# Transitive Inference in Humans (*Homo sapiens*) and Rhesus Macaques (*Macaca mulatta*) After Massed Training of the Last Two List Items

Greg Jensen, Yelda Alkan, Fabian Muñoz, Vincent P. Ferrera, and Herbert S. Terrace Columbia University

Transitive inference (TI) is a classic learning paradigm for which the relative contributions of experienced rewards and representation-based inference have been debated vigorously, particularly regarding the notion that animals are capable of logic and reasoning. Rhesus macaque subjects and human participants performed a TI task in which, prior to learning a 7-item list (ABCDEFG), a block of trials presented exclusively the pair FG. Contrary to the expectation of associative models, the high prior rate of reward for F did not disrupt subsequent learning of the entire list. Monkeys (who each completed many sessions with novel stimuli) learned to anticipate that novel stimuli should be preferred over F. We interpret this as evidence of a task representation of TI that generalizes beyond learning about specific stimuli. Humans (who were task-naïve) showed a transitory bias to F when it was paired with novel stimuli, but very rapidly unlearned that bias. Performance with respect to the remaining stimuli was consistent with past reports of TI in both species. These results are difficult to reconcile with any account that assigns the strength of association between individual stimuli and rewards. Instead, they support sophisticated cognitive processes in both species, albeit with some species differences.

Keywords: cognition, transitive inference, task representation, rhesus macaques

Supplemental materials: http://dx.doi.org/10.1037/com0000065.supp

Transitive inference (TI) is a fundamental process in propositional logic and has been studied by psychologists for over a century (Burt, 1911). If A > B, and B > C, and all three items belong to an ordered list, then the transitive property of the ">" operator permits the conclusion that A > C. Thus, TI provides a formal definition for what it means to "know" that a set of items are ordered. Nevertheless, accounts of TI performance based on stimulus-reward associations have proven remarkably persistent because associative and cognitive accounts of TI learning often make qualitatively similar predictions. Here, we tested a paradigm in which stimulus-reward associations are predicted to substantially impair TI learning.

The first animal study of TI was reported by McGonigle and Chalmers (1977). They presented squirrel monkeys (*Saimiri sciureus*) with adjacent pairs of stimuli from a five-item list (i.e., for a list of stimuli ABCDE, subjects were initially trained on AB, BC, CD, and DE). The "correct" choice was whichever item came earlier in the list. Once subjects reliably selected correct items, they were tested on all 10 possible pairs. Despite having never seen the nonadjacent pairings previously, subjects not only selected the correct items with high accuracy, but did so at rates comparable to those of 4-year-old human children.

Since this initial study, accurate performance on TI tasks has been reported in a growing range of species, and the capacity for TI in animals appears to be ubiquitous among vertebrates (Vasconcelos, 2008). Some form of TI has been reported in at least 20 species (Jensen, 2017), suggesting that TI taps into cognitive faculties that are well preserved across an evolutionarily wide range of species.

Two major behavioral features have been regularly reported in the TI literature. One is the *terminal item effect* (Wynne, 1997). The first item in a TI list is always rewarded, while the last is never rewarded, so these items are less ambiguous than the nonterminal items. This appears to yield a boost to their accuracy. The other is the *symbolic distance effect* (D'Amato & Colombo, 1990). Adjacent pairs typically evoke the lowest response accuracy, whereas items two positions apart (e.g., BD) have higher accuracy, and so on, such that the pair comprising the two terminal items has the highest accuracy. In both cases, higher accuracy tends to be correlated with lower reaction times (RTs), so terminal item effects and symbolic distance effects are sometimes also reported with respect to reaction time (RT; McGonigle & Chalmers, 1992).

Despite this rich literature of empirical results, the field remains divided on how TI itself is performed. The cognitive approach to explaining TI in animals rests on the premise that subjects form a representation of the list (Jensen, Muñoz, Alkan, Ferrera, & Ter-

This article was published Online First March 23, 2017.

Greg Jensen, Department of Psychology, Columbia University; Yelda Alkan and Fabian Muñoz, Department of Neuroscience, Columbia University; Vincent P. Ferrera, Department of Neuroscience and Department of Psychiatry, Columbia University; Herbert S. Terrace, Department of Psychology, Columbia University.

This work was supported by U.S. National Institute of Mental Health, grant NIH-MH081153 awarded to Vincent P. Ferrera and Herbert S. Terrace. We thank Carolina Montes, Orly Morgan, Grant Spencer, and Nathan Tumazi for their assistance collecting the human data, and Brendan Villalobos for his help in programming the human task.

Correspondence concerning this article should be addressed to Greg Jensen, Department of Psychology, Columbia University, 1190 Amsterdam Avenue MC 5501, 406 Schermerhorn Hall, New York, NY 10027. E-mail: greg.guichard.jensen@gmail.com

race, 2015; Zentall, 2001). Several lines of evidence are consistent with the cognitive approach. For example, subjects can transfer prior serial learning between TI and a serial task with qualitatively different task demands (Jensen, Altschul, Danly, & Terrace, 2013). In addition, the serial position of each stimulus can be dissociated from its reward magnitude without disrupting subsequent TI (Gazes, Chee, & Hampton, 2012).

Many studies of TI in animals reject the cognitive interpretation in favor of models of associative learning (Allen, 2006; Couvillon & Bitterman, 1992). Under the logic of this approach, subjects can infer the order of stimuli by determining how strongly each stimulus is associated with the amount and/or probability of reward. Retrospective calculation of stimulus-reward association cannot explain successful inferences based on traditional adjacent-pair training, because the frequency of reward is equal across all nonterminal stimuli. Nevertheless, more elaborate associative models have been proposed, such as value transfer theory (von Fersen, Wynne, Delius, & Staddon, 1991), which posits that associative strength can "leak" between stimuli (such that stimulus A, being a "proven winner," imparts some of its value to B). This, in principle, could permit TI-like effects in some scenarios without necessitating a full-blown representation of the ordered list.

Another proposed approach for explaining TI is to implement a "configural component" (Siemann & Delius, 1998; Wynne, 1995, 1997). Configural models of reinforcement simultaneously consider the associative strength of isolated stimuli (the value of stimulus B), as well as the associative strength of stimuli within the context of a specific stimulus pairing (the value of stimulus B in the context of the pair BC). A purely configural model would treat each stimulus pairing as an unrelated experiment, whereas traditionally associative models would be purely "elemental" (assigning an associative strength to each stimulus in general, with no configural component). Those proposed as candidates for TI take a weighted contribution of the elemental and configural associative strengths.

Lazareva and Wasserman (2012) performed an important experiment to rigorously test several associative models of TI, including the value transfer and configural approaches. They argued that previous studies of TI had guaranteed that all adjacent pairs received equal exposure during training. However, the empirical literature had not established whether performance would be disrupted if some pairs were presented more often than others. If, for example, the pair DE was presented much more frequently than the pair BC, then the selection of D would be paired with reward delivery more often than the selection of B. For certain training scenarios, the associative models they tested predict that, when presented with a subsequent pairing of BD, subjects should favor D over B because of D's greater association with reward. In general, presenting one pair repeatedly and reinforcing the choice of one item from that pair should drive that item's "reward value" toward ceiling.

Configural models are designed to prevent this from occurring, they remain agnostic about entirely novel configural pairings. Therefore, although a correctly implemented configural model will not favor D over B as a function of massed presentation of DE, neither would such a model favor B over D. At best, response accuracy during the first presentation of BD would occur at chance. The configural model proposed by Wynne (1995) consistently fails to exceed chance even before massed trials (Lazareva & Wasserman, 2012; Figure 4), whereas the model proposed by Siemann and Delius (1998) has mixed success before massed

trials, and only consistently fails afterward (Lazareva & Wasserman, 2012; Figure 5).

