL, Eichenbaum HB. Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*. 2002; 5(5):458–462.
- Gillan DJ. Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology-Animal Behavior Processes*. 1981; 7(2):150–164.
- Goswami U. Transitive relational mappings in 3 year olds and 4 year olds: The analogy of Goldilocks and the 3 bears. *Child Development*. 1995; 66(3):877–892.
- Greene AJ, Gross WL, Elsinger CL, Rao SM. An fMRI analysis of the human hippocampus: Inference, context, and task awareness. *Journal of Cognitive Neuroscience*. 2006; 18(7):1156–1173. [PubMed: 16839289]
- Greene AJ, Spellman BA, Dusek JA, Eichenbaum HB, Levy WB. Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans. *Memory & Cognition*. 2001; 29:893–902. [PubMed: 11716062]
- Grosenick L, Clement TS, Fernald RD. Fish can infer social rank by observation alone. *Nature*. 2007; 445:429–432. [PubMed: 17251980]
- Hampton RR, Hampstead BM, Murray EA. Selective hippocampal damage in rhesus monkeys impairs spatial memory in an open-field test. *Hippocampus*. 2004; 14(7):808–818. [PubMed: 15382251]
- Hampton RR, Shettleworth SJ. Hippocampal lesions impair memory for location but not color in passerine birds. *Behavioral Neuroscience*. 1996; 110(4):831–835. [PubMed: 8864273]
- Heckers S, Zalesak M, Weiss AP, Ditman T, Titone D. Hippocampal activation during transitive inference in humans. *Hippocampus*. 2004; 14:153–162. [PubMed: 15098721]
- Hogue ME, Beaugrand JP, Lague PC. Coherent use of information by hens observing their former dominant defeating or being defeated by a stranger. *Behavioural Processes*. 1996; 38:241–252.
- Holmes KJ, Lourenco SF. Common spatial organization of number and emotional expression: A mental magnitude line. *Brain and Cognition*. 2011; 77(2):315–323. [PubMed: 21839568]

- Lavenex PB, Lavenex P. Spatial memory and the monkey hippocampus: Not all space is created equal. *Hippocampus*. 2009; 19(1):8–19.10.1002/hipo.20485 [PubMed: 18727046]
- Lazareva, OF. Transitive inference in non-human animals. In: Zentall, TR.; Wasserman, EA., editors. *The Oxford Handbook of Comparative Cognition*. 2. Oxford: Oxford University Press; 2012.
- Lazareva OF, Smirnova AA, Bagozkaja MS, Zorina ZA, Rayevsky VV, Wasserman EA. Transitive responding in hooded crows requires linearly ordered stimuli. *Journal of the Experimental Analysis of Behavior*. 2004; 82(1):1–19. [PubMed: 15484868]
- Lazareva OF, Wasserman EA. Effect of stimulus orderability and reinforcement history on transitive responding in pigeons. *Behavioural Processes*. 2006; 72(2):161–172. [PubMed: 16460886]
- Lazareva OF, Wasserman EA. Transitive inference in pigeons: Measuring the associative value of stimuli B and D. *Behavioural Processes*. 2012; 89(3):244–255. [PubMed: 22178714]
- Maclean EL, Merritt DJ, Brannon EM. Social complexity predicts transitive reasoning in prosimian primates. *Animal Behaviour*. 2008; 76:479–486. [PubMed: 19649139]
- McGonigle BO, Chalmers M. Are Monkeys Logical. *Nature*. 1977; 267(5613):694–696. [PubMed: 406574]
