Learning decouples accuracy and reaction time for rapid decisions in a transitive inference task

Abbreviated Title: Modeling rapid decisions in transitive inference

Fabian Munoz[1,2], Greg Jensen[3], Maxwell Shinn[4], Yelda Alkan[1,2], John D. Murray[5], Herbert S. Terrace[6], Vincent P. Ferrera[*,1,2,7,8]

1. Department of Neuroscience, Columbia University Medical Center, New York, NY, 10032,
2. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, 10027
3. Reed College, Portland, OR, 97202
4. UCL Queen Square Institute of Neurology, University College London, London WC1E 6BT, UK
5. Department of Psychology, Dartmouth University, Hanover, NH 03755
6. Department of Psychology, Columbia University, New York, NY 10027
7. Department of Psychiatry, Columbia University Medical Center, New York, NY, 10032
8. Kavli Institute for Brain Sciences, Columbia University, New York, NY, 10027

*Corresponding author email address: vpf3@cumc.columbia.edu

Keywords: sequential sampling, drift-diffusion model, drift rate, collapsing boundary, leaky integration

Abstract

Transitive inference (TI) is a cognitive process in which decisions are guided by internal representations of abstract relationships. While the mechanisms underlying transitive learning have been well studied, the dynamics of the decision-making process during learning and inference remain less clearly understood. In this study, we investigated whether a modeling framework traditionally applied to perceptual decision-making—the drift diffusion model (DDM)—can account for performance in a TI transfer task involving rapid decisions that deviate from standard accuracy and response time (RT) patterns. We trained three macaque monkeys on a TI transfer task, in which they learned the implied order of a novel list of seven images in each behavioral session. Monkeys indicated their decisions with saccadic eye movements. Consistent learning of the list structure was achieved within 200–300 trials per session, with asymptotic accuracies reaching approximately 80–90%. Behavioral performance exhibited a symbolic distance effect, with accuracy increasing as the ordinal distance between items grew. Notably, RTs remained relatively stable across learning, despite improvements in accuracy. We applied a generalized DDM implementation (PyDDM; Shinn et al., 2020) to jointly fit accuracy and RT data. Model fits were achieved by incorporating both an increasing evidence accumulation rate and a collapsing decision bound, successfully capturing the RT

distribution shapes observed during learning. These findings suggest that decision-making during serial learning and transfer in a TI task can be characterized by a "variable collapsing bound" DDM. Our results highlight a distinct dynamical regime of the DDM framework, extending its applicability to cognitive domains involving symbolic reasoning and serial relational learning.

Introduction

Transitive inference (TI) assesses the ability to learn and generalize relationships that obey the principle of transitivity—for example, if A > B and B > C, it follows that A > C. This cognitive ability is typically tested by training subjects on novel lists comprising 5–9 easily discriminable items, such as images, and examining their performance on untrained item pairs. Extensive empirical evidence suggests that TI performance is supported by the formation of an internal representation of the list's ordinal structure (Prado et al., 2012; Bonnefond et al., 2014; Jensen et al., 2015). However, the mechanisms by which such internal representations are translated into behavioral responses remain poorly understood. In particular, it is unclear whether the decision-making processes underlying TI share properties with those involved in perceptual decision-making.

In the present study, we examined the decision dynamics of nonhuman primates performing a TI task. In each session, monkeys were presented with novel sets of images in paired comparisons and were required to infer their implicit order. While performance accuracy increased with learning, response times (RTs) remained fast and relatively constant, despite the

absence of strong time constraints. These findings suggest a dissociation between accuracy improvements and temporal aspects of decision-making, raising questions about how internal representations are accessed and acted upon during learning.

TI differs fundamentally from traditional perceptual decision-making tasks in that it relies on internally constructed rather than externally available evidence. Consequently, variations in performance are more closely tied to abstract task variables than to sensory stimulus properties. One such variable is symbolic distance (SD)—defined as the difference in rank between two items in the inferred list. Greater symbolic distance is strongly associated with higher accuracy, a hallmark of TI performance. This reliance on internally guided inference suggests that the dynamics of TI decisions may diverge from the patterns observed in perceptual tasks, where decision times typically vary systematically with stimulus strength.