To compare these model simulations to real performance, Lazareva and Wasserman first trained pigeons on four adjacent pairs in a five-item list, presenting each pair with equal frequency. They then presented only the pair DE for hundreds of trials before they assessed TI performance. They also performed computer simulations to train and test whether associative models could perform TI following "massed trials" training to one pair. The results were clear: Pigeons who received massed DE trials showed equivalent performance to other animals who were trained equally on all pairs, and both groups consistently favored B over D. However, each associative model's performance was disrupted by massed presentation of DE, with performance either reduced to chance or displaying a heavy initial bias toward choosing D.

Lazareva and Wasserman's results make clear that the associative models they tested were conceived with a narrow set of task designs in mind. TI is ordinarily trained with equalized frequencies for all stimulus pairs, but this constraint shouldn't be expected in real-world choice scenarios. It is implausible that animals evolved with perfect counterbalancing as a feature of the environment, so models of transitive inference should be resistant to unequal presentation frequencies. A similar point is made by an earlier paper (Weaver, Steirn, & Zentall, 1997) that demonstrated that value transfer models could not explain successful TI in pigeons when all pairs were trained uniformly, but terminal items yielded rewards 50% of the time.

The development of more robust models of TI requires new experiments that manipulate previously uniform task parameters. By definition, although associative models seek to maximize positive outcomes and minimize negative ones, they only come to optimal conclusions when training carefully balances exposure to all outcomes. A key benefit of representation-based learning is that representations can decouple the relationship between stimuli and their individual reward histories, and this renders their inferences robust against biased sampling.

Although humans, rhesus monkeys, and pigeons all display serial learning, they do not do so in identical ways (Scarf & Colombo, 2008). Although some studies of TI in birds have yielded TI performance comparable to rhesus monkeys (e.g., Bond, Wei, & Kamil, 2010), similar cognitive functions in birds and mammals may depend on wholly divergent brain architectures (Güntürkün & Bugnyar, 2016). Lazareva and Wasserman's demonstration that pigeon performance is not disrupted by massed trials has yet to be tested in primates.

It is also unclear from the literature how humans and monkeys might compare under mass trial manipulations. Studies of TI, including those that compare multiple primate species (examples include Merritt & Terrace, 2011; Jensen et al., 2015), ordinarily present stimulus pairs in an evenly balanced fashion. The consequences of deliberately presenting some pairs more than others remain largely unexplored. This has contributed to the difficulty in distinguishing between associative and cognitive explanations for performance. The two accounts predict different consequences of training using an uneven balance of stimulus pairs, so such manipulations should help us evaluate the relative merits of each account.

We performed an experiment in which both rhesus macaques and human participants learned 7-item lists. In both cases, training began with massed presentation of the last two list items, pair FG. Because F was always correct when paired with G, the initial reward likelihood associated with F approached 100%. After a block of FG pairs, subjects were trained on all pairs that included either F or G (to test what effect FG training had), as well as the adjacent pairs AB, BC, CD, and DE. Then, in a final phase, all 21 possible pairs were tested.

To date, no study of TI has deliberately overtrained a stimulus pair in advance of a traditional adjacent-training-all-pairs-testing task design. As such, our study is not a replication of Lazareva and Wasserman's manipulation, but rather is a complementary design investigating a related question about cognitive versus associative interpretations of TI. If association with reward was a substantial influence, then subjects should develop a bias toward choosing F. Such a bias should disrupt performance during the second phase of the experiment, subsequent learning during that phase, or both.

### Method

#### **Animal Subjects**

Data were collected from three adult male rhesus macaques (*Macaca mulatta*): *F*, H, and L. All subjects had prior experience performing transitive inference tasks, as described by Jensen and colleagues (2015). Subjects were housed at the New York Psychiatric Institute and its Department of Comparative Medicine, under the oversight of Columbia University's Institutional Animal Care and Use Committee (protocol AAAN7101). All operations were in accordance with the recommendations of the Guide for the Care and Use of Laboratory Animals of the National Institutes of Health.

#### **Animal Apparatus**

Subjects were seated in an upright primate chair during the experiment. Head movements were restrained by head post, and eye movements were recorded using a monocular scleral search coil (Judge, Richmond, & Chu, 1980; Robinson, 1963; CNC Engineering, Seattle, WA). Eye movements (saccades, followed by fixation) were used by the animal to signal which stimulus was selected during each trial. Unless otherwise indicated, the apparatus was identical to that used by Teichert Yu, and Ferrera (2014).

## **Human Participants**

Participants were 33 college undergraduates (21 women, 12 men) who earned course credit. The experiment was approved by Columbia University's Institutional Review Board (protocol AAAA7861), confirming to the guidelines for human research set forth by the American Psychological Association.

# **Human Apparatus**

Participants performed the experiment on a personal computer (iMac MB953LL/A), with responses made via a mouse-and-cursor interface.

## Procedure

At the start of the session, the task specified a list order for a set of seven novel stimuli, here denoted by the letters A through G (Figure 1, top). Subjects never observed all seven stimuli simultaneously; instead, each trial displayed a pair of list items. When



*Figure 1.* Task and experimental design. Top row, Example of the stimuli in a seven-item transitive inference list. Because B > C > D, C is a correct answer when paired with D (or any stimulus of lower rank), but is incorrect when paired with B (or any stimulus of higher rank). Bottom left, Sequence of events in a single trial. After an intertrial interval, a blue start stimulus appears. Selecting it brings up a pair of images from the implied list. If the correct item is chosen, a green check mark is displayed, followed by reward delivery. If the incorrect item is chosen, a red X is displayed, followed by a 2-s time-out. Bottom right, Design of experimental phases. In Phase 1, subjects are presented with only FG; in Phase 2, all adjacent pairs, as well as all pairs including F or G, are trained; in Phase 3, all pairs are trained. Pairs within the gray triangle were all stimuli that excluded F and G. See the online article for the color version of this figure.

stimuli were presented, a response was considered "correct" if it was made to the earlier item in this list. The structure of each trial is depicted in Figure 1 (left): Following a response to a start stimulus, subjects were then presented with the list pairing. Feedback regarding whether the response was correct (green check mark) or incorrect (red X) was given immediately. Animal subjects also received water via a tube for each correct response.

An experimental session always consisted of a fixed number of trials. Animal subjects completed 760 trials in a session, arranged into three phases (Figure 1, right). During Phase 1, the pair FG was presented 40 times. During Phase 2, 15 of the possible pairings were presented: All "adjacent pairs" (those having a symbolic distance of 1), and all other pairs that included either F or G (regardless of symbolic distance). The set of all pairings of F with a novel stimulus are hereafter denoted as "xF," and all novel pairings with G are denoted as "xG." The aim of this phase was to train the ordering of stimuli A through E relative to F and G, as well as to test performance on F-and-G-related pairs. Each of the 15 pairs was presented 20 times, resulting in a total of 300 trials. Finally, in Phase 3, all 21 pairings were presented. Again, each pair was presented 20 times, so Phase 3 lasted 420 trials. In all phases, the positions of the stimuli were randomly counterbalanced, such that the correct response was on each side of the screen half the time. Each subject completed multiple sessions (34 sessions for F, 76 sessions for H, and 20 sessions for L), with a new set of stimuli learned during each of those sessions. Monkeys performed differing numbers of sessions because of scheduling constraints.

Figure 1 (right) shows how all 21 pairs are organized with respect to two predictors. The first, symbolic distance, is wellrepresented in the literature: It is the absolute difference between stimulus ranks. The second is "joint rank," which is the sum of the stimulus ranks (here, low ranks correspond to earlier list items). For example, the pair BD has a distance of 2 and a joint rank of 6. There are several motivations for describing pairs in terms of joint rank. First, each stimulus pair corresponds uniquely with a particular combination of distance and joint rank (e.g., no pair other than BD has a distance of 2 and a joint rank of 6). This gives each pair a specific coordinate within the predictor space. Second, joint rank is strictly orthogonal to distance. The largest joint rank (FG, joint rank of 13) and the smallest (AB, joint rank of 3) both have a symbolic distance of 1. This permits distance and joint rank to be used simultaneously as predictors without a collinearity confound. The intuitive interpretation of these two descriptors is that distance corresponds to the relative dispersion of stimuli along the number line, whereas joint rank corresponds to their absolute position within the list.

