Positional Inference in Rhesus Macaques

Greg Jensen^{1,2,5}, Vincent P. Ferrera^{1,2,4}, and Herbert S. Terrace^{3,4}

¹Dept. of Neuroscience, Columbia University

²Zuckerman Mind Brain Behavior Institute, Columbia University

³Dept. of Psychology, Columbia University

⁴Dept. of Psychiatry, Columbia University

⁵Corresponding author (email: greg.guichard.jensen@gmail.com)

ABSTRACT

Understanding how organisms make transitive inferences is critical to understanding their general ability to learn serial relationships. In this context, transitive inference (TI) can be understood as a specific heuristic that applies broadly to many different serial learning tasks, which have been the focus of hundreds of studies involving dozens of species. In the present study, monkeys learned the order of 7-item lists of photographic stimuli by trial and error, and were then tested on “derived” lists. These derived lists combined stimuli from multiple training lists in ambiguous ways. We found that subjects displayed strong preferences when presented with novel test pairs. These preferences were helpful when test pairs had an ordering congruent with their ranks during training, but yielded consistently below-chance performance when pairs had an incongruent order relative to training. This behavior can be explained by the joint contributions of transitive inference and another heuristic that we refer to as “positional inference.” Positional inferences play a complementary role to transitive inferences in facilitating choices between novel pairs of stimuli. The theoretical framework that best explains both transitive and positional inferences is a spatial model that represents both the position and uncertainty of each stimulus. A computational implementation of this framework yields accurate predictions about both correct responses and errors for derived lists.

Keywords: derived list, transitive inference, positional inference, serial learning, symbolic distance effect

Serial learning refers to the ability of an organism to learn the underlying order of a set of items and to exploit that knowledge to make inferences about their relative position in that order (Terrace, 2010, 2012). As it is usually defined, such learning arises from experience and feedback, rather than from a comparison of the sensory features of the stimuli. For example, to order a collection of stones from smallest to largest, a subject could compare the size of the stones directly without reasoning about the stones as an ordered set. By contrast, learning the order of the Latin alphabet cannot be explained by mere stimulus comparisons. Nothing about the names of letters or the symbols that represent them suggests that Z has an obvious “lastness.” Its position is thus arbitrary relative to sensory properties, as is its close cousin Zeta being the 6th item of the Greek alphabet. Serial learning is necessary for such arbitrary orderings to be encoded.

Many simple organisms are able to organize their behavior without relying on serial learning at all. For example, single-celled organisms may display complex and adaptive behaviors in response to changing chemical gradients. However, the overwhelming majority of complex organisms that have been tested are able to behave in accord with implied orderings that are not signaled by explicit spatial or temporal cues (Jensen, 2017). Such behavior often requires the ability to make inferences about transitive serial relationships. Since the first non-human demonstration of transitive inference in squirrel monkeys (McGonigle and Chalmers, 1977), evidence of such inferences has been reported in dozens of species.

Although the scope and sophistication of serial learning varies, some form has been reported in every vertebrate species that has been tested (Jensen, 2017). Recently, fish (Hotta et al., 2020) and wasps (Tibbetts et al., 2019) have displayed evidence of being able to not only learn arbitrarily ordered sequences of discriminable colors, but to also exploit that knowledge to make transitive inferences about novel stimulus pairings. For example, wasps trained to favor red over blue, as well as blue over yellow, were given a choice between red and yellow. They were more likely to choose red, apparently demonstrating the transitive rule that $\text{red} > \text{blue} > \text{yellow}$ implies $\text{red} > \text{yellow}$.

Associative Proposals for Serial Learning

Although the accumulated evidence shows that many non-human animals make choices consistent with transitive inference, competing mechanisms have been proposed to explain these choices, inspired by different theoretical frameworks. One such proposal is that model-free stimulus-outcome associations can explain seemingly transitive reasoning. Although the most basic associationist accounts of reinforcement are ruled out for non-trivial cases of serial learning (Lashley, 1951), more complex reinforcement-based accounts can give rise to serial learning under some conditions. (Wynne, 1995, 1997) went so far as to declare the matter settled in favor of reinforcement and association, at least with respect to the most common experimental protocol for transitive inference, and this view still enjoys some support (reviewed by Vasconcelos, 2008).

During the last decade, however, many new experimental designs have pushed associative predictions past their breaking point. Lazareva and Wasserman (2006, 2012) trained pigeons using a “bias reversal” procedure in which the pair DE was trained separately for an extended period following training on adjacent pairs from the list ABCDE. Ordinarily, pigeons trained on adjacent pairs infer that $B > D$. However, as a consequence of the bias reversal procedure, more food was delivered as a consequence of choosing D than of choosing other stimuli (both in trials immediately before testing and overall across the entire learning history). Thus, even sophisticated associative models championed by Wynne predicted low (or even below-chance) accuracy on the critical pair BD, a potential “reversal” of the bias to favor D. Despite this additional training, however, pigeons chose B over D during the test trials. Monkeys show a similar resistance to bias reversal (Jensen et al., 2017). As such, it appears that as a general rule, bias reversal procedures succeed in changing the preference of associative models, but fail to change the preferences of actual organisms, making such associations a poor candidate mechanism for transitive inference. The viability of associative models has been further undermined by experiments in which larger rewards were given after responding to stimuli that are correct less often. Although sensitive to that manipulation, monkeys are resilient to its effects and still perform consistently above chance (Gazes et al., 2012; Jensen et al., 2019).

Although purely associative mechanisms appear to be ruled out by recent results, the manifestation of a transitivity-consistent preference at test is not sufficient to determine the mechanism underlying that preference. To do that we must turn to other features of TI performance. The most widely discussed feature is the “symbolic distance effect” (SDE), a common result in which some novel pairings appear to be easier than others (D’Amato and Colombo, 1990). As the number of intervening steps between list items grows (i.e. the number of steps between symbols increases), performance tends to increase (and reaction times decrease). For example, given an SDE, response accuracy for the pair BE is expected to be greater than that for the pair BD because the stimuli are more widely spaced within the list.

Measuring the SDE provides a more continuous basis for assessing TI than the dichotomous judgments of whether critical pairs are above chance. This is especially true when organisms learn longer lists, since these have multiple critical pairs. While a 5-item list has only one critical pair (BD), a 6-item list has two critical pairs whose symbolic distance is 2 (BD and CE) and a third whose symbolic distance is 3 (BE).

Models of Stimulus Position

Although cognitive models of serial learning have generally emphasized “representations” of list order to account for list learning, the features of such representations vary considerably from one model to the next. At a minimum, such a representation must encode ordinal rank in order to exploit transitivity (Terrace, 1986) and cognitive models are often framed as “ordinal” to remain broadly agnostic and avoid claims based on propositional logic (Terrace, 2012). More recent proposals argue that behavior may be better explained by a representation in which stimuli are represented using a scalar metric, such as position along something like a continuous number line (Jensen et al., 2013). The advantage of this approach is to introduce the possibility that items are represented not only in terms of position but also in terms of their uncertainty within that continuum. If item position is “fuzzy,” then the differing degrees of overlap between items can help to explain the symbolic distance effect, as well as the resilience of serial learning to bias reversal procedures (Jensen et al., 2015).

Evidence consistent with a scalar representation of list items comes from experiments in which subjects learned different lists and were then required to make judgments of the relative position of items drawn from those lists. For example, Merritt and Terrace (2011) trained monkeys using 5-item and 9-item lists and then tested them on pairs in which one item was drawn from each of those lists, rewarding choices to the item that appeared earlier in its original ordering. If subjects were given a choice between items that had the same absolute rank during training (e.g. the third item in a 5-item list vs. the third item in a 9-item list), subjects favored the item from the longer list, suggesting that their judgment depended on an item’s relative position. Monkeys also preferred the 5th-ranked item from a 9-item list over the 4th-ranked item from a 5-item list. This result is consistent with a representation that rescales the distances between items as a function of list length, something that could not be encoded with a strictly ordinal representation.