In many two-alternative choice paradigms involving nonhuman primates, RTs often exceed one second and are inversely correlated with stimulus strength for both dynamic (Roitman & Shadlen, 2002; Churchland et al., 2008; Ding & Gold, 2010) and static stimuli (Ratcliff et al., 2003; Middlebrooks & Schall, 2014). These behavioral metrics are closely linked to the temporal evolution of neural activity across cortical and subcortical regions, including the lateral intraparietal area (Roitman & Shadlen, 2002), prefrontal cortex (Kim & Shadlen, 1999; Kiani et al., 2014), caudate nucleus (Ding & Gold, 2010), and superior colliculus (Jun et al., 2021). These observations have informed the development of computational models, particularly the drift-diffusion model (DDM), which provide principled explanations for the coupling between accuracy and RT in perceptual decisions.

Recent work highlights the existence of decision processes across multiple timescales (Murray et al., 2014; Soltani et al., 2021), including rapid decisions in naturalistic contexts, where speed is critical for survival. Indeed, studies of oculomotor behavior reveal that primates can make highly accurate decisions within 200–300 milliseconds, with frontal eye field (FEF) activity discriminating targets within this timeframe (Cohen et al., 2010; Stanford et al., 2010; Shankar et al., 2011).

The Drift Diffusion Model (DDM)—a widely used framework grounded in sequential sampling theory (Wald, 1945; Ratcliff, 1978)—treats decisions as a process of noisy evidence accumulation until a response threshold is reached. This approach captures key empirical regularities in accuracy and RT across many perceptual tasks. However, its application to tasks such as TI, where performance is not directly tied to sensory input, remains contentious. The qualitative differences between TI and perceptual paradigms motivate a critical examination of whether the DDM can meaningfully capture decision behavior in this domain.

To address this, we applied a generalized drift-diffusion model (GDDM) using PyDDM (Shinn et al., 2020) to simultaneously fit accuracy and RT data from monkeys learning novel TI problems. The GDDM provided high-fidelity fits to both behavioral measures. Model estimates indicated that both the drift rate (reflecting the strength of evidence accumulation) and boundary separation (reflecting response caution) varied as a function of learning, with drift rate further modulated by symbolic distance. These findings contrast with those from other domains, where constant decision boundaries are typically sufficient (e.g., Voskuilen et al., 2017) and suggest that TI behavior is best captured by a "variable collapsing bound" model. Our

results thus extend the DDM framework to a novel cognitive domain, illuminating the temporal dynamics of decision-making based on internal structure rather than perceptual input.

Methods

Subjects

Subjects were three male rhesus macaques (F, H, and L), 13-14 years old and weighing 8-10 kg at time of experiment. The research was conducted in accordance with U.S. Department of Health and Human Services (National Institutes of Health) guidelines for the care and use of laboratory animals and was approved by the Institutional Animal Care and Use Committee at Columbia University and the New York State Psychiatric Institute. Monkeys were prepared for experiments by surgical implantation of a post for head restraint and a recording chamber to give access to prefrontal cortex. All surgery was performed under general anesthesia (isoflurane 1-3%) and aseptic conditions. Subjects had prior experience performing transitive inference tasks (Jensen et al. 2015).

Visual Stimuli and Eye Movements

The subjects were seated in an upright primate chair during the experiment and responded to visual stimuli by making saccadic eye movements. Stimuli were generated and controlled by a CRS VSG2/3F video frame buffer and displayed on a CRT monitor with a resolution of 1280 x

1024 pixels at 60 Hz. The viewing distance was 60 cm.

Visual stimuli consisted of 140-by-130 pixel color photographs (7° by 8° visual angle) and small squares that served as fixation and eye movement targets. To control the retinal position of visual stimuli, monkeys were required to maintain fixation before the stimuli were presented. The fixation point was a small red square (0.5° visual angle). NHP were allowed to initiate a saccade response any time after stimulus onset.