Human participants experienced the same procedure, albeit with half as many trials in each condition. Thus, they experienced Phase 1 for 20 trials, Phase 2 for 150 trials, and Phase 3 for 210 trials. All learning occurred in one session.

It is important to emphasize that our procedure did not make use of "probe trials" in which a response preference was elicited without subsequently providing feedback. Throughout our procedures, a correct response was always rewarded. As such, our data cannot be divided into "training" and "test" periods. Instead, subjects learned continuously, and we sought to measure that continuous learning using a time-series analysis.

## Results

To evaluate how estimated performance changed as a function of learning, we used Gaussian process (GP) regression (Rasmussen & Williams, 2006). GP regression is a highly flexible nonlinear estimation technique that is well suited to time series analysis. It has been called a categorical analysis with an infinite number of categories, arrayed along some continuum (McElreath, 2016). One continuum of interest is time: Response accuracy for each pair changes over time. Orthogonal to time are the continua of symbolic distance and joint rank: Response accuracy is also expected to change in some way as a function of these variables.

GP regression is performed by estimating the extent to which every observation covaries with every other, given some prior metric for comparing the distance between observations. Each observation should influence other nearby observations (e.g., that occur at similar times) more than they should do so for observations that are distant. That said, observations should also influence one another as a function of their similarity (Lucas, Griffiths, Williams, & Kalish, 2015). Although such an analysis is not possible using classical methods (because of irreducible uncertainty regarding a solution for the joint contributions of distance and similarity), Bayesian methods make GP regression possible by imposing a strong prior belief that similarity and distance are related (Gelman et al., 2013). This permits robust time-series analysis to be performed without constraining the data to a particular functional form. This approach avoids difficulties encountered in past studied of TI, in which the assumptions of the regression model distort estimates of performance. For example, using logistic regression to model TI learning (Jensen et al., 2013) imposes the assumption that performance (a) increases or decreases monotonically, and (b) must reach 0.0 or 1.0 at the limit for any nonzero slope. Because monkeys tend to make errors regardless of amount of training, logistic regression is guaranteed to provide a poor estimate given a sufficiently long time series.

GP regression, contrastingly, depends on relatively few assumptions, instead allowing the data to dictate the form taken by the time series estimate. The chief constraint is that GPs are presumed to be smooth (i.e., differentiable without discontinuity). Beyond this constraint, one can imagine the model estimate as the posterior distribution of the relative density of all possible smooth curves, conditioned on the data and the informative prior. Although a full GP model can be computationally prohibitive to fit, requiring runtime  $O(n^3)$  for *n* observations, we took advantage of the "expectation propagation" approximation (Tolvanen, Jylänki, & Vehtari, 2014) implemented in the GPstuff toolbox (Vanhatalo et al., 2013) to greatly accelerate estimation. The scripts implementing this analysis in MATLAB (Version 8.4.0, The MathWorks Inc., Natick, MA, 2014) are included in the online supplemental material.

Because models were fit using Bayesian methods, the results that are reported are generally posterior probability distributions, which represent the "state of belief" of the analysis after taking the data into account. Our "credible intervals" each represent the middle X% percentile ranges in the posterior probability distribution for the estimated quantity. The interpretation of the credible intervals reported for parameters is slightly different from the more traditional "confidence intervals" invoked under the classical statistical paradigm, because credible intervals take the prior into account. We do not report p values, nor use null hypothesis significance tests, because our objective is to obtained estimated measures of performance.

We first estimated how response accuracy and RT changed over the course of the experiment, doing so independently for each stimulus pair in each phase. Responses were coded as *correct* (1) or *incorrect* (0), such that the probability of a correct response could be represented in terms of log-odds (with support from negative to positive infinity). The GP regression then fit performance in log-odds space. These model fits were converted back to probabilities using the logistic transformation, yielding a continuous probability of correct responses that was smooth with respect to time. Thus, response accuracy was modeled as a smoothly changing probability of a correct response. Reaction times were presumed to be approximately Gaussian on a log scale, and fit accordingly.

Figure 2 (top row) shows response accuracy (estimated the GP regression and subsequently averaged across subjects using a bootstrapping procedure) for monkeys choosing pairs that included either stimulus F (olive, dashed) or stimulus G (cyan, solid). During the 40 trials of Phase 1 (presentation of FG only), response accuracy on the pair FG rapidly improved, reaching approximately 90% correct selection of F. At the start of Phase 2, however, accuracy for FG dropped to 70%. Phase 2 also began with above-chance performance for all trials that paired either F or G with a novel stimulus (A thru E). This is unsurprising in the case of xG

pairings because choosing G did not yield a reward in any phase of the task. However, consistent above-chance responding on xF pairs during Phase 2 entails *avoidance* of F, despite F having yielded reliable rewards during Phase 1. For example, the first time EF was presented (Figure 2, top right, trial #41), E was chosen 60% of the time despite that it had never been previously reinforced or even presented to the subject. Similar results were obtained for all the novel stimuli. Thus, prior preference for F while it was paired with G did not prevent subjects from rapidly learning to not choose F when paired with the novel stimuli A through E.

In Phase 3, all 21 stimulus pairs were presented. By the end of Phase 3, a very pronounced symbolic distance effect was evident: Accuracy for Distance 1 pairs tended to be lowest, whereas the highest accuracy was associated with the Distance 5 pairs. Figure 2 (bottom row) shows mean response accuracy for pairs composed of the first five items in the list, A through E. These are color-coded by symbolic distance (red, solid = Distance 1, blue, dashed = Distance 2, green, dotted = Distance 3, violet, dot-dashed = Distance 4). Adjacent pair performance rose above chance in the case of AB, but remained approximately at chance levels during Phase 2 for the other pairs. However, despite low performance, a symbolic distance effect was observed at the transfer to all pairs in Phase 3.

Figure 3 presents the log RTs, following the same format as Figure 2. As in Figure 2, models were fit for each animal sepa-



*Figure 2.* Proportion of correct responses by monkeys for stimulus pairs over successive trials. Shaded areas represent 99% credible intervals. Top row: Performance on stimulus pairs that included F and/or G for Distances 1 through 5. Performance on all such pairs exceeded chance levels, and performance on xG pairs consistently exceeded that of xF pairs. Despite rapidly climbing to nearly 89.0% accuracy, the pair FG dropped immediately to 69.6% accuracy at the start of Phase 2. Bottom row: Performance on pairs composed of items A through E. Performance on adjacent pairs generally remained low during Phase 2, but improved somewhat in Phase 3. In addition, the six novel pairs in Phase 3 showed enhanced performance, consistent with a symbolic distance effect. See the online article for the color version of this figure.



*Figure 3.* Monkey reaction times (RTs; log-transformed) over successive trials. Shaded areas represent 99% credible intervals. Top row: Log RTs for pairs that included F and G for Distances 1 through 5. Reaction times were generally stable, with slightly faster times observed at larger symbolic distances. Bottom row: Log RTs for pairs composed of items A through E. Consistent RTs were observed across all such pairs. See the online article for the color version of this figure.

rately, then subsequently averaged using bootstrapping. Monkeys tended to respond very quickly,  $\exp(-1.35) \approx 0.26$  s, as is typical for oculomotor choice RTs. Although some very mild differences were evident (the pair AF evoked faster responses than the pair EF), it seems reasonable to conclude that RTs were near a lower floor for the fastest times that subjects could in principle achieve.

Figure 4 presents mean response accuracy for human participants. Unlike monkeys, participants acquired a stronger preference for F during Phase 1. Despite this, they immediately discounted F at the beginning of Phase 2. They then showed rapid acquisition, racing to ceiling on all xF and xG pairs. Performance also rapidly rose above chance levels for adjacent pairs in Phase 2, and a symbolic distance effect was not evident. If anything, performance on adjacent pairs tended to slightly exceed the Distance 2 pairs.

Figure 5 presents the log RTs of participants. Humans responded much more slowly than monkeys,  $exp(0.5) \approx 1.65$  s, but unlike monkeys they became systematically faster over the course of the session. Signs of a symbolic distance effect were more evident in the human RTs than in their response accuracy. In particular, participants tended to respond more quickly to larger symbolic distances.