- McGonigle BO, Chalmers M. Monkeys are rational! *Quarterly Journal of Experimental Psychology Section B-Comparative and Physiological Psychology*. 1992; 45(3):189–228.
- Merritt DJ, Casasanto D, Brannon EM. Do monkeys think in metaphors? Representations of space and time in monkeys and humans. *Cognition*. 2010; 117:191–202. [PubMed: 20846645]
- Merritt DJ, Terrace HS. Mechanisms of inferential order judgments in humans (*Homo sapiens*) and rhesus monkeys (*Macaca mulatta*). *J Comp Psychol*. 2011; 125(2):227–238. [PubMed: 21341909]
- Miller HJ Jr. Matching based hedonic scaling in the pigeon. *Journal of Experimental Analysis of Behavior*. 1976; 26:335–347.
- Montgomery KC. Concerning the use of analysis of variance on latency data. [Note]. *American Journal of Psychology*. 1953; 66(1):131–135. [PubMed: 13030859]
- Moses SN, Brown TM, Ryan JD, McIntosh AR. Neural system interactions underlying human transitive inference. *Hippocampus*. 2010; 20(8):894–901. [PubMed: 20054816]
- Moses SN, Ostreicher ML, Ryan JD. Relational framework improves transitive inference across age groups. *Psychological Research-Psychologische Forschung*. 2010; 74(2):207–218.
- Moyer RS, Landauer TK. Time required for judgments of numerical inequality. *Nature*. 1967; 215:1519–1520. [PubMed: 6052760]
- Nagode JC, Pardom JV. Human hippocampal activation during transitive inference. *Learning and Memory*. 2002; 13(7):939–944.
- Paz-y-Miño G, Bond A, Kamil AC, Balda RP. Pinyon jays use transitive inference to predict social dominance. *Nature*. 2004; 430:778–781. [PubMed: 15306809]
- Pears R, Bryant P. Transitive inferences by young children about spatial position. *British Journal of Psychology*. 1990; 81:497–510.
- Piaget, J. *Logic and psychology*. New York: Basic Books, Inc; 1960.
- Piaget, J.; Inhelder, B. *The child's conception of space*. New York: W.W. Norton; 1967.
- Prado J, Van der Henst JB, Noveck IA. Spatial associations in relational reasoning: Evidence for a SNARC-like effect. *Quarterly Journal of Experimental Psychology*. 2008; 61(8):1143–1150.
- Previtali P, de Hevia MD, Girelli L. Placing order in space: the SNARC effect in serial learning. *Experimental Brain Research*. 2010; (201):599–605.
- Rapp PR, Kansky MT, Eichenbaum H. Learning and memory for hierarchical relationships in the monkey: Effects of aging. *Behavioral Neuroscience*. 1996; 110(5):887–897. [PubMed: 8918992]
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Prokasy, AHBWF., editor. *Classical conditioning II: Current research and theory*. New York: Appelton-Century-Crofts; 1972. p. 64–99.
- Roberts WA, Phelps MT. Transitive inference in rats: A test of the spatial coding hypothesis. *Psychological Science*. 1994; 5(6):368–374.
- Russell J, McCormack T, Robinson J, Lillis G. Logical (versus associative) performance on transitive reasoning tasks by children: Implications for the status of animals' performance. *Quarterly Journal*