Merritt and Terrace’s approach can be viewed as an example of a broader class of protocols for probing serial memory called “derived lists.” First used by Ebbinghaus (1913), a derived list consists of stimuli assembled from lists that were trained independently of one another. If each item in a derived list comes from a different training list, items in this new list should have no association with one another. It follows that subjects’ initial choices on a derived list must depend on additional assumptions. A subject’s initial preferences when presented with between-list pairs should therefore provide information about the representations of both training lists that were encoded during training.

For example, Chen et al. (1997) trained monkeys using the simultaneous chain procedure (Terrace, 2005). After learning four 4-item lists, monkeys were presented with derived lists consisting of one item taken from each of the training lists. Monkeys learned these derived lists more rapidly if each item’s new rank matched the rank it held during training. If a monkey had learned four lists, ABCD, EFGH, IJKL, and MNOP, then its learning rate for the derived list ENCL was rapid, since each stimulus was in the same position it had held during training. However, monkeys had more difficulty learning the derived list DKFM, in which each item occupied a different rank than was implied during training. Along the same lines, Gazes et al. (2014) demonstrated a similar effect in a transitive inference paradigm, showing that when the order of a set of stimuli was changed from training to testing, accuracy for “incongruent” between-list pairs was lower than it was to “congruent” pairs.

Kao et al. (2020) systematically extended the derived list approach. Human participants learned the order of five lists, each containing five items. They were then tested using derived lists that varied how much each item deviated from the positions implied during training. Derived lists that preserved the relative ordering of its items yielded higher accuracy than lists whose relative ordering contradicted the training order. Put another way, participant performance could be predicted on the basis of how congruous or incongruous the training order was with the testing order for each derived list. The results led Kao and colleagues to identify two forms of inference. The first is the familiar “transitive” inference, which

allows participants to exploit the transitive property of a strictly ranked list. The second is a “positional” inference, which occurs not on the basis of relative ranks of stimuli, but rather on the overall relationship of one list of items to another.

To understand what we mean by a positional inference, consider a participant that is unfamiliar with any Latin alphabet. This participant is given two lists containing the letters ABCDE and FGHIJ and are tasked with choosing the “earlier letter” in pairs drawn from each list. If the transitive property holds, then a participant will conclude that $B > D$ and $G > I$. However, transitivity does not provide information about how to evaluate the pair GC the first time it is presented. If the participant relies on a positional inference that the two training lists have the same “start” and “end” positions, then *lacking any other information*, they would initially favor G over C, because G appears earlier in one list than C does in the other. By that same assumption, the participant should also favor B over H, since any item in the second list position should win out over any item in the third list position. This initial bias is, of course, easily overcome with training; anyone familiar with the Latin alphabet has learned that C comes before G. Kao and colleagues argue that positional inferences are a heuristic default in ambiguous situations. The transitive inference should also be understood in these terms: it is a heuristic that arises given the assumption that stimuli relate to one another in an ordinal fashion. However, transitivity is not always an appropriate assumption. TI is not recommended when learning to play rock-paper-scissors.

In summary, we use the term “positional inference” to refer to the assumptions that inform whatever preferences subjects display when presented with stimulus pairs that did not belong to the same ranked lists during training. As such, we use “transitive inference” to refer only to preferences that rely on the transitive property of ranked lists. If, for example, a subject was exposed to the relationships $A > B$, $B > C$, $D > E$, and $E > F$, we assert that the assumptions underlying transitive inference provides no logical basis for a judgment about the AF. Subjects who display a preference for A over F (or vice versa) would have to rely on additional information. We refer to the information on which such a preference is based as “positional inference” because it reveals how the subjects aligns the position of list ABC overall, relative to the list DEF.

The present study sought to replicate the findings of Kao and colleagues in rhesus monkeys. Since humans learn ordered lists more rapidly than monkeys (Jensen et al., 2015), monkeys were given a week of training prior to testing. Experiment 1 examined derived lists that preserved the ranks of stimuli, in which pairs were entirely congruent with the order implied by training. Experiment 2 examined the consequences of using derived lists whose orderings conflicted in some cases with positional inference, potentially disrupting performance.

EXPERIMENT 1: DERIVED LISTS WITH MAINTAINED ORDINAL POSITIONS

Although it seems intuitively reasonable that subjects would make both transitive and positional inferences, the latter possibility has not been systematically studied in animals across all possible pairs in a derived list. The purpose of Experiment 1 was to create a derived list that contains some stimulus pairs at test that require transitive (but not positional) inferences, some that require positional (but not transitive) inferences, and some that require both. By testing these simultaneously, we can determine whether subjects rely on positional information when determining how to relate stimuli combined from different training lists relate to one another. We also measured response accuracy for these pairs to compare it to the accuracy of within-training-list pairs that subjects should be able to solve based on transitive inference alone.

Figure 1. Description of the overall experimental procedure in Experiment 1. **(Top)** Depiction of a 7-item list ABCDEFG, consisting of randomly selected stock images. Apart from visual inspection to ensure that stimuli were not confusable, stimuli followed no particular pattern in their appearance. At no stage do subjects see all seven items at once. **(Left)** Representation of a single experimental trial. Stimuli are presented to scale, relative to the tablet computer’s screen. Each trial begins with a blue start stimulus, followed by the presentation of a pair of list stimuli. If the subject touches the image with the lower rank (C, in this case, on the right side of the screen), a reward is delivered. Otherwise, a time-out occurs. Image position on screen is counterbalanced across trials. **(Center)** A depiction of the stimulus pairs used during training and testing. Training consists of just the six adjacent pairs from a 7-item list. Testing consists of all pairs drawn from a 6-item list. **(Right)** Description of training and testing sequence. The first four sessions consist of 100 trials of training each, alternating between two 7-item lists. On the fifth day, 288 trials of training (split across the two lists) are followed by 420 trials of testing (split across the two resulting derived lists).

Methods

Subjects

Subjects were four adult rhesus macaques (*Macaca mulatta*), identified here as N, O, R, and S. All subjects had previous experience with serial learning procedures, including transitive inference. However, TI using derived lists was novel for all subjects prior to Experiment 1.

Subjects were housed individually in a colony room at the Animal Care Facility of the New York State Psychiatric Institute (NYSPI), along with approximately two dozen other macaques. Because correct responses were rewarded with water, subjects were fluid-restricted to 300mL of water per day, or however much they were able to obtain by performing the task, whichever was greater (with additional water supplemented after each session as needed). Typical performance yielded between 200mL and 300mL, whereas perfect performance could yield as much as 500mL. Monkeys were also given a ration of biscuits (provided before the session) and fruit (provided after the session).

The study was carried out in accordance with the guidelines provided by the *Guide for the Care and Use of Laboratory Animals* of the National Institute of Health (NIH). The study was overseen by NYSPI’s Department of Comparative Medicine (DCM) and was approved by the Institutional Animal Care and Use Committees (IACUC) at Columbia University and NYSPI.

Apparatus

Subjects performed the task using a tablet computer mounted in their home cage. The tablet, running Windows 8.1, had a 10.1" HD display (1266×768 resolution) which both presented stimuli and provided a capacitive touch screen interface to record responses. All tasks were programmed in JavaScript and run using the Google Chrome browser, set in kiosk mode.

In order to deliver rewards, the tablet was connected to a solenoid valve by way of an Arduino Nano interface, which opened the valve for fixed intervals. Rewards were delivered via a stainless steel spigot below the tablet. Each valve opening (a "drop") delivered 0.25mL of water. When subjects received multiple drops, the valve opened and closed that many times in rapid succession to ensure that a consistent volume of liquid was being delivered. Unless otherwise noted, this device was identical to that described by Tanner et al. (2017).