The transition at the start of Phase 2 shows several surprising effects with respect to F. At the start of Phase 2, monkeys appeared to consistently avoid F, despite consistent reinforcement of its selection in Phase 1. In addition, monkey response accuracy on FG appeared to drop abruptly on a single trial, from near 90% at the end of Phase 1 to near 70% at the start of Phase 2. Humans showed neither of these effects. Their FG trials remained close to ceiling, whereas all other xF pairs appeared to begin near chance (albeit

with considerable uncertainty). The patterns of responding to F suggest different processes at work in these two species.

To get a more precise understanding of this transition, we examined choices based on order of appearance at this transition, as shown in Figure 6. The first four points report the estimated response accuracy for FG during the last four trials of Phase 1 for both monkeys (Figure 6, left) and humans (Figure 6, right). The next six points in each plot report the *n*th instance in Phase 2 of FG (black, dotted), any xF pair (olive, dashed), or any xG pair (cyan, dot-dashed). For example, the first xF pair in Phase 2 might be the pair AF or BF and might not take place until the third or fourth trial. Thus, the first xF trial in Phase 2 is the first unambiguous case where some other stimulus should be favored over F.

Figure 6 (left) shows that monkeys chose F less than half the time on the very first xF trial. It also shows that the drop in FG performance between phases is indeed an abrupt step downward. Contrastingly, humans were about 80% likely to choose F during the first xF pairing, and about 20% likely to choose G during the first xG pairing. This preference for F was very rapidly corrected. Thus, whereas monkeys had a prior expectation that F should be avoided in favor of novel stimuli (because of their prior experience with the task), humans had no such expectation and thus had to instead rapidly "unlearn" the associated value of F.

A "novel stimulus bias" might help explain why subjects tend to succeed at xF and xG pairings. To test for such a bias without a reward confound, we examined only trials (a) in which stimuli A, B, C, D, or E were presented for the first time at the start of Phase 2, and b in which another of those five stimuli had previously been presented. Monkeys appeared to display a bias toward the novel



*Figure 4.* Proportion of correct responses by humans for stimulus pairs over successive trials. Shaded areas represent 99% credible intervals. Top row: Performance on stimulus pairs that included F and/or G for Distances 1 through 5. Performance on FG rapidly rose to ceiling levels, as did performance on pairs including G. Other pairs including F began at chance level at the start of Phase 2, but also rapidly rose to ceiling. Bottom row: Performance on pairs composed of items A through E. Accuracy on adjacent pairs rose steadily on Phases 2 and 3, and novel nonadjacent pairs were above chance (but below adjacent pairs) at the start of Phase 3. See the online article for the color version of this figure.

stimulus in the pairing (179 out of 293 trials across all three monkeys), p(novel) = 0.611, 99% confidence interval (CI) 0.534 to 0.682, as did humans (74 out of 125 trials across participants), p(novel) = 0.592, 99% CI 0.477 to 0.699. Therefore, given two relatively novel stimuli, both species tended to favor the stimulus that was entirely novel. Monkeys may have used the same strategy in xF pairings, in which they favored novel stimuli at about the same rate (as indicated in Figure 6).

On the other hand, however, monkeys may have avoided F because "task representation" (Stoet & Snyder, 2003), reflecting an understanding of the task structure, accrued over the course of learning many lists. To investigate this possibility, we performed a logistic regression for each subject, consisting of only trials from the first 30 trials of Phase 2 across all lists. Response accuracy was evaluated in terms of an intercept and two predictors: trial number within a list, and the order in which lists were learned. The intercept in this analysis estimated response accuracy on the first trial of Phase 2 in the first list. At the start of the study, subject F responded to initial xF trials at chance levels, p(correct | xF) =0.488, 99% CI 0.428 to 0.548. However, subject H responded above chance,  $p(\text{correct} \mid \text{xF}) = 0.717, 99\%$  CI 0.678 to 0.754, as did subject L,  $p(\text{correct} \mid \text{xF}) = 0.639, 99\%$  CI 0.564 to 0.70. All three subjects consistently favored the initial novel stimulus during their last session, subject F: p(correct | xF) = 0.572, 99% CI 0.512to 0.630; subject H:  $p(\text{correct} \mid \text{xF}) = 0.711, 99\%$  CI 0.671 to 0.749; subject L:  $p(\text{correct} \mid \text{xF}) = 0.615, 99\%$  CI 0.538 to 0.68. The analysis is included in the online supplemental material.

GP regression is a continuous time-series analysis that permits estimates of not only response accuracy, but also the rate at which performance changes. Figure 7 shows the rate at which monkeys improved at the start of Phases 2 and 3. These were measured by converting the probabilities in Figure 2 to log-odds, then obtaining the average slope of the resulting curves during the first block of 30 trials. At the start of Phase 2 (Figure 7, left), monkeys displayed an apparent symbolic distance effect with respect to the *rate* of learning, an effect not previously reported in the literature. In addition, learning was generally faster for xF pairs than it was for xG pairs, or for the entirely novel pairs (AB, BC, DC, DE). By the start of Phase 3 (Figure 7, right); however, learning rates were consistently slow. Only novel pairs had a learning rate that excluded zero from the 99% credible interval.

Figure 8 shows the corresponding learning rate analysis for human participants, which were very different than those of monkeys. First, human learning rates were about an order of magnitude faster. Second, no symbolic distance effect was evident. Instead, xF pairs were learned very rapidly at the start of Phase 2 (consistent with Figure 6) and had stabilized by the start of Phase 3. Another difference was that novel pairs in Phase 3 seemed to be learned faster when paired with stimulus close to A. This is likely an example of a terminal item effect, which is only seen at one end of the list because the other end has already received considerable training.

To better understand how the symbolic distance effect and the terminal item effect might manifest in these data, a more



*Figure 5.* Human reaction times (RTs; log-transformed) over successive trials. Shaded areas represent 99% credible intervals. Top row: Log RTs for pairs that included F and G for Distances 1 through 5. Participants were much slower than monkeys, but RTs steadily improved over the course of the experiment. Bottom row: Log RTs for pairs composed of Items A through E. Reaction times to pairs that included Item A sped up most during learning, whereas pairs closer to the "middle" of the list (such as DE) sped up very little. See the online article for the color version of this figure.

complex GP model combined several predictors. This permitted estimates to be fit for all stimulus pairs. Trial number was used as a predict to make a precise inference of performance during the first trial of Phase 3. In addition, each pair was coded for its symbolic distance and its joint rank. As noted previously, the orthogonality of distance and joint rank permits the two to be used simultaneously as predictors across all pairs. In addition, estimates may also be obtained for the interpolated regions



*Figure 6.* Probability of choosing F and G during a pair's last four presentations during Phase 1 and first six presentations during Phase 2. Whiskers represent the 99% credible interval for the parameter, and boxes represent the 80% credible interval. The probability of choosing F in the context of the pair FG (black, dotted) is contrasted against the probability of choosing F in the context of all other xF pairs (olive, dashed). Also shown is the probability of choosing G in the context of all other xG pairs (cyan, dot-dashed). Left: Mean selection probability for rhesus macaques. Monkeys show a drop in performance between the last presentation of FG in Phase 1 and its first presentation in Phase 2. In addition, monkeys avoid choosing F even on the first xF pairing. Right: Mean selection probability for human participants. Although humans avoid G in xG pairings, they show a brief preference for F during the first presentation of xF in Phase 2. This preference has been fully reversed after only four presentations. See the online article for the color version of this figure.



*Figure 7.* Estimates of the rate of change of response accuracy for monkeys, measured in log-odds units of probability per trial. Pairs are sorted by symbolic distance and by joint rank. Boxes represent 80% credible intervals, while whiskers represent 99% credible intervals. Individual pairs are color-coded identically to Figures 2–5. Left: Learning rate at the start of Phase 2. Pairs including stimulus F tended to show faster improvement, as did pairs with greater symbolic distance. Right: Learning rate at the start of Phase 3. Novels pairs showed elevated learning, but all learning rates were low compared with Phase 2. See the online article for the color version of this figure.

between stimuli, even though no stimulus pairing resides in those regions. This facilitates understanding of the underlying estimated function.