of Experimental Psychology Section B-Comparative and Physiological Psychology. 1996; 49(3): 231–244.

- Schwarz W, Keus IM. Moving the eyes along the mental number line: Comparing SNARC effects with saccadic and manual responses. *Perception & Psychophysics*. 2004; 66(4):651–664. [PubMed: 15311664]
- Shaki S, Fischer MH. Reading space into numbers - a cross-linguistic comparison of the SNARC effect. *Cognition*. 2008; 108(2):590–599. [PubMed: 18514179]
- Siemann M, Delius JD. Influences of task concreteness upon transitive responding in humans. *Psychological Research-Psychologische Forschung*. 1996; 59(2):81–93.
- Siemann M, Delius JD. Algebraic learning and neural network models for transitive and non-transitive responding. *European Journal of Cognitive Psychology*. 1998; 10(3):307–334.
- Siemann M, Delius JD, Dombrowski D, Daniel S. Value transfer in discriminative conditioning with pigeons. *Psychological Record*. 1996; 46(4):707–728.
- Siemann M, Delius JD, Wright AA. Transitive responding in pigeons: Influences of stimulus frequency and reinforcement history. *Behavioural Processes*. 1996; 37(2–3):185–195.
- Smith C, Squire L. The importance of awareness in transitive inference. *Journal of Cognitive Neuroscience*. 2005:239–239.
- Spiers HJ, Maguire EA. The neuroscience of remote spatial memory: A tale of two cities. [Review]. *Neuroscience*. 2007; 149(1):7–27.10.1016/j.neuroscience.2007.06.056 [PubMed: 17850977]
- Tarpy, R. Contemporary learning theory and reasearch. New York: McGraw–Hill Companies, Inc; 1997.
- Treichler FR, Raghanti MA. Serial list combination by monkeys (*Macaca mulatta*): test cues and linking. *Animal Cognition*. 2010; 13:121–131. [PubMed: 19533183]
- Treichler FR, Raghanti MA, Van Tilburg DN. Linking of serially ordered lists by macaque monkeys (*Macaca mulatta*): List position influences. *Journal of Experimental Psychology-Animal Behavior Processes*. 2003; 29(3):211–221. [PubMed: 12884680]
- Treichler FR, Raghanti MA, Van Tilburg DN. Serial list linking by macaque monkeys (*Macaca mulatta*): List property limitations. *Journal of Comparative Psychology*. 2007; 121(3):250–259. [PubMed: 17696651]
- Treichler FR, Van Tilburg D. Concurrent conditional discrimination tests of transitive inference by macaque monkeys: List linking. *J Exper Psychol: Anim Behav Proc*. 1996; 22(3):105–117.
- Treichler FR, Van Tilburg D. Training requirements and retention characteristics of serial list organization by macaque monkeys. *Animal Cognition*. 1999; 2:235–244.
- Treichler FR, VanTilburg D. Concurrent conditional discrimination tests of transitive inference by macaque monkeys: List linking. *Journal of Experimental Psychology-Animal Behavior Processes*. 1996; 22(1):105–117. [PubMed: 8568492]
- Van der Jeugd A, Goddyn H, Laeremans A, Arckens L, D’Hooge R, Verguts T. Hippocampal involvement in the acquisition of relational associations, but not in the expression of a transitive inference task in mice. *Behavioral Neuroscience*. 2009; 123(1):109–114. [PubMed: 19170435]
- Vasconcelos M. Transitive inference in non-human animals: An empirical and theoretical analysis. *Behavioural Processes*. 2008; 78:313–334. [PubMed: 18423898]
- von Fersen L, Wynne CDL, Delius JD, Staddon JER. Transitive inference formation in pigeons. *Journal of Experimental Psychology-Animal Behavior Processes*. 1991; 17:334–341.
- Wei BM, Kehmeier S, Schloegl C. Transitive inference in free-living greylag geese, *Anser anser*. *Animal Behaviour*. 2010; 79:1277–1283.
- Woocher FD, Glass AL, Holyoak KJ. Positional discriminability in linear orderings. *Memory and Cognition*. 1978; 6(2):165–173.
- Wynne CDL. Reinforcement accounts for transitive inference performance. *Animal Learning & Behavior*. 1995; 23(2):207–217.
- Wynne CDL. Pigeons transitive inferences: Tests of simple accounts of a complex performance. *Behavioral Processes*. 1997; 39:95–112.
- Wynne, CDL. A minimal model of transitive inference. In: Wynne, CDL.; Staddon, JER., editors. *Models of action: Mechanisms for adaptive behavior*. Mahwah, NJ: Erlbaum; 1998. p. 269-307.

- Zalesak M, Heckers S. The role of the hippocampus in transitive inference. *Psychiatry Research-Neuroimaging*. 2009; 172(1):24–30.
- Zentall TR, Sherburne LM. The transfer of value in simultaneous discriminations: Implications for cognitive and social processes. *Brain and Values*. 1998:323–336.

Adjacent Training Pairs




Figure 1.

Examples of clip art stimuli presented as adjacent training pairs (*left*) and critical non-adjacent internal test pairs (*right*) used in experiment 1. During training, correct selection of the S+ in a given pair resulted in an auditory reinforcer paired with a food reinforcer on 75% of trials. On trials containing test pairs all choices resulted in an auditory reinforcer only.