This apparatus was mounted in a Lexan frame, which fit securely into the space created by opening the door to the subject's home cage. At the start of each trial, a solid blue square was presented in the center of the screen, in order to focus the subject's attention and to direct their hand toward a consistent center point. Touching it initiated the next trial. All experimental stimuli were 250x250 pixel images, presented to the right and left of the start stimulus.

Procedure

Subjects made responses by touching stimuli that appeared on the tablet screen. Stimuli consisted of photographic images, drawn at random from a large bank of stock photos. A "list" of stimuli consisted of seven such images (denoted as "ABCDEFGH"). The images were assigned a rank order by the experimenter, such that stimuli closer to the beginning of the list (i.e. with a lower rank) were preferable to stimuli later in the list. For example, given the pair AB, stimulus A would be the correct answer, and touching it would yield a three-drop reward, whereas touching stimulus B would be treated as incorrect. However, given the pair BC, stimulus B would be correct. Figure 1 (top) gives an example of such a 7-item list, with ">" denoting each item dominating those later in the list.

Subjects never saw all seven items at once. Instead, on any given trial, they were presented with pairs of images, as depicted in Figure 1 (left). Following the start stimulus, subjects touched one of the two images. If the touched image had the lower rank, a reward was delivered; otherwise, a cue was provided that indicated that the response was incorrect, along with a 2s time-out before the next trial. Subjects had to infer the list order strictly on the basis of this feedback; no other temporal or spatial cues were provided that could be used to infer which response was correct.

Sessions consisted either of "training" or "testing" phases. During training, only the adjacent pairs of a list were presented; during testing, all pairs were presented. In order to keep the presentation of pairs uniform over time, sessions were split up into blocks, during which each pair was presented twice in order to counterbalance their spatial positions. So, for example, in a 7-item list, a block of 12 trials would present each of the 6 training pairs twice each (once with the correct response on the left, and once with it on the right). These 12 trials were randomly permuted in each block. Similarly, during testing, all pairs were presented in a similarly counterbalanced fashion. Figure 1 (center) presents a diagram of the pairs used in training and testing, as well as the "symbolic distance" for each pair (i.e. the difference between the ranks of the two stimuli).

Experiment 1 spanned 21 weeks, with each week consisting of four training days and one testing day. Figure 1 (right) gives a schematic outline of the phases of training and testing presented during each week, which are described below.

The first four days of each week consisted of training sessions lasting 1000 trials, or until the subject ceased responding (subjects completed the training sessions on most days). On Day 1, subjects were

presented with only the adjacent pairs from a 7-item list, in which each item was an arbitrary stock photograph, chosen at random prior to the start of training:

List 1: $A_1B_1C_1D_1E_1F_1G_1$

Here, the subscript “1” indicates that these are the stimuli from the first training list. As noted above, the session was divided into blocks of 12 trials, presenting each of the six pairs twice (to counterbalance for position on screen). The order of these 12 trials within each block was randomized. This ensured that subjects had uniform exposure to all adjacent pairs.

On Day 2, subjects were introduced to a second 7-item list, consisting of a new set of images never before seen by the subject:

List 2: $A_2B_2C_2D_2E_2F_2G_2$

We distinguish this new set of stimuli by giving them the subscript “2” to indicate that they come from the second training list. As before, subjects completed up to 1000 responses, using List 2, divided into randomized 12-trial blocks in order to keep exposure to the adjacent pairs as uniform as possible. Days 1 and 2 are collectively referred to as “first training,” since they represent the subjects’ first exposure to each list.

On Day 3, subjects were again trained on List 1 for a full session of 1000 trials, and on Day 4, subjects were again trained on List 2 for 1000 trials. Days 3 and 4 are collectively referred to as “second training.”

Day 5 was the testing day and consisted of three phases. During the “refresher” phase, subjects were first presented with 144 trials of the adjacent pairs from List 1, then were presented with 144 trials of adjacent pairs from List 2 (all arranged into 12-trial blocks that counterbalanced for stimulus position). These served to remind subjects of both lists. Following this refresher, a new 6-item list was constructed using stimuli from the two training lists:

List 3: $A_1B_2C_1D_2E_1F_2$

Hereafter, we describe lists assembled in this way as “derived lists,” since they consist of familiar stimuli assembled in a new way. Subjects were presented with all 15 pairs (adjacent and non-adjacent) from this new list, arranged into blocks of 30 trials to counterbalance for position and randomize order of presentation. List 3 was trained for a total of 8 such blocks, totaling 240 trials. These 240 trials are collectively referred to as “first test.” Note that because stimulus G was omitted from our derived lists, each of the six stimuli used to construct them had at least some past history of yielding rewards.

Finally, the last phase of the testing day introduced another list derived from the training lists:

List 4: $A_2B_1C_2D_1E_2F_1$

As with List 3, this derived list was trained for 240 trials, organized into randomized and counterbalanced 30-trial blocks. The List 4 phase is referred to as “second test.”

Figure 2 diagrams the different varieties of pair that subjects might encounter during Experiment 1. In addition to being divided into “adjacent pairs” (which are the only variety to appear during training) and “non-adjacent pairs” (which are among those presented at test), stimulus pairs can also be “within-list” or “between-list.” In this study, the training phase consisted exclusively of “within-list adjacent pairs.” The testing phase, contrastingly, excluded pairs of this type entirely, and consisted instead of a mixture

Figure 2. Characterization of the different varieties of stimulus pairs subjects were presented with during training and testing. **(Left)** Any particular pair could be characterized as “within-list” if both stimuli belonged to the same training list, or as “between-list” if both stimuli belonged to separate training lists. Additionally, any pairs could be characterized as “adjacent” if the difference in ranks between the two was equal to 1, or as “non-adjacent” if the difference in ranks was greater than 1. An example of each type of pair is presented. **(Right)** Schematic depiction of the four example pairs. Note that although $C_2 > E_2$ is implied by the transitive property of ranks, transitivity alone is not sufficient to conclude that $C_2 > E_1$. Such a conclusion depends on additional assumptions, and we collectively refer to these assumptions regarding how List 1 and List 2 relate to one another as “positional inference.”

of the other three varieties. For example, the pair B_1C_1 never appeared during testing, since our derived lists are constructed by alternating between training lists. During testing, the only relationship required to express a preference when given within-list non-adjacent pairs is transitive inference, as that pair of stimuli can be connected by applying the rules of ordinal rank to the information obtained during the training phase. All other pairs during testing, however, are between-list pairs, and so require that subjects make additional assumptions about how Lists 1 and 2 relate to one another as a whole. We collectively refer to these additional assumptions as “positional inference.”

To describe performance during training and testing, the probability of a correct response was estimated using logistic regression. During training, only an intercept and a slope (with respect to trial number) were used. During testing, performance was predicted using trials, symbolic distance, the trial-distance interaction, and “mixture” (set to 0 for within-list pairs, and set to 1 for between-list pairs). In equation form, the model at test takes the following form:

$$p(\text{Correct}) = (1 + \exp(-(\beta_0 + \beta_t \cdot t + \beta_D \cdot D + \beta_{tD} \cdot tD + \beta_M \cdot M)))^{-1} \quad (1)$$

Here, β_0 refers to the intercept, β_t to the slope with respect to trial t , β_D to the slope with respect to symbolic distance D , β_{tD} to the interaction of trials and distances, and β_M to the effect of mixture M (dummy coded as described above). These parameters were fit using multi-level models, yielding both population- and subject-level estimates, which was implemented in the Stan programming language (Carpenter et al., 2017). The code for these models is provided in the electronic supplement.