Eye position was recorded using a monocular scleral search coil (CNC Engineering, Seattle WA) and digitized at a rate of 1 kHz (Judge et al. 1980). Subjects expressed their choices by making eye movements from the central fixation point to peripheral target stimuli. Eye velocity was computed offline by convolving eye position with a digital filter. The filter was constructed by taking the first derivative of a temporal Gaussian, $G(t)$, such that $\frac{dG}{dt} = -k \cdot t \cdot exp\left(\frac{t^2}{\tau^2}\right)$, where $\tau = 8$ msec and k is a constant that sets the filter gain to 1.0. This filter does not introduce a time shift between the position input and velocity output, but adds temporal uncertainty to the velocity estimates. Horizontal and vertical eye velocities ($h'(t)$ and $v'(t)$, respectively) were combined to estimate radial eye speed $r'(t)$ using the formula $r'(t) = \sqrt{h'(t)^2 + v'(t)^2}$.

Transitive Inference Task

Prior to the beginning of each behavioral session, a set of 7 photographs never before presented to the subjects was selected from a database of over 2500 images. A different set of images was

used for each behavioral session. These stimuli were randomly assigned a unique rank, indicated by the letters A thru G, to create an ordered list (**Fig. 1A**). No explicit information about stimulus rank was presented to the subjects. That is, the rank assigned to each stimulus was never displayed, nor was there any information about serial order conveyed by the spatial or temporal presentation of the stimuli.

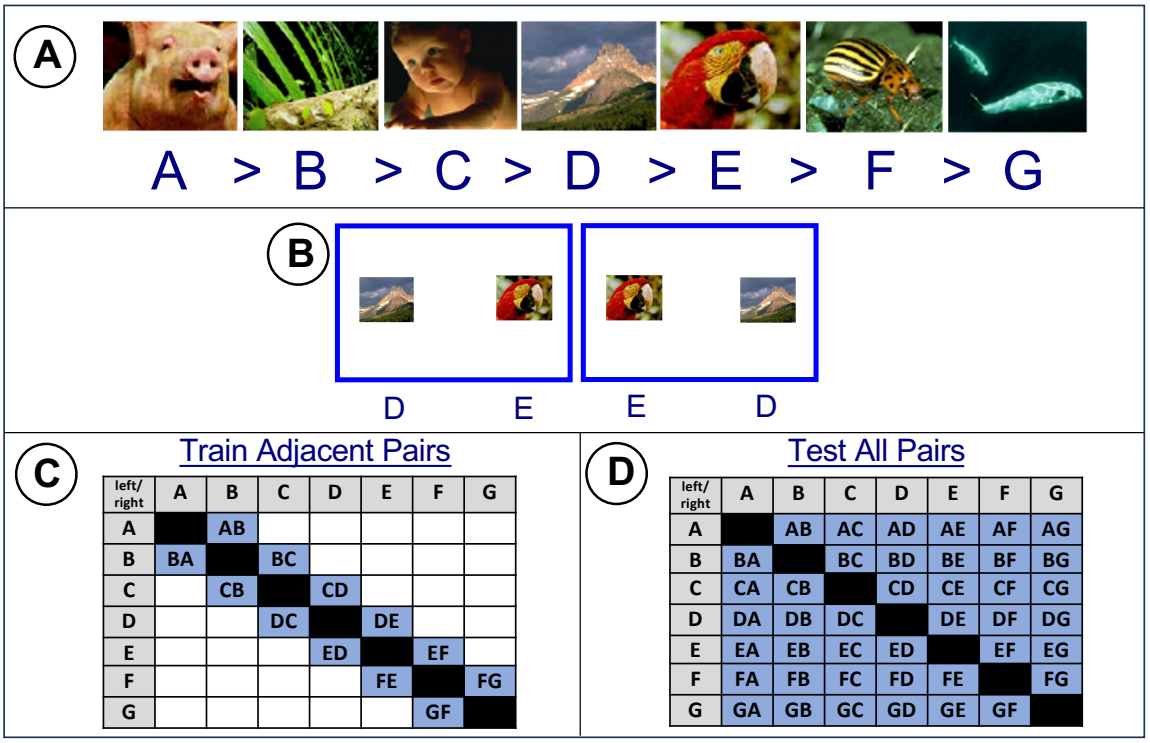

left/ right	A	B	C	D	E	F	G
A		AB					
B	BA		BC				
C		CB		CD			
D			DC		DE		
E				ED		EF	
F					FE		FG
G						GF	

left/ right	A	B	C	D	E	F	G
A		AB	AC	AD	AE	AF	AG
B	BA		BC	BD	BE	BF	BG
C	CA	CB		CD	CE	CF	CG
D	DA	DB	DC		DE	DF	DG
E	EA	EB	EC	ED		EF	EG
F	FA	FB	FC	FD	FE		FG
G	GA	GB	GC	GD	GE	GF	