Figure 9 shows the estimated probability of a correct response (top row) and the log RT (bottom row) for all monkeys at the start of Phase 3. The color coding of each of the pairs is consistent with that used in Figures 2 through 5, and the novel pairs at this point in the experiment are AC, BD, CE, AD, BE, and AE. Both response accuracy and RT show a consistent distance effect in both familiar and unfamiliar pairs, with larger

symbolic distances associated with higher accuracy and lower RTs. The added benefit of previously training (in Phase 2) the pairs including F and G is evident in the asymmetry of these curves. The case for a terminal item effect in the monkey data are mixed, however. While pairs including the terminal A and G appeared to yield slightly higher accuracy, log RTs seemed closer to linear, with FG being the slowest pair despite being the most extensively trained.

Figure 10 depicts response accuracy and RTs for human participants at the start of Phase 3. As in Figures 4 and 5, accuracy was



*Figure 8.* Estimates of the rate of change of response accuracy for humans, measured in log-odds units of probability per trial. Pairs are sorted by symbolic distance and by joint rank. Boxes represent 80% credible intervals, while whiskers represent 99% credible intervals. Individual pairs are color-coded identically to Figures 2–5. Left: Learning rate at the start of Phase 2. Pairs including stimulus F tended to show much faster improvement. Right: Learning rate at the start of Phase 3. Novels pairs showed elevated learning, and most learning rates were higher than those observed in monkeys. See the online article for the color version of this figure.

Monkey Performance (Phase 3 Start)



*Figure 9.* Estimates of response accuracy (top row) and log reaction times (RTs; bottom row) in monkeys during the first trial of Phase 3. Pairs are sorted by symbolic distance and by joint rank. Gray shaded areas correspond to the 99% credible interval, interpolated between items. Individual pairs are color-coded identically to Figures 2–5 and also include a horizontal bar to denote 80% credible intervals. Response accuracy displayed both a symbolic distance effect and a terminal item effect, as well as asymmetry resulting from the additional training on stimuli F and G. Log RTs displayed a symbolic distance effect, but did not show a clear terminal item effect. See the online article for the color version of this figure.

consistently higher, and RTs slower, than in monkeys. In addition, however, a much clearer pattern of terminal item effects was evident (especially for the pair AB). Although a clear distance effect was evident in the RTs, the evidence for such an effect on accuracy was much more equivocal. As in Figure 4, performance on adjacent pairs tended to marginally exceed that of corresponding Distance 2 pairs. However, Distance 4 pairs tended to outperform Distance 3 pairs, and Distance 3 pairs tended to outperform Distance 2 pairs, suggesting a mild distance effect among the six novel pairs. The effect of massed training of FG was also evident in human RTs, such that xF and xG RTs were consistently faster than other stimulus pairs.

Data and analyses are provided in the online supplemental material.

#### Discussion

Learning and decision-making are thought to depend on an ability to predict the reward value of alternative actions. Reinforcementbased theories critically depend on the history of experienced rewards and seek to explain choices made by nonhuman animals in terms of associations between stimuli, actions, and outcomes. As an alternative, cognitive factors appear to also contribute to animal decision making, whereby representational mechanisms both facilitate and constrain an organism's ability to infer relationships between stimuli and outcomes. The TI paradigm can help disambiguate these contributions, because explaining TI performance in terms of stimulusreward associations is especially difficult. Our results support cognitive interpretations of TI performance.

Modifying the manipulation introduced by Lazareva and Wasserman (2012), we investigated whether TI performance was disrupted by additional training on a particular stimulus pair. According to associative models, when the pair FG is overtrained, it should drive the association of stimulus F with reward toward ceiling. Our procedure differed from that of Lazareva and Wasserman (2012) in several ways: We presented massed trials at the outset rather than after learning, we always rewarding correct responses, and we did not use correction trials. If TI depends on stimulus-reward associations, the reward history of F should lead to it being favored in pairings such as EF and DF. Even though we repeatedly presented FG trials in advance of ordinary TI learning in both rhesus monkeys and humans, we found no evidence of a Human Performance (Phase 3 Start)



*Figure 10.* Estimates of response accuracy (top row) and log reaction times (RT; bottom row) in humans during the first trial of Phase 3. Pairs are sorted by symbolic distance and by joint rank. Gray shaded areas correspond to the 99% credible interval, interpolated between items. Individual pairs are color-coded identically to Figures 2–5 and also include a horizontal bar to denote 80% credible intervals. Human response accuracy showed a pronounced terminal item effect, but a weak symbolic distance effect. Human RTs, on the other hand, displayed both effects. See the online article for the color version of this figure.

bias for choosing F in subsequent phases of learning. This result is consistent with the lack of disrupted learning in pigeons reported by Lazareva and Wasserman (although see Lazareva, Kandray, & Acerbo, 2015, for a case in which massed DE trials disrupted BD performance in a subset of subjects).

Following FG training, rhesus macaques displayed a surprising pattern of behavior. Although novel stimuli paired with G were more likely to be selected than novel stimuli paired with *F*, all xF and xG pairings yielded performance above chance. In other words, despite a learning history in which F had only ever yielded reward, subjects systematically *avoided* F in favor of novel stimuli. This result, evident in GP estimates of response accuracy (Figure 2), was also found in the accuracy of the very first xF trial of Phase 2 (Figure 6). Every session was performed with novel stimuli, so subjects could not have had a prior reward association for any specific stimulus in the F position.

This result may reflect a systematic bias toward novel stimuli. This is consistent with the finding that rhesus macaques find novel stimuli to be more salient, and make biased eye movements toward them (Ghazizadeh, Griggs, & Hikosaka, 2016). Because we provided massed presentations of FG, this bias was functional. However, a second explanation, not mutually exclusive with the first, is that subjects developed a task representation (Stoet & Snyder, 2003), sometimes called a "task set" (Sakai, 2008). Despite having never seen stimulus F before the current session, subjects had learned that stimuli *like* F (i.e., those sharing the characteristic of having appeared during the initial massed trials) were likely to be of lower rank than novel stimuli seen later in the session.

The present results are consistent with monkeys having benefitted from a task representation, but do not on their own rule out a novelty account. A logistic analysis of response accuracy in terms of both trial number and session number yielded mixed results: Two monkeys began at chance on xF pairs, whereas the third began at chance but increasingly avoided F as session number increased. A satisfactory dissociation of the task representation hypothesis from the novelty hypothesis is beyond the scope of this study, because a subject pool of three is too small to make strong statements about the naïve performance. The current experiment also cannot distinguish the relative influence of a combination of the two factors.

Another surprising effect was that despite achieving approximately 90% accuracy overall during FG training, accuracy on the 242

pair dropped to around 70% as soon as other stimulus pairs began to be presented. This effect was evident even when comparing, on a session-by-session basis, the last FG trial of Phase 1 to the first FG trial of Phase 2. While this could potentially be another manifestation of a task representation, another possibility is that the maintenance of a stimulus pair and inferential reasoning both draw upon the same working memory resources (Halford, Cowan, & Andrews, 2007). This is consistent with past studies that have found that rhesus monkeys perform only slightly above chance on adjacent pairs after hundreds of training trials (Jensen et al., 2013, 2015). Despite only needing to learn the correct response for 6 adjacent pairs, monkeys seemed unable to memorize the correct choice for each pair. Past estimates of the capacity of visual short-term memory (STM) suggest that rhesus macaques are limited to between one and two items (Elmore et al., 2011). Thus, poor performance on pairings that animals have seen many times over may reflect the unavailability of memorization strategies, leaving more implicit mechanisms to perform the necessary inferences. The present result (wherein FG is selected with high accuracy only when presented in a block of identical trials) suggests that performance in Phase 1 was facilitated by a memory trace of the previous trial.