Derived Lists: First Trial Response Accuracy

Figure 3. Estimated population response accuracy on the first trial of each phase in Experiment 1. Between-list pairs (white circles) consisted of one stimulus taken from each of the training lists, whereas within-list pairs (black circles) drew both images from the same training list (even if that pair had not been seen before). Boxes represent the 80% credible interval, and whiskers represent the 95% credible interval.

Results

Figure 3 shows the estimated population probability of a correct response on the first trial of each phase. At the start of first training, subjects were consistently close to chance, whereas at the start of second training, accuracy was generally close to 70%. This performance was stable during the refresher trials on the testing day. At test, a clear distance effect was observed: Distance 1 pairs had the lowest accuracy, and Distance 5 pairs had the highest. Performance for between-list pairs was broadly comparable to that of within-list pairs, consistent with subjects treating rank as interchangeable between lists.

Figure 4 shows the estimated population- and subject-level parameters describing performance on the first trial of the testing phases. Although the population estimates remain relatively uncertain (on account of the small number of subjects), subject-level estimates are much more precise (on account of the large number of trials completed by each subject). In general, performance on trial zero was clearly above chance, as demonstrated by the large intercept, and a robust symbolic distance effect was also seen. A small effect of mixture was also observed in most subjects, suggesting that accuracy was slightly (but not substantially) lower for between-list pairs compared to within-list.

Discussion

Overall, our results replicate the basic finding that Kao and colleagues reported for derived lists in which all stimuli retained their training rank. Although three out of four subjects displayed a slight disadvantage on between-list pairs (i.e. those that combine stimuli from across both training lists), this effect was generally small compared to the symbolic distance effect. Furthermore, since performance on non-adjacent pairs was high regardless of whether test pairs used within-list or between-list pairs, subjects were successfully able to exploit both transitive and positional relationships simultaneously.

Performance Parameters

Figure 4. Logistic regression parameter estimates for the population (black circles) and for each subject (white diamonds), describing performance on the first trial of each test phase of Experiment 1. Boxes represent the 80% credible interval, and whiskers represent the 95% credible interval.

EXPERIMENT 2: DERIVED LISTS WITH CHANGED ORDINAL POSITIONS

The results of Experiment 1 suggest that subjects assumed that the ranks of each stimulus in each training list could be translated directly to comparisons with any other stimuli. If this is true, then changing the order of the stimuli during testing should specifically penalize those pairs whose new ordering differs from that implied by training, even when the stimuli included in that pair come from different training lists. For example, given the stimuli E_2 and F_1 , a subject making a positional inference would assume that $E_2 > F_1$. We will henceforth refer to this kind of scenario as a “congruent” pair. If the testing order rearranged the stimulus positions so that $F_1 > E_2$, subjects should make correct responses less often. This will be referred to as an “incongruent” pair. Furthermore, subjects making a transitive inference about the stimulus positions (e.g. comparing B_2 and F_1) should be especially predisposed to favor B_2 over F_1 because of its larger symbolic distance, and thus be *more* likely to respond incorrectly if the testing list reverses the order of this pair. The more the testing list differs from the training list, the greater the expected disruption.

In this experiment, we use testing list with different stimulus orders than those implied during training. If subjects learn the absolute positions of stimuli during training and use these to infer the ordering of novel pairings at test, then larger changes to the stimulus ordering should result in a correspondingly large disruption at the outset of testing.

Methods

Subjects & Apparatus

Subjects, apparatus, housing, and care were the same as in Experiment 1.

Procedure

The degree that testing orderings differed from those used during training was measured using a metric called “transposition distance” (TD) (Jaro, 1989). TD can be thought of as the number of adjacent

Figure 5. Task structure for derived lists with scrambled positions. **(Left).** Schematic demonstration of how the two training lists are recombined into two testing lists (one being TD=5, and one being TD=11). **(Right).** Listing of all ten test orderings, with an indication of whether each of the fifteen testing pairs retain the order implied by training (white box, check mark) or has its ordering flipped relative to training (gray box, X mark). TD=5 lists have five flipped pairs, and TD=11 lists have eleven flipped pairs.

transpositions of items it would take to get from one ordering to another. For example, going from ABCDE to BACED would require two transpositions: One flipping A&B, and another flipping D&E. Since only adjacent transpositions are allowed, going from ABCDE to CABDE would also require two transpositions: first flipping B&C, then flipping A&C. This metric is useful for comparing rank orderings, since it formalizes the intuition that ABCDE is a ranking that “more closely resembles” EABCD (in which only E changes its rank relative to the other list elements) than it does EDCBA (in which the entire list order is reversed). As in Experiment 1, test lists consisted of six items. However, the ordering of those lists differed from the training order by either TD = 5 (i.e. 5 adjacent transpositions) or TD = 11 (i.e. 11 adjacent transpositions).

The study proceeded on cycles of 5 days per week, doing so for 12 weeks. The first four days of each week were identical to the training days in Experiment 1: Two 7-item lists were trained, on alternating days, for up to 1000 trials per session. Additionally, the refresher on Day 5 was also identical to Experiment 1: 144 additional trials for adjacent pairs from each of the training lists.

Testing with derived lists began on trial 289 of Day 5. The derived lists presented to subjects did not retain the item order used during training. Using the same stimuli as List 3 in Experiment 1, the order of the stimuli is rearranged into one of the ten following orders:

TD=5 Option 1: FABCDE

TD=5 Option 2: CABFED

TD=5 Option 3: AEDBCF

TD=5 Option 4: BCEAFD

TD=5 Option 5: ADCFBE

TD=11 Option 1: FCDEAB

TD=11 Option 2: EFADCB

TD=11 Option 3: DECFBA

TD=11 Option 4: BFEDCA

TD=11 Option 5: EDFBAC

Derived Lists: First Trial Response Accuracy

Figure 6. Estimated population response accuracy on the first trial of each phase in Experiment 2. Between-list pairs (white circles) consisted of one stimulus taken from each of the training lists, whereas within-list pairs (black circles) drew both images from the same training list (even if that pair had not been seen before). Boxes represent the 80% credible interval, and whiskers represent the 95% credible interval.

It was important to have several options for each TD to ensure that the effects we were observing were not idiosyncratic to any particular permutation of the items. This also ensures that, as much as possible, pairs representing every symbolic distance are presented as both “training-consistent” and “training-inconsistent” pairs. Using this selected test ordering, subjects completed 240 trials during which all 15 stimulus pairs were presented in 30-trial blocks, randomized and counterbalanced as in Experiment 1.

Finally, during the last phase of testing, List 4 was presented, with the order of its stimuli changed. If List 3 received a TD=5 stimulus ordering, then List 4 received a TD=11 ordering, and vice versa. This ensured that each testing phase used each TD once but did so in an unpredictable order.

Figure 5 (left) gives a visual description of one way in which the two training lists might have been reassembled into two testing lists. As in Experiment 1, the stimulus G was not used, so every test stimulus had been paired with rewards at some point during training. Of the twelve remaining stimuli, half were used to build a TD=5 list, and the other half were used to build a TD=11 list, the order of each being drawn at random from the options listed above.

Figure 5 (right) displays whether each of the 15 testing pairs has an order that is reversed from the order implied by training. This also provides an alternative way of calculating TD. If one tallies how test pairs are incongruent with the training order, the count of flipped pairs will equal the transposition distance.

Results

Figure 6 shows estimated response accuracies on the first trial of each phase. As in Experiment 1, response accuracy was roughly at chance levels at the start of first training, but had risen to a stable level by the start of second training, which is maintained during the testing day refresher. In the TD=5 derived list,

Figure 7. Logistic regression parameter estimates for the population (black circles) and for each subject (white diamonds), describing performance on the first trial of each test phase of Experiment 2. Boxes represent the 80% credible interval, and whiskers represent the 95% credible interval.

a positive distance effect was obtained, but overall performance was lower than in the TD=0 scenario used in Experiment 1. In the TD=11 lists, responding was initially below chance, and a negative distance effect was observed, with more widely spaced stimuli resulting in lower overall accuracy. This is consistent with the observation that pairs whose symbolic distance was large during training are more likely to be incongruent in the TD=11 testing order. As a direct consequence, a subject relying on positional inferences for its preferences would be expected to display below-chance accuracy when presented with large-symbolic-distance pairs in that condition. Unlike Experiment 1, between-list pairs appeared to slightly outperform within-list pairs.