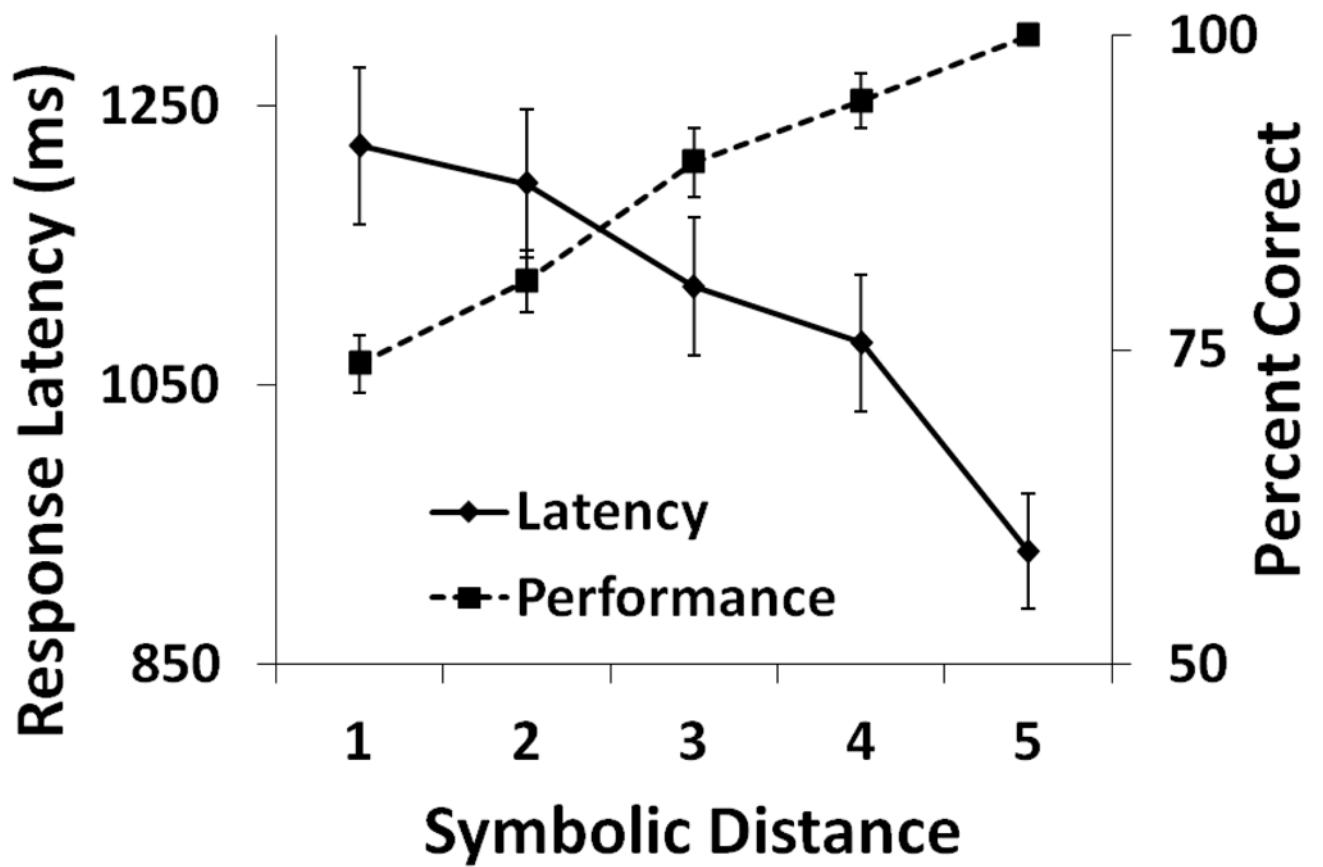


Figure 2. Average accuracy and average median response latency on all non-adjacent test pairs (end anchor and internal) in Experiment 1 by symbolic distance. Error bars represent standard error of the mean.




Figure 3. SDE for response latency (left) and accuracy (right), sorted by the first item in the pair. Stars (*) in the legend indicate pairs for which the repeated measures ANOVA was significant (Table 2). Because this analysis controls for the contribution of the end anchors, it indicates that the overall SDE is not driven entirely by performance on the end anchors.

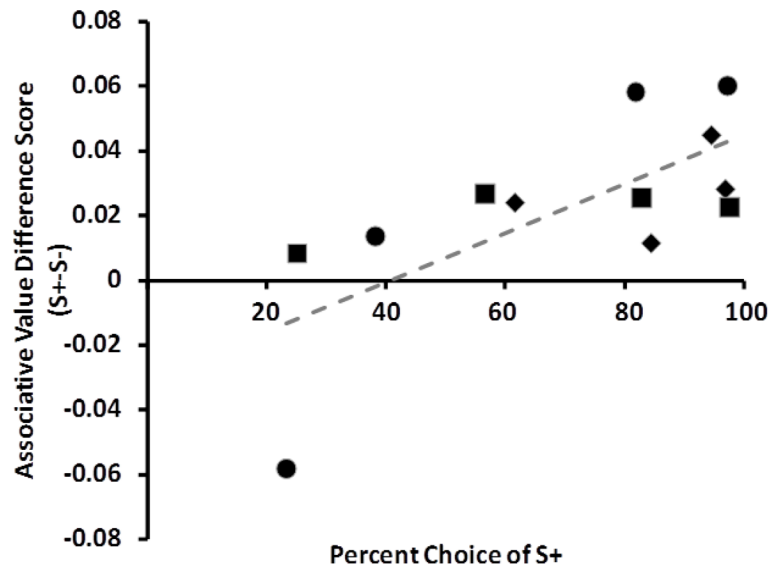


Figure 4.