A traditional five-item TI task was embedded in our design, with subjects being trained on adjacent pairs AB, BC, CD, and DE during Phase 2 before being tested on all pairs in Phase 3. The pattern of transfer displayed by the monkeys resembles that seen in previous experiments (Jensen et al., 2015). Phase 2 was unusual, however, in its continued inclusion of all xF and xG pairs, in addition to the adjacent pairs. Although there is no logical reason why exposure to these pairs should facilitate or inhibit learning of the order of A through E (since neither F nor G were informative about the relative positions of the remaining stimuli), it is worth considering whether some interaction might nevertheless have occurred. A possible future experimental design might be to split Phase 2 into two periods: An xF/xG only period (to test the effect of massed trials) and an A-through-E period (in anticipation of testing in the final phase). In particular, training A-through-E before presenting xF/xG pairs could provide a way to disentangle the novelty/task set ambiguity of the present study.

Our human participants displayed a different pattern of behavior from the monkeys. Superficially, this pattern more closely resembled the expectations from associative models: Participants favored F on the very first xF pairing, whereas they avoided G on the first xG pairing. Humans adapted very rapidly, however, correctly avoiding F by the fourth xF presentation. The speed of human adaptation to xF pairs was much faster than any other learning displayed in this study, which is especially notable given that participants were task-naïve. Associative models with very fast learning rates are prone to unstable behavior, on account of their oversensitivity to strings of good and bad luck. Furthermore, the humans' very rapid learning did not appear to be a feature of the transition from Phase 2 to Phase 3, as would be expected if rapid learning was a mere consequence of an associative learning rate.

In both the monkey and human cases, the stimuli A through E effectively offered a second, embedded TI task, with adjacent-pair training during Phase 2 and all-pair testing during Phase 3. Humans and monkeys both successfully inferred the relationships between nonadjacent test pairs from this set with no evidence of disruption because of overtraining of F. This embedded task also permitted measurement of bias for novel stimuli, which humans displayed at rates comparable to that of the monkeys. In the case of xF, the reward history of F appears to have mitigated or overridden that bias on the first trial.

Overall, the present human and monkey results display both consistencies and differences with Lazareva and Wasserman's pigeon data. On the one hand, FG training did not disrupt learning of stimuli A through E. Monkeys never favored F, despite initial rewards earned from FG. They appeared to exploit a task representation to anticipate F's undesirability relative to novel stimuli. Beyond that anticipation, their performance otherwise displays classical symbolic distance and terminal item effects. One of the most curious effects was a drop in the accuracy of FG choice at the start of Phase 2.

Human behavior at the start of Phase 2 did not resemble that of monkeys. Participants very briefly favored F, but they overcame their initial bias almost immediately. They also did not display drops in response accuracy to any pairings following phase transitions.

It would be inappropriate to compare monkey and human performance directly in this study. The two species made responses using a different modality (eye movements for the monkeys vs. mouse clicks for the humans). Past human work suggests that humans responded similarly in binary choice paradigms across different response modalities (Szumska, van der Lubbe, Grzeczkowski, & Herzog, 2016), but it is unknown whether monkeys similarly generalize. Training also differed with respect to their overall experience with the task (monkeys completing dozens of sessions with different lists vs. a single session for humans). Human performance would very likely look different if participants had learned at least 20 lists apiece. A matched study, in which humans and nonhumans experience identical procedures, would be a worthy design for a future study. Nevertheless, if either species had relied on a stimulus-reward-based associative strategy, then the high rate of reward for stimulus F should have yielded a preference for F.

#### **Time-Series Analysis of TI**

Traditionally, studies of TI train subjects extensively on adjacent pairs, then "test" them using nonadjacent pairs. The reported duration of training in such studies is not consistent from one article to the next. Many labs end training based on a performance criterion, so subjects may experience training periods of dramatically different lengths. Because the aim of a behavioral test is to assess performance during a cross-section in time, all that can be ascertained during test trials is response accuracy and RT at the moment of the test.

In this study, we instead approached TI performance as a time series, with the expectation that subjects would continue to accrue information over the course of successive trials, and that performance would improve accordingly. Using GP regression, we could construct a much more detailed picture of how learning unfolded for each stimulus pair in both species.

Rather than report performance as a function of stimulus pair, Figure 7 depicts the rate of change for monkey response accuracy, doing so in terms of log-odds units of probability. As far as we are aware, the only method previously used to measure learning rate in studies of TI has been "trials to criterion" (e.g., Gazes et al., 2012), making ours the first study to examine the first derivative (i.e., instantaneous change) of response accuracy directly. Figure 7 reveals an intriguing possibility: A symbolic distance effect for learning rate. The implication is that, when measured in terms of log-odds, stimulus pairs separated by larger symbolic distances are not only expected to have higher accuracy, but *also* to improve more rapidly as additional training unfolds. Realistically, this effect cannot last: As performance reaches ceiling, learning rates drop to zero. Nevertheless, there is no evidence that prior overtraining of F impaired the rate at which monkeys learned to avoid F.

Figure 8 depicts the rate of change for human performance. The difference in scale compared with monkey learning is striking, with the mean learning rate as much as 10 times faster for humans. The pattern of learning was qualitatively different between the two species: At the start of Phase 2, humans made steady progress on most pairs (including all novel adjacent pairs), and also made very rapid progress on all xF pairs. Although it is intuitively obvious that humans and monkeys differ in their rate of learning (cf. visual comparisons of performance in Figures 2–4), the analysis of the learning rate provides a more quantitative account of that difference. This difference is even more striking given the lack of experience humans had with the task relative to the monkeys, a reflection of the "profound discontinuity" between human and nonhuman cognition described by Penn, Holyoak, and Povinelli (2008).

Comparisons of learning rates also provide a new way to compare the predictions of different models of TI. This, in turn, can be used to place much more severe constraints on models that seek to explain behavior. Additional studies will be needed to determine how learning rate effects manifest in different species and in various experimental preparations.

### Task Awareness and Task Representation

Task "awareness" has been a major focus in the human study of TI (Greene, Spellman, Dusek, & Eichenbaum, 2001; Lazareva & Wasserman, 2010; Martin & Alsop, 2004), and it has been proposed as an explanation of the difference between human and monkey TI performance (Moses, Villate, & Ryan, 2006). In most cases, task awareness is defined in terms of verbal report: If a participant can deduce (and subsequently report) that the stimuli belong to an ordered hierarchy, they are considered to be "aware;" all others are "unaware." Those participants who could not verbally articulate the structure of the task also perform poorly on the inferential test trials (Frank, Rudy, Levy, & O'Reilly, 2005; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005; however, see Greene et al., 2001; Siemann & Delius, 1996).

Unfortunately, the way awareness has been defined does little to clarify the underlying cognitive mechanisms of TI. Although many reports link awareness and successful inference for test pairs in humans (Lazareva & Wasserman, 2010; Libben & Titone, 2008), it is unclear whether the two can be dissociated and, if so, which precedes the other. It is also unclear how to reconcile these findings with reports of TI in unaware humans (Greene et al., 2001; Siemann & Delius, 1996;). Moses, Villate, Binns, Davidson, and Ryan (2008) argue that TI depends not on a single brain region or cognitive faculty, but rather on "cognitive integrity" across multiple subsystems.

While it is true that monkeys differ from "aware" humans in a variety of respects, they also differ from "unaware" humans in an important way: the monkeys display transitive inference, whereas unaware humans, at least according to most reports, do not. This raises an important question: If a human's task awareness is linked to making successful transitive inferences, are analogous mechanisms involved in nonhuman TI? The "awareness" label in the human literature relies on verbal report, so macaques by definition do not qualify. Nevertheless, even if macaques have a subjective experience of TI that falls far short of human experience, it may be that some measurable distinction exists between their experience and the total absence of awareness. A range of theoretical constructs regarding animal cognition have been proposed whose functions may overlap with the task awareness reported in humans.