Figure 7 describes these effects in terms of the population- and subject-level parameters describing performance on trial zero. In the TD=5 scenario, a smaller intercept was observed than in Experiment 1, but it and its corresponding distance effects were nevertheless positive. In the TD=11 scenario, however, both the intercept and the distance effects tended to be negative, with subject-level parameters consistently excluding chance from the 95% credible interval. While mixture effects tended to be a bit above zero (contrary to Experiment 1), they were also relatively small.

Discussion

The high levels of performance in Experiment 1 were consistent with a reliance on positional inferences in the derived lists used in that experiment enhanced performance. By contrast, Experiment 2 used lists in which positional inferences were a liability. The disruption in performance under the TD=11 condition is not surprising because subjects tested in that condition would be incorrect for a majority of the test pairs if they executed the strategy they had used in Experiment 1. However, it is also important that there were only small *differential* effects for between-list pairs, suggesting that subjects treated between-list pairs similarly to within-list pairs, with varying performance determined primarily by rank during training in both cases. In other words, subjects began the test phase in both experiments with positional inference as a default assumption, even though this default assumption resulted in below-chance performance for the TD=11 condition. Our results are consistent with those obtained by (Kao et al., 2020).

THEORETICAL ANALYSIS: INFERRING REPRESENTATIONS FROM PERFORMANCE

Although regression models like those provided above can give a good description of the patterns in observed data, and can even allow the relative effect size of various contributing associations to be measured, they should always be understood to be nothing more than descriptive. McElreath (2020) describes such regressions as “geocentric” as an allusion to Ptolemy’s model for calculating the positions of planets in the night sky, which effectively relied on a primitive form of Fourier analysis and as such could give an arbitrarily precise correct prediction, provided the analyst was willing to introduce enough epicyclic parameters. As such, even regression models that make correct predictions with high accuracy usually rely on mathematical machinery that does not resemble the underlying mechanism in any of its particulars.

Although the behavioral data alone is not sufficient to reconstruct the particular computations that subjects were making, building a model that directly implements our theory should, if the theory has any merit, provide a means of seeing whether performance during training is able to make specific predictions about performance during testing. That is, whereas the regressions described above describe response accuracy during testing, implementing a generative model should allow test performance to be predicted on the basis of the training data alone.

Methods

Although the betasort algorithm (Jensen et al., 2015) is computationally efficient at simulating behavior, it is computationally costly for the purposes of statistical inference. Even though generating a random response (given some set of parameters) is fast, computing the likelihood associated with that response involves computing an integral that has no closed-form solution and no off-the-shelf optimized approximation (Jensen et al., 2015, Equation 9). Although hill-climbing can be used to obtain reasonable parameter estimates for single subjects, estimation of multi-level model parameters using Monte Carlo methods is prohibitive.

As a computationally efficient alternative, betasort performance can be approximated as a series of overlapping Gaussian distributions (Jensen et al., 2019). The position of the stimulus A can be described by two parameters, a location μ_A and a scale σ_A , such that the choice of stimulus A occurs when it receives a larger random draw from that Gaussian distribution than stimulus B. Since the difference between random draws is itself Gaussian, the probability of choosing A vs. B can be determined using the value at zero of the integral of the standard Gaussian distribution, denoted by $\Phi(x)$:

$$p(\text{Choose } A | A, B) = \Phi\left(\frac{\mu_B - \mu_A}{\sqrt{\sigma_A^2 + \sigma_B^2}}\right) = \int_0^\infty \mathcal{N}\left(x | \mu = \mu_B - \mu_A, \sigma = \sqrt{\sigma_A^2 + \sigma_B^2}\right) dx \quad (2)$$

Since $\Phi(x)$ is a well-understood and well-optimized function, it can reliably be deployed to convert parameters μ_i and σ_i for all stimuli i into likelihoods. That allows the posterior distributions of those parameters to be inferred by Monte Carlo methods.

Figure 8 depicts this logic visually. Each probability of a correct response is uncertain, denoted as the “estimated response proportion” (left column), represented by beta distributions. The density of those distributions at each point is presumed to result from the $\Phi(x)$ function, given a standard score that depends on two μ terms and two σ terms, denoted as the “estimated difference distributions” (center column). Those μ and σ terms, in turn, collectively belong to an “estimated representation” of the stimulus positions along a scalar continuum. Using a numerical procedure, these parameters can be fit

Figure 8. Schematic depiction of how response accuracies are used to infer a subject's representation of position. These representations are described in terms of μ and σ terms (with each stimulus being assigned one of each). These parameters are estimated simultaneously based on all relevant pairs (which, during training, consists of the 6 adjacent pairs). Since each response accuracy can be described in terms of the overlap between position distributions, the accuracies can be used to estimate those positions numerically, using an MCMC sampler.

simultaneously, allowing us to infer the contents of the representation from the observed error patterns for the various stimulus pairs (see Jensen et al., 2019, for another implementation of this method).

In addition to using μ_i and σ_i to describe the representation of each stimulus position, the betasort model uses a “random response” parameter, which represents the odds that a response made by a subject will be completely arbitrary. We will denote this as r . Such random responses may occur if, for example, the subject makes some proportion of its responses while distracted, or alternatively may represent an exploratory strategy. Consequently, the full equation for a correct response to A given AB would be:

$$p(\text{Choose } A | A, B) = \frac{r}{2} + (1 - r) \Phi \left(\frac{\mu_B - \mu_A}{\sqrt{\sigma_A^2 + \sigma_B^2}} \right) \quad (3)$$

This is effectively a mixture distribution, where r governs the mixture of trials that depend on the representation with those that do not. Note that this model is strictly descriptive: it describes response accuracy at a given point during learning, but is agnostic about the mechanism by which that learning arose.

In order to derive a complete representation of performance from the training data, 15 parameters are required: one μ_i and one σ_i for each of the seven stimuli, as well as an additional parameter r governing impulsive responses. Any single trial, in turn, depends on five of those parameters. The overlap of parameters across trials provides constraint to help avoid overfitting. Finally, since μ_i and σ_i have no natural units, the additional requirements that $\mu(\mu_A, \mu_B, \mu_C, \mu_D, \mu_E, \mu_F, \mu_G) = 0$ and $\sigma(\mu_A, \mu_B, \mu_C, \mu_D, \mu_E, \mu_F, \mu_G) = 1$ are imposed, in order to keep values centered and prevent their values from becoming arbitrarily large or small.

Figure 9. Inferred representation of relative stimulus positions following training in both experiments. (Left). Estimates of relative stimulus positions after training for each subject. White violin plots represent uncertainty about the mean, whereas gray violin plots take the variance of the estimates into account and this represent the total overall uncertainty about the stimulus position. (Right). Estimated probability of a “guess” response for each subject. Boxes represent the 80% credible interval, and whiskers represent the 95% credible interval.

Experiments 1 and 2 made use of identical training procedures, and performance during second training and during the refresher on Day 5 were representative of stable performance for all subjects. Consequently, if we assume that performance was stable after the first session, then for 33 pairs of training lists, each subject provided two counts of successes and attempts for each adjacent pair. For example, during the first week of Experiment 1, Subject 1 was presented with the pair AB 195 times during second training and 48 times during the refresher session, and correctly chose A 173 times and 42 times, respectively.