Average percent choice of the S+ during object discrimination acquisition and reversal plotted by associative value difference score. Each point represents the average of all data from one day of measurement (3 measurement sessions, 24 object discrimination trials). The first two days of measurement and the last two days of measurement are depicted for initial training (diamonds), first (squares) and second (circles) reversal phases. For each phase the left-most point on the x axis corresponds to the first measurement day, and points progress rightward through training to the final criterial point. The trend line is indicated with a dashed line. There was a significant relationship between choice and pair difference score in the object discrimination task.

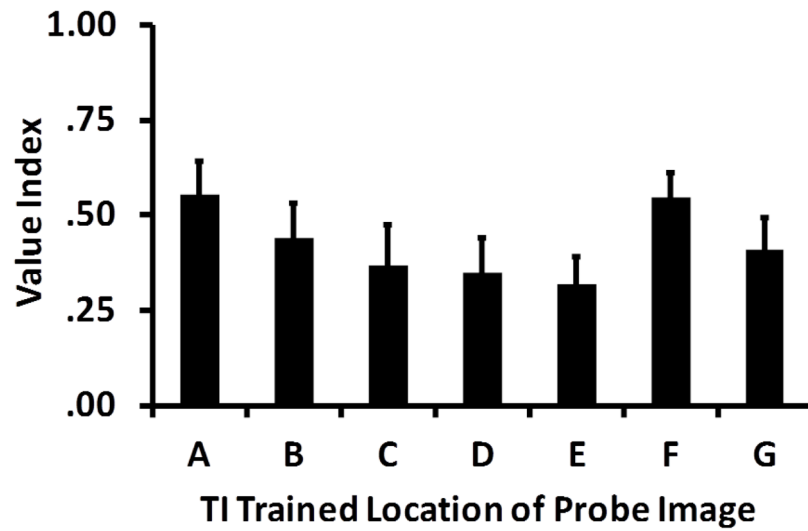



Figure 5. Average measured value index for each TI stimulus when presented as a probe in Experiment 2. Value indexes reflect the proportion of touches to the TI probe image during the 30s presentation in the concurrent RI format. Associative value accounts predict a linear decrease in values from the highest ranked item, A, to the lowest ranked item, G. Error bars are standard errors.






Figure 7. Average total errors to criterion for the 6 premise pairs in the Congruent (solid black) and Incongruent (solid grey) conditions. Premise pairs were learned in the order they are shown from left to right (i.e. pair FG was learned first and AB last). Arrow indicates change in number of reinforcers during training. Reinforcement of pairs CD, BC, and AB in the Incongruent condition is in conflict with the implied order of the stimuli. * indicates a significant difference between conditions for that pair (paired-samples t-tests, $p < .05$). Error bars are standard errors.




Figure 8. Accuracy on the 13 premise pairs in the first intermixed training session of Experiment 4. Accuracy on all pairs except pair HI remained above chance, and only pairs FG, GT, and TU showed significant decrements in accuracy from the last session of re-familiarization (significant difference on a paired samples t-test indicated by *). Error bars indicate standard errors.




Figure 9.

Average accuracy and response latency on all between list test pairs (end anchor and internal) in Experiment 2 by symbolic distance. With the exception of distance 1, accuracy follows the pattern expected if monkeys used transitive inference, latency did not differ systematically across the symbolic distances. Error bars represent standard error of the mean. * indicates above chance accuracy according to one sample t-test.

Training and test phases in Experiment 1. The pairs presented in each phase are shown in columns, with the total number of trials per session in each phase is shown in the bottom row. Subjects were moved on to the next phase after performing over 80% correct in a single session on all pairs present in that phase.

Table 1

Premise pair training phase											
	1	2	3	4	5	6	7	8	9	10	11
Pairs presented	FG		FG	FG	FG	FG	FG	FG	FG	FG	FG
		EF	EF		EF	EF	EF	EF	EF	EF	EF
				DE	DE		DE	DE	DE	DE	DE
						CD	CD	CD	CD	CD	CD
								BC	BC	BC	BC
										AB	AB
Number of trials per session	25	25	50	25	75	25	100	25	125	25	150

Table 2

Statistical results of Repeated Measures ANOVA for SDE by first item as shown in Figure 3. The first item in the pairs included in each analysis is presented on the far left. There was a significant increase in performance with increasing symbolic distance for all first images, and a significant decrease in response latency with increasing symbolic distance for pairs in which A was correct.

First item	Response latency		Percent correct	
	F	P	F	P
A	2.79	.04*	8.65	<.01*
B	1.33	.28	11.93	<.01*
C	3.34	.05	46.37	<.01*
D	1.77	.21	11.88	.01*