The observed above-chance performance on xF pairs, in which animals avoided F despite its reward history, is consistent with the claim that monkeys benefited from some "task representation" (Sakai, 2008; Stoet & Snyder, 2003). Task representations, in this context, operate as superordinate frameworks within which current learning is embedded. Thus, as subjects learn how to respond to the current session's pair FG, they may also learn how F and G fit into a broader context learned during previous sessions. However, above-chance xF performance is also consistent with a bias toward novel stimuli, which may or may not be innate. One way to distinguish novelty-seeking from task representation would be to replicate Lazareva & Wasserman's (2012) design directly, because their result does not suffer from a stimulus novelty confound. Another, noted previously, is to provide some training of A-through-E pairs without providing a transitive link to either F or G.

The comparative study of task representations has yet to reach a consensus regarding either terminology or theoretical focus. The study of how animals learn overarching task demands has also been described as "rule representation" (Bunge & Wallis, 2008), "abstract concept learning" (Wright & Katz, 2007), or simply "task context" (Asaad, Rainer, & Miller, 2000). Although the evidence suggests representations more abstract than those proposed by associative models, the available evidence falls considerably short of a demonstration of "task awareness," which as defined in humans is explicitly metacognitive (Terrace & Son, 2009). Although several lines of behavioral evidence are consistent with hypothesized metacognitive faculties in animals (Smith, Couchman, & Beran, 2014), whether metacognition can be inferred from behavior alone may be formally undecidable (Clark & Hassert, 2013).

As a consequence, building a bridge between task awareness and task representation will require a better understanding of the implicit mechanisms of human task representation. When the evidence is approached from this perspective, there are good reasons to think that humans differ from macaques with respect to how task representations influence behavior. Humans, for example, are adept at "task switching" paradigms. These require participants to make choices according to one of several possible rules (Kiesel et al., 2010). When faced with task switching, rhesus macaques have difficulty suppressing the inappropriate rule. Instead, both behavioral (Avdagic, Jensen, Altschul, & Terrace, 2014) and neurophysiological (Klaes, Westendorff, Chakrabarti, & Gail, 2011) evidence suggest that macaques consider both rules simultaneously, resulting in interference. One pair of studies makes this species difference especially clear. Stoet and Snyder (2003) demonstrated that, over the course of tens of thousands of trials, macaques had difficulty using a cue to switch from one task representation to another. Humans performed an identical procedure for an equally lengthy training period, and never displayed interference between task sets (Stoet & Snyder, 2007).

In this spirit, future comparative studies of TI should place greater emphasis on exploring how task representations are constructed, maintained, and deployed by different species. The first step in doing so is to study how task representation contributes to TI performance in humans without falling back on self-report.

#### Conclusion

Initial overtraining that repeatedly associated a single stimulus with reward did not impair monkeys' or humans' ability to avoid choosing that stimulus when appropriate in the context of TI learning. Instead, both species used information gleaned from the task to improve their performance, albeit in different ways. Monkeys performed above chance with respect to F and G being paired with other stimuli, either because of an innate bias for novelty or their prior experience with the task. Humans, on the other hand, were task-naïve and showed an initial preference for F which was almost immediately replaced with avoidance. This initial preference for F was present despite an apparent bias toward novel stimuli. In addition to performance, our GP regression permitted an analysis of the learning rate, which differed dramatically in monkeys and humans. Although the relative association of stimuli with reward is a useful strategy under some circumstances, it does not appear to be how transitive inference is performed. Instead, cognitive models that relate stimuli to one another along some ordinal dimension are needed to account for the full range of published TI results.

#### References

- Allen, C. (2006). Transitive inference in animals: Reasoning or conditioned associations? In S. Hurley & M. Nudds (Eds.), *Rational animals*? (pp. 175–185). New York, NY: Oxford University Press.
- Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology*, 84, 451–459.
- Avdagic, E., Jensen, G., Altschul, D., & Terrace, H. S. (2014). Rapid cognitive flexibility of rhesus macaques performing psychophysical task-switching. *Animal Cognition*, 17, 619–631. http://dx.doi.org/10 .1007/s10071-013-0693-0
- Bond, A. B., Wei, C. A., & Kamil, A. C. (2010). Cognitive representation in transitive inference: A comparison of four corvid species. *Behavioural Processes*, 85, 283–292. http://dx.doi.org/10.1016/j.beproc.2010.08.003
- Bunge, S. A., & Wallis, J. D. (Eds.). (2008). Neuroscience of rule-guided behavior. New York, NY: Oxford University Press.
- Burt, C. (1911). Experimental tests of higher mental processes and their relation to general intelligence. *Journal of Experimental Pedagogy*, 1, 93–112.
- Clark, K. B., & Hassert, D. L. (2013). Undecidability and opacity of metacognition in animals and humans. *Frontiers in Psychology*, 4, 171. http://dx.doi.org/10.3389/fpsyg.2013.00171
- Couvillon, P. A., & Bitterman, M. E. (1992). A conventional conditioning analysis of "transitive inference" in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 18, 308–310. http://dx.doi.org/ 10.1037/0097-7403.18.3.308

- D'Amato, M. R., & Colombo, M. (1990). The symbolic distance effect in monkeys (*Cebus apella*). Animal Learning & Behavior, 18, 133–140. http://dx.doi.org/10.3758/BF03205250
- Elmore, L. C., Ma, W. J., Magnotti, J. F., Leising, K. J., Passaro, A. D., Katz, J. S., & Wright, A. A. (2011). Visual short-term memory compared in rhesus monkeys and humans. *Current Biology*, 21, 975–979. http://dx.doi.org/10.1016/j.cub.2011.04.031
- Frank, M. J., Rudy, J. W., Levy, W. B., & O'Reilly, R. C. (2005). When logic fails: Implicit transitive inference in humans. *Memory & Cognition*, 33, 742–750. http://dx.doi.org/10.3758/BF03195340
- Gazes, R. P., Chee, N. W., & Hampton, R. R. (2012). Cognitive mechanisms for transitive inference performance in rhesus monkeys: Measuring the influence of associative strength and inferred order. *Journal of Experimental Psychology: Animal Behavior Processes*, 38, 331–345. http://dx.doi.org/10.1037/a0030306
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). New York, NY: CRC Press.
- Ghazizadeh, A., Griggs, W., & Hikosaka, O. (2016). Ecological origins of object salience: Reward, uncertainty, aversiveness, and novelty. *Frontiers in Neuroscience*, 10, 378. http://dx.doi.org/10.3389/fnins.2016 .00378
- Greene, A. J., Spellman, B. A., Dusek, J. A., Eichenbaum, H. B., & Levy, W. B. (2001). Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans. *Memory & Cognition*, 29, 893–902. http://dx.doi.org/10.3758/BF03196418
- Güntürkün, O., & Bugnyar, T. (2016). Cognition without cortex. Trends in Cognitive Sciences, 20, 291–303. http://dx.doi.org/10.1016/j.tics.2016 .02.001
- Halford, G. S., Cowan, N., & Andrews, G. (2007). Separating cognitive capacity from knowledge: A new hypothesis. *Trends in Cognitive Sciences*, 11, 236–242. http://dx.doi.org/10.1016/j.tics.2007.04.001
- Jensen, G. (2017). Serial learning. In J. Call, G. M. Burghardt, I. M. Pepperberg, C. T. Snowden, & T. Zentall (Eds.), APA Handbook of comparative cognition: Perception, learning, and cognition (Vol. 2 pp. 385–409). Washington, DC: American Psychological Association. http://dx.doi.org/10.1037/0000012-018
- Jensen, G., Altschul, D., Danly, E., & Terrace, H. (2013). Transfer of a serial representation between two distinct tasks by rhesus macaques. *PLoS ONE*, 8, e70285. http://dx.doi.org/10.1371/journal.pone.0070285
- Jensen, G., Muñoz, F., Alkan, Y., Ferrera, V. P., & Terrace, H. S. (2015). Implicit value updating explains transitive inference performance: The betasort model. *PLoS Computational Biology*, *11*, e1004523. http://dx .doi.org/10.1371/journal.pcbi.1004523
- Judge, S. J., Richmond, B. J., & Chu, F. C. (1980). Implantation of magnetic search coils for measurement of eye position: An improved method. *Vision Research*, 20, 535–538. http://dx.doi.org/10.1016/0042-6989(80)90128-5
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching: A review. *Psychological Bulletin*, 136, 849–874. http://dx.doi.org/10 .1037/a0019842
- Klaes, C., Westendorff, S., Chakrabarti, S., & Gail, A. (2011). Choosing goals, not rules: Deciding among rule-based action plans. *Neuron*, 70, 536–548. http://dx.doi.org/10.1016/j.neuron.2011.02.053
- Lazareva, O. F., Kandray, K., & Acerbo, M. J. (2015). Hippocampal lesion and transitive inference: Dissociation of inference-based and reinforcement-based strategies in pigeons. *Hippocampus*, 25, 219–226. http://dx.doi.org/10.1002/hipo.22366
- Lazareva, O. F., & Wasserman, E. A. (2010). Nonverbal transitive inference: Effects of task and awareness on human performance. *Behavioural Processes*, 83, 99–112. http://dx.doi.org/10.1016/j.beproc.2009.11.002
- Lazareva, O. F., & Wasserman, E. A. (2012). Transitive inference in pigeons: Measuring the associative values of Stimuli B and D. *Behav*-