The 15 parameters needed to describe each subject’s representation of a given list were fit using a multi-level model implemented in the Stan programming language (Carpenter et al., 2017), whose code is provided in the electronic supplement. Parameters for these models depended only on the second training and refresher data; neither data from first training or from testing contributed to the estimates. Additionally, in order to provide empirical data against which to compare model predictions, response accuracy for each of the 15 testing pairs was fit separately using only the testing data, in order to give a pairwise estimate of trial zero response accuracy at test. By comparing how well the model predicts testing performance (based only on the training data), as compared to the empirical estimate of observed testing performance, we can assess whether Equation 3 gives a good characterization of what animals learned during training.

Results

Figure 9 (left) shows the inferred stimulus positions for each of the four subjects. Here, the uncertainty for the mean μ_i is shown as a white violin plot, whereas the full positional uncertainty for each item (i.e. taking σ_i into account) is shown as a gray violin plot (Hintze and Nelson, 1998). Note that the model itself was given no information about the ordering of the stimuli. Thus the consistent uniformity with which the stimuli were ordered reflects subjects’ performance. However, item F tended to have somewhat higher uncertainty than its position would otherwise imply (particularly in subjects N and R). Consequently, the

pair EF tended to have the lowest expected accuracy of any pair.

Figure 9 (right) shows the random choice parameter r estimated for each subject. This suggests that subjects N, R, and S behaved as though they were guessing around 14% of the time during training (albeit with a lot of uncertainty in the estimate), whereas subject O behaved as though he was guessing around 8% of the time.

The models described by the parameters in Figure 9 represent estimates of the monkeys' knowledge of ordinal position. They can be used to make precise predictions about how likely subjects are to make correct responses on their first choice during testing with derived lists. Comparing these predictions with the testing data provides a further validation of the model. To obtain an empirical estimate of this response probability directly from the observed data, a logistic regression was performed for each of the testing pairs in each of the list orders. With one TD=0 list, five TD=5 lists, and five TD=11 lists, there were 165 separate regressions for each subject, yielding a predicted accuracy for each pair. Because these regressions were often based on a small number of trials, these estimates of observed response accuracy were often highly uncertain.

In order to describe the pattern of which pairs were likely to be above or below chance at testing, it is important to distinguish between the symbolic distance in the new testing order, and the symbolic distance implied by the old training order. For example, consider the testing list order FABCDE and its first adjacent pair, FA. Under the new order, FA has a symbolic distance of 1. However, subjects presented with the pair FA on the first trial of testing would be expected to overwhelmingly favor A over F, not only because $A > F$ was implied by their training, but also because that same training had implied that the symbolic distance between A and F would be 5, which (based on the pattern observed in Figure 2) should be an especially "easy" discrimination. We will use the term "symbolic distance discrepancy" to refer to the gap between stimuli implied by training, using negative values for testing pairs that are incongruent with the training order. Accordingly, pair FA has a symbolic distance discrepancy of -5. Its magnitude is 5 because training implied that A was five ranks before F, and the value is negative because F is now the correct choice, rather than A, making FA an incongruous between-list pair. By contrast, the pair BC in this list order has a value of 1 for both its symbolic distance and its symbolic distance discrepancy. The two stimuli are still adjacent, and the order implied by training remains valid during testing, making BC a congruous between-list pair.

Figure 10 (left) depicts the predicted response accuracy for 165 stimulus pairs (pooled across all 11 testing orders), sorted by their symbolic distance discrepancy. These predictions were based on our Bayesian model of stimulus positions, whose probabilities are calculated using Equation 3. If stimuli flipped their order in the transition from training to testing (i.e. if their discrepancy values were negative), the model predicts that responding should be below chance. The larger the discrepancy, the further from chance behavior should be: Discrepancies of 5 should display the highest accuracies, while discrepancies of -5 should display the lowest accuracies. In addition to the accuracies for individual pairs, an aggregate uncertainty of the mean for all pairs with a particular discrepancy is presented as a horizontal black line. These aggregate means were weighted as a function of the precision of the estimate (i.e. the inverse of the estimate's variance) for each pair, so that highly uncertain pairs contributed less to the mean than very certain pairs.

Figure 10 (right) depicts the estimated response accuracies for each of these pairs, based on 165 separate logistic regressions. Although these estimates are individually uncertain, the overall pattern is clear: Low accuracies for negative discrepancies, high accuracies for positive discrepancies, and values approaching their respective asymptotes as the discrepancies get further from zero. As with the model estimates, aggregate weighted means were obtained for all pairs with a particular discrepancy, plotted as horizontal black lines.

Figure 10. Response accuracy for all pairs on the first trial of Experiment 2 as a function of symbolic distance discrepancy. Points represent each of 165 pairs (15 per lists across 11 lists), and whiskers represent the 95% credible interval for each pair. Additionally, black horizontal lines represent weighted means of the accuracy of pairs with that symbolic distance discrepancy, and gray shaded areas represent the 95% credible interval for those means. (Left). Response accuracy predicted by the Bayesian model of stimulus position depicted in Figure 8. (Right). Response accuracy observed in subjects, estimated by performing a separate logistic regression on each of the pairs.

The weighted means of predicted vs. observed response accuracy that are depicted in Figure 10 are plotted against one another in Figure 11. From this, it is clear that although there is a good overall concordance between predicted and observed accuracies (Spearman's ρ of at least 0.945, based on bootstrapped estimates), the relationship is not precisely linear. Instead, the predicted accuracies tend to bunch toward the extremes relative to the observed accuracies, resulting in a sigmoid deviation from the identity line.

Figure 11. Association between predicted vs. observed performance, based on the ten weighted means of pairwise response accuracy depicted in Figure 10. Whiskers represent the 95% credible interval for the weighted means. Bootstrapped estimates of Spearman's ρ confirms that these correspond to a high level of concordance between observation and prediction (no mean ρ_s less than 0.945, and no credible interval reaching lower than 0.867).

Discussion

Although this additional Bayesian modeling exercise may appear redundant with the regression model depicted in the earlier results sections (particularly Figures 3 and 6), Figures 10 and 11 contributes fundamentally different information. Logistic regression (in this study, and in general) is best understood as a descriptive tool, even when analysts use it to make predictions. Even a model founded on spurious correlations can make good predictions if a phenomenon's underlying causal structure remains unchanged. What our additional analysis aimed to do was to make a strong set of assumptions ("monkeys behave as if they represent stimulus position on a continuum with Gaussian noise") and fit a model founded on that assumption using only the training data. That model was then used to predict, with considerable accuracy, performance when subjects were presented with derived lists whose item order differed from that implied by training. Not only was this model effective at describing distance effects under test conditions, but it reliably described the degree to which the combination of transitive and positional inferences would yield below-chance responding at test (Figure 11).

The "symbolic distance discrepancy" approach is reminiscent of (and partly inspired by) the list-linking work initially spearheaded by Treichler and Van Tilburg (1996). In perhaps the most ambitious design undertaken by the group, Treichler et al. (2007) first trained monkeys on 5-item lists. They then presented subjects with 15-item derived lists, constructed by stringing each 5-item list together in sequence (so, all the items in one list, followed by all the items in the second, and ending with all the items in the third). For example, a monkey might learn ABCDE, FGHIJ, and LKMNO as three separate lists and (in the study's control condition) transition directly to being tested on pairs from the list ABCDEFGHIJKLMNO. Since the control condition gave no indication that $B > F$, the natural intuition arising from a positional inference would be that F should be preferred (since it was the first item in FGHIJ). This is precisely the pattern that Treichler and colleagues report, with pairs that have a negative symbolic distance discrepancy (termed a "serial gap distance" in that study) displaying below-accuracy responding.