ioural Processes, 89, 244-255. http://dx.doi.org/10.1016/j.beproc.2011 .12.001

- Libben, M., & Titone, D. (2008). The role of awareness and working memory in human transitive inference. *Behavioural Processes*, 77, 43– 54. http://dx.doi.org/10.1016/j.beproc.2007.06.006
- Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic Bulletin & Review*, 22, 1193–1215. http://dx.doi.org/10.3758/s13423-015-0808-5
- Martin, N., & Alsop, B. (2004). Transitive inference and awareness in humans. *Behavioural Processes*, 67, 157–165. http://dx.doi.org/10.1016/ j.beproc.2004.03.017
- McElreath, R. (2016). Adventures in covariance. In *Statistical rethinking* (pp. 410–419). New York, NY: CRC Press.
- McGonigle, B. O., & Chalmers, M. (1977). Are monkeys logical? *Nature*, 267, 694–696. http://dx.doi.org/10.1038/267694a0
- McGonigle, B. O., & Chalmers, M. (1992). Monkeys are rational! *Quarterly Journal of Experimental Psychology*, 45, 189–228.
- Merritt, D. J., & Terrace, H. S. (2011). Mechanisms of inferential order judgments in humans (*Homo sapiens*) and rhesus monkeys (*Macaca mulatta*). Journal of Comparative Psychology, 125, 227–238. http://dx .doi.org/10.1037/a0021572
- Moses, S. N., Villate, C., Binns, M. A., Davidson, P. S. R., & Ryan, J. D. (2008). Cognitive integrity predicts transitive inference performance bias and success. *Neuropsychologia*, 46, 1314–1325. http://dx.doi.org/ 10.1016/j.neuropsychologia.2007.12.009
- Moses, S. N., Villate, C., & Ryan, J. D. (2006). An investigation of learning strategy supporting transitive inference performance in humans compared to other species. *Neuropsychologia*, 44, 1370–1387. http://dx .doi.org/10.1016/j.neuropsychologia.2006.01.004
- Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31, 109–130. http://dx.doi.org/10.1017/ S0140525X08003543
- Rasmussen, C. E., & Williams, C. K. I. (2006). Gaussian processes for machine learning. Cambridge, MA: MIT Press.
- Robinson, D. A. (1963). A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Transactions on Biomedical Electronics*, 10, 137–145. http://dx.doi.org/10.1109/TBMEL .1963.4322822
- Sakai, K. (2008). Task set and prefrontal cortex. Annual Review of Neuroscience, 31, 219–245. http://dx.doi.org/10.1146/annurev.neuro.31 .060407.125642
- Scarf, D., & Colombo, M. (2008). Representation of serial order: A comparative analysis of humans, monkeys, and pigeons. *Brain Research Bulletin*, *76*, 307–312. http://dx.doi.org/10.1016/j.brainresbull.2008.02 .022
- Siemann, M., & Delius, J. D. (1996). Influences of task concreteness upon transitive responding in humans. *Psychological Research*, 59, 81–93. http://dx.doi.org/10.1007/BF01792429
- Siemann, M., & Delius, J. D. (1998). Algebraic learning and neural network models for transitive and non-transitive responding. *Journal of Cognitive Psychology*, 10, 307–334. http://dx.doi.org/10.1080/713 752279
- Smith, C., & Squire, L. R. (2005). Declarative memory, awareness, and transitive inference. *The Journal of Neuroscience*, 25, 10138–10146. http://dx.doi.org/10.1523/JNEUROSCI.2731-05.2005

- Smith, J. D., Couchman, J. J., & Beran, M. J. (2014). Animal metacognition: A tale of two comparative psychologies. *Journal of Comparative Psychology*, *128*, 115–131. http://dx.doi.org/10.1037/a0033105
- Stoet, G., & Snyder, L. H. (2003). Executive control and task-switching in monkeys. *Neuropsychologica*, 41, 1357–1364. http://dx.doi.org/10.1016/ S0028-3932(03)00048-4
- Stoet, G., & Snyder, L. H. (2007). Extensive practice does not eliminate human switch costs. *Cognitive, Affective & Behavioral Neuroscience*, 7, 192–197. http://dx.doi.org/10.3758/CABN.7.3.192
- Szumska, I., van der Lubbe, R. H., Grzeczkowski, L., & Herzog, M. H. (2016). Does sensitivity in binary choice tasks depend on response modality? *Consciousness and Cognition*, 43, 57–65. http://dx.doi.org/ 10.1016/j.concog.2016.05.005
- Teichert, T., Yu, D., & Ferrera, V. P. (2014). Performance monitoring in monkey frontal eye field. *The Journal of Neuroscience*, 34, 1657–1671. http://dx.doi.org/10.1523/JNEUROSCI.3694-13.2014
- Terrace, H. S., & Son, L. K. (2009). Comparative metacognition. Current Opinion in Neurobiology, 19, 67–74. http://dx.doi.org/10.1016/j.conb .2009.06.004
- Tolvanen, V., Jylänki, P., & Vehtari, A. (2014, September). Expectation propagation for nonstationary heteroscedastic Gaussian process regression. Paper presented at the 2014 IEEE International Workshop on Machine Learning for Signal Processing, Reims, France.
- Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., & Vehtari, A. (2013). GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14, 1175–1179.
- Vasconcelos, M. (2008). Transitive inference in non-human animals: An empirical and theoretical analysis. *Behavioural Processes*, 78, 313–334. http://dx.doi.org/10.1016/j.beproc.2008.02.017
- von Fersen, L., Wynne, C. D. L., Delius, J. D., & Staddon, J. E. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 334–341. http://dx.doi.org/ 10.1037/0097-7403.17.3.334
- Weaver, J. C., Steirn, J. N., & Zentall, T. R. (1997). Transitive inference in pigeons: Control for differential value transfer. *Psychonomic Bulletin & Review*, 4, 113–117. http://dx.doi.org/10.3758/BF03210782
- Wright, A. A., & Katz, J. S. (2007). Generalization hypothesis of abstractconcept learning: Learning strategies and related issues in *Macaca mulatta*, *Cebus apella*, and *Columba livia*. *Journal of Comparative Psychology*, 121, 387–397. http://dx.doi.org/10.1037/0735-7036.121.4 .387
- Wynne, C. D. L. (1995). Reinforcement accounts for transitive inference performance. Animal Learning & Behavior, 23, 207–217. http://dx.doi .org/10.3758/BF03199936
- Wynne, C. D. L. (1997). Pigeon transitive inference: Tests of simple accounts of a complex performance. *Behavioural Processes*, 39, 95–112. http://dx.doi.org/10.1016/S0376-6357(96)00048-4
- Zentall, T. R. (2001). The case for a cognitive approach to animal learning and behavior. *Behavioural Processes*, *54*, 65–78. http://dx.doi.org/10 .1016/S0376-6357(01)00150-4

Received October 30, 2016 Revision received January 23, 2017 Accepted January 24, 2017