Our results expand upon this by demonstrating that the discrepancy between order implied during training and order required during test holds for a more general reordering of stimuli, and not just for concatenated lists. It also shows that this pattern of both successful inference and reliably unsuccessful error at the start of testing is consistent with a scalar model of item position.

GENERAL DISCUSSION

The present study made use of the “derived list” approach to serial learning in order to shed light on a fundamental question: What assumptions do subjects make when choosing among stimuli that may be related serially? Transitive inference depends on the assumption that stimuli have at least an ordinal relationship, and that choices consistent with that assumption are well documented in a wide variety of species. Between-list pairs, however, are strictly ambiguous if subjects rely exclusively on using the transitive property to assemble ordinal rankings. A subject’s preference when choosing one of a between-list pair of stimuli reveals an assumption, often a heuristic or some other kind of default strategy. In Experiment 1, we harnessed the assumptions underlying positional inferences to obtain above-chance performance for between-list pairs by presenting subjects with lists that strictly preserved their rankings across lists (Figure 3). Subjects displayed both high levels of accuracy and symbolic distance effects, consistent with using both the transitive property of ordinal ranks for within-list pairs and the assumption that ranks were interchangeable across lists for between-list pairs (Figure 4). In Experiment 2, we presented subjects with derived lists that were reordered, such that there was a mix of congruent and incongruent pairs. The patterns of preference displayed by subjects, both above *and* below chance (Figures 6 and 10), were consistent with subjects relying on positional inferences at the start of testing, even after many sessions in which testing reordered the stimuli. In other words, subjects persisted in making assumptions that impaired performance, even in the TD=11 condition when they would have been better off guessing. Performance overall was described fairly well by a model in which subjects represented stimulus position as overlapping distributions along a scalar continuum (Figures 9, 10, and 11).

These results build directly from past work using derived lists and list reordering (Chen et al., 1997; Gazes et al., 2014; Kao et al., 2020), as well as related work in which familiar lists are linked to form longer lists, thereby resulting in longer derived lists with many new pairings of stimuli (e.g. Treichler and Van Tilburg, 1996). List-linking experiments achieve their linkage by training a “linking pair.” For example, lists ABCDE and FGHIJ might be linked by repeated presentations of the pair EF. However, many such experiments include control conditions without linking pair training, and these control conditions effectively produce as many incongruent pairs (e.g. FA) as they do congruent ones (e.g. BC). Incongruent pairs in control conditions generally produce below-chance performance (Treichler and Raghanti, 2010; Gazes et al., 2014; Wei et al., 2014). Additionally, when Treichler et al. (2007) pushed the list-linking procedure to the point of trying to link three 15-item lists into a single 45-item list, their results were somewhat mixed, suggesting that many of the errors on this enormous list were the result of subjects defaulting back to positional inferences, as their errors were consistent with the assumption that stimulus ranks were interchangeable between lists. Though less studied than transitive inferences, this body of work suggests that positional inferences may reflect another set of heuristics favored by organisms in serial learning contexts.

From Phenomenon to Model

Transitivity is a property of rank-ordered sets that allows subjects to infer that if $A > B$ and $B > C$, then $A > C$. Behavioral performance consistent with an ability to make transitive inferences has been demonstrated in monkeys and a host of other species (Jensen, 2017). Most experimental efforts to study transitive inference have sought existence proofs. Although successful choices of critical test pairs imply an ability to order list items, as well as some ability to compare items based on their relative position within the list, there are many ways such a preference could arise. Transitivity only requires a strict ordinal ranking, and only operates within a set. For example, if every member of one school’s track team ran a race and their rank was recorded based on their finishing order, merely knowing the rank-ordering of another set of athletes at a different school would give you no insight into which school’s team would win if all of

the athletes ran in a single race. To make comparisons across different sets requires additional knowledge. In the example of a footrace, knowing an athlete's speed (a scalar measurement with a well-defined measurement scale and zero point) is more informative than their rank (an ordinal comparison to some other athlete).

An ordinal ranking may arise without a full-fledged representation. Efforts to explain serial learning in associative terms (von Fersen et al., 1991; Wynne, 1995; Reid, 2009) have yielded associative models that display choices consistent with “transitive inference” given very specific experimental designs and training histories. However, when experimental designs are changed, employing different amounts of training (Lazareva and Wasserman, 2006, 2012; Jensen et al., 2017) or size of rewards (Gazes et al., 2012; Jensen et al., 2019), associative models usually fail to make correct predictions about subjects' behavior.

Instead, subjects appear to employ a representational framework that is sufficiently flexible to encode the ordering of items under a wide variety of conditions. This feature of serial learning is presumably necessary for it to have adaptive value because the natural world does not present animals with stimuli in carefully balanced experimental designs. Our view is that although it is useful to refer to “transitive inference” and “positional inference” as forms of preference that emerge as a function of feedback, neither should be presumed to necessarily be an explicit cognitive process in its own right. Rather, both transitive and positional inference are phenomena that emerge from an underlying representation. We argue here and elsewhere (Jensen et al. 2015) that subjects represent stimulus position with an error term along a scalar continuum. Although this model was inspired by transitive-inference and list-linking experiments, it is also proper to say that such a model predicts both transitive and positional inference. Furthermore, this model also predicts the degree of preference for both within-pairs and between-list pairs as a function of symbolic distance.

Although studies of TI provide ample evidence that subjects behave as if they represent items along a scalar continuum, this continuum is a theoretical construct that we infer from behavior. The feedback that guides this behavior, a series of correct/incorrect cues in response to choices, are themselves uninformative and relative. As such, subjects who see two items that have never been paired must rely on some additional assumptions to make an educated guess. Although subjects have the option of responding randomly, it's clear that they don't (Figures 2, 5, 9). If the items are connected by a series of past linkages, then a transitive inference may be a safer bet than responding randomly. If two stimuli do not appear even indirectly related, a further assumption is required. We use the term “positional inference” to refer to choices that rely on an item's overall position within a list rather than its relationship to specific items. If a stimulus would “win most matchups” among the stimuli to which it can be connected via transitive inference, a positional inference further assumes that it will probably “win most matchups” in general. Kao et al. (2020) frame this as a judgment of absolute position, in contrast to the judgments of relative position that arise from the assumption of transitivity.

Just as positional inference may be a useful default that subjects can learn to subvert, however, transitivity is also not the only available default assumption. Gross and Greene (2007) presented participants with two simple conditions. One condition required participants to make a transitivity judgment, learning that $A > B > C$. The second condition required a transverse judgment, namely that $A > B > C > A$, similar to a game of Rock-Paper-Scissors in which each stimulus dominates one case and is dominated by another in a closed loop. When participants were then introduced to novel stimuli and trained on the pairs XY and YZ, whether they defaulted to a transitive or a transverse interpretation of test pair XZ depended on which condition they had originally been exposed to. Participants who learned that $A > C$ defaulted to the transitive assumption that $X > Z$, whereas participants who learned that $C > A$ defaulted to the transverse assumption that $Z > X$. Non-human animals are also able to learn to rely on transverse patterning (Couvillon and Bitterman, 1996; Alvarado and Bachevalier, 2005), although this appears to

require considerably more training than typical transitive inference paradigms. This more flexible capacity to understand complex webs of relationship between stimuli that do not necessarily resolve to a strict ordering is sometimes referred to as “configural” learning (see Jacobs, 2006, for review). While the present findings do not undermine claims of transverse inference or other configural learning phenomena, our data nevertheless suggest that transitive and positional inferences may be default heuristics that subjects employ under conditions of uncertainty. Truly transverse relationships are rare in nature. Stimuli can often be ranked or ordered under a wide variety of circumstances, whether those relationships map directly to sensory intensity or spatial position. Serial rankings can also reflect more abstract relationships such as hierarchies of social dominance or foods with different caloric densities.

In conclusion, the present evidence is not merely consistent with the claim that animals rely on both transitive and positional inferences to make decisions, but that both the successes and failures of subjects in our experiments map onto a single framework. Our model represents each stimulus with an uncertain scalar position along an abstract continuum, and this efficiently explains all patterns of preferences described in this study.

ACKNOWLEDGMENTS

We thank Yelda Alkan, David Freshwater, Aliza Gross, Katherine Liu, and Grant Spencer for assistance with data collection.

Funding

This work was supported by US National Institute of Mental Health, grant numbers NIH-MH081153 & NIH-MH111703 awarded to Vincent Ferrera and Herbert Terrace.

AUTHOR CONTRIBUTIONS

GJ, VPF, and HST conceived the experiments. GJ and VPF acquired data. GJ wrote the task software and performed analyses. GJ, VPF, and HST wrote the paper.

REFERENCES

- Alvarado, N. C. and Bachevalier, J. (2005). Comparison of the effects of damage to the perirhinal and parahippocampal cortex on transverse patterning and location memory in rhesus macaques. *Journal of Neuroscience*, 25:1599–1609.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76:1–32.
- Chen, S., Swartz, K. B., and Terrace, H. S. (1997). Knowledge of the ordinal position of list items in rhesus monkeys. *Psychological Science*, 8:80–86.
- Couvillon, P. A. and Bitterman, M. E. (1996). Transverse patterning in pigeons. *Animal Learning & Behavior*, 24:410–422.
- D’Amato, M. R. and Colombo, M. (1990). The symbolic distance effect in monkeys (cebus apella). *Animal Learning & Behavior*, 18:133–140.
- Ebbinghaus, H. (1885—1913). *Memory: A Contribution to Experimental Psychology*. Columbia University Teacher’s College Bureau of Publications.

- Gazes, R. P., Chee, N. W., and 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.
- Gazes, R. P., Lazareva, O. F., Bergene, C. N., and Hampton, R. R. (2014). Effects of spatial training on transitive inference performance in humans and rhesus monkeys. *Journal of Experimental Psychology: Animal Learning and Cognition*, 40:477–489.
- Gross, W. L. and Greene, A. J. (2007). Analogical inference: The role of awareness in abstract learning. *Memory*, 15:838–844.
- Hintze, J. L. and Nelson, R. D. (1998). Violin plots: A box plot-density trace synergism. *American Statistician*, 52:181–184.
- Hotta, T., Ueno, K., Hataji, Y., Kuroshima, H., Fujita, K., and Kohda, M. (2020). Transitive inference in cleaner wrasses (*Labroides dimidiatus*). *PLOS ONE*, 15:e0237817.
- Jacobs, L. F. (2006). From movement to transitivity: The role of hippocampal parallel maps in configural learning. *PLOS ONE*, 17:99–110.
- Jaro, M. A. (1989). Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida. *Journal of the American Statistical Association*, 84:414–420.
- Jensen, G. (2017). Serial learning. In Call, J., Burghardt, G. M., Pepperberg, I. M., Snowdon, C. T., and Zentall, T., editors, *APA Handbook of Comparative Psychology: Perception, Learning, and Cognition*, pages 385–409. American Psychological Association.
- Jensen, G., Alkan, Y., Ferrera, V. P., and Terrace, H. S. (2019). Reward associations do not explain transitive inference performance in monkeys. *Science Advances*, 5:eaaw2089.
- Jensen, G., Alkan, Y., Muñoz, F., Ferrera, V. P., and Terrace, H. S. (2017). Transitive inference in humans (*Homo sapiens*) and rhesus macaques (*Macaca mulatta*) after massed training of the last two list items. *Journal of Comparative Psychology*, 131:231–245.
- Jensen, G., Altschul, D., Danly, E., and Terrace, H. S. (2013). Transfer of a serial representation between two distinct tasks by rhesus macaques. *PLOS ONE*, 8:e70285.
- Jensen, G., Muñoz, F., Alkan, Y., Ferrera, V. P., and Terrace, H. S. (2015). Implicit value updating explains transitive inference performance: The betasort model. *PLOS Computational Biology*, 11:e1004523.
- Kao, T., Jensen, G., Michaelcheck, C., Ferrera, V. P., and Terrace, H. S. (2020). Absolute and relative knowledge of ordinal position. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.
- Lashley, K. S. (1951). The problem of serial order in behavior. In Jeffress, L. A., editor, *Cerebral Mechanisms in Behavior*, pages 112–131. Wiley.
- Lazareva, O. F. and Wasserman, E. A. (2006). Effect of stimulus orderability and reinforcement history on transitive responding in pigeons. *Behavioural Processes*, 72:161–172.
- Lazareva, O. F. and Wasserman, E. A. (2012). Transitive inference in pigeons: Measuring the associative value of stimuli b and d. *Behavioural Processes*, 89:244–255.
- McElreath, R. (2020). *Statistical Rethinking, Second Edition*. CRC Press, Boca Raton, FL, USA.
- McGonigle, B. O. and Chalmers, M. (1977). Are monkeys logical? *Nature*, 267:694–696.
- Merritt, D. J. and 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.
- Reid, A. K. (2009). Resistance to change within heterogeneous response sequences. *Journal of Experimental Psychology: Animal Behavior Processes*, 35:293–311.
- Tanner, N., Jensen, G., Ferrera, V. P., and Terrace, H. S. (2017). Inferential learning of serial order of perceptual categories by rhesus monkeys (*Macaca mulatta*). *Journal of Neuroscience*, 37:6268–6276.
- Terrace, H. S. (1986). A nonverbal organism's knowledge of ordinal position in a serial learning task. *Journal of Experimental Psychology: Animal Behavior Processes*, 12:203–214.

- Terrace, H. S. (2005). The simultaneous chain: A new approach to serial learning. *TRENDS in Cognitive Sciences*, 9:202–210.
- Terrace, H. S. (2010). The comparative psychology of serially organized behavior. *Comparative Cognition & Behavior Reviews*, 4:39–73.
- Terrace, H. S. (2012). The comparative psychology of ordinal behavior. In Zentall, T. R. and Wasserman, E. A., editors, *Oxford Handbook of Comparative Cognition*, pages 615–651. Oxford University Press, Oxford, UK.
- Tibbetts, E. A., Agudelo, J., Pandit, S., and Roijjas, J. (2019). Transitive inference in polistes paper wasps. *Biology Letters*, 15:20190015.
- Treichler, F. R. and Raghanti, M. A. (2010). Serial list combination by monkeys (*Macaca mulatta*): Test cues and linking. *Animal Cognition*, 13:121–131.
- Treichler, F. R., Raghanti, M. A., and Van Tilburg, D. N. (2007). Serial list linking by macaque monkeys (*Macaca mulatta*): List property limitations. *Journal of Comparative Psychology*, 121:250–259.
- Treichler, F. R. and Van Tilburg, D. N. (1996). Concurrent conditional discrimination tests of transitive inference by macaque monkeys: List linking. *Journal of Experimental Psychology: Animal Behavior Processes*, 22:105–117.
- Vasconcelos, M. (2008). Transitive inference in non-human animals: An empirical and theoretical analysis. *Behavioural Processes*, 78:313–334.
- von Fersen, L., Wynne, C. D. L., Delius, J. D., and Staddon, J. E. R. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17:334–341.
- Wei, C. A., Kamil, A. C., and Bond, A. B. (2014). Direct and relational representation during transitive list linking in pinyon jays (*Gymnorhinus cyanocephalus*). *Journal of Comparative Psychology*, 121:250–259.
- Wynne, C. D. L. (1995). Reinforcement accounts for transitive inference performance. *Animal Learning & Behavior*, 23:207–217.
- Wynne, C. D. L. (1997). Pigeon transitive inference: Tests of simple accounts of a complex performance. *Behavioural Processes*, 39:95–112.