Figure 1. **Transitive Inference Transfer Task.** This task tests subjects' ability to learn the implied serial order of a set of images from pairwise presentations. **A.** Set of 7 randomly selected images arranged in an ordered list (A>B>C>D>E>F>G). A novel set of images was used for each behavioral session. **B.** Choice trials with spatial position counterbalanced. **C.** Adjacent pairs used for the training phase of the session **D.** All pairs used for the transfer phase.

On each trial (**Fig. 1B**), the subject first fixated a small target in the center of the screen. Following a random delay (0.4 sec to 1.2 sec, positively skewed with a mean of 0.5 sec), the fixation point was extinguished, and two stimuli appeared equidistant from the center of the screen in opposite directions. To receive a reward, subjects had to select the stimulus with the

lower rank by making a saccadic eye movement to the selected stimulus within 1.5 sec of stimulus onset and fixating for a random interval (0.4 to 0.6 sec, uniformly distributed). When this interval elapsed, auditory feedback was delivered indicating that the response was either correct (high tone, 880 Hz) and would be rewarded, or was incorrect (low tone, 440 Hz). To receive the reward on correct trials, subjects had to maintain fixation for another interval (0.35 to 0.65 sec, uniformly distributed), after which the screen went dark and fluid rewards (juice or water drops) were delivered through a sipper tube.

Each session was divided into three phases: single target presentations, training, and testing. During the single target phase, only one stimulus, selected at random, was shown during each trial. The stimulus was shown on the left or right side of the screen and the NHP responded by making a saccadic eye movement to fixate the stimulus. There were approximately 5 blocks of single target presentations, or about 70 trials with left/right counterbalancing. Single target trials provided estimates of simple reaction times.

In the training phase of the task, adjacent pairs of stimuli were presented. Each pair comprised two stimuli whose ordinal positions were adjacent in the list (A/B, B/C, C/D, etc, **Fig. 1C**). To receive a reward, the subject was required to make a saccadic eye movement to the stimulus with the lower rank (e.g. A from A/B). For a 7-item list, there are 6 adjacent pairs. With positional counterbalancing, this resulted in training blocks that were 12 trials in length. This phase generally lasted up to 20 blocks (i.e. 240 trials in total).

In the testing phase of the task, all possible combinations of the 7 items were presented (**Fig. 1D**). The subjects were still required to saccade to the lower ranking stimulus. This

phase tested the subjects' ability to transfer knowledge about the list order gained during the adjacent pairs phase to novel, non-adjacent pairs. For a 7-item list, there are 21 possible pairs (6 adjacent, 15 non-adjacent). With counterbalancing of the stimulus positions, this resulted in blocks of 42 trials. This phase generally lasted for 5-10 blocks (240-420 trials).

Each pair of stimuli could be identified in terms of the ranks of the individual stimuli they included. Pairs could be also described in terms of "symbolic distance," a metric that is calculated by taking the difference between the individual stimulus ranks (D'Amato & Colombo 1990). For example, since the pair B/C consists of items that are adjacent in the list (i.e. ranks 2 and 3, respectively), the symbolic distance is 1, while the pair A/G (ranks 1 and 7) has a symbolic distance of 6.

Sequential Probability Ratio Test (SPRT)

The Sequential Probability Ratio Test (Wald 1945, Wald & Wolfowitz 1948) forms the theoretical basis of Drift Diffusion Models (DDMs). We used the SPRT to simulate evidence accumulation using parameters estimated directly from the behavioral data. This involved two assumptions: First that the subject's relative preference for each item is represented internally as a likelihood distribution on a one-dimensional scale and, second, that the mean and variance of each item's likelihood can be modeled as a binomial probability distribution whose parameters are related to choice probability.

To estimate the parameters of the binomial likelihood distribution for each item, every choice of that item was scored as a success and every time it appeared but was unchosen was

counted as a failure. The binomial distribution for each item thus represented the likelihood of choosing that item when it was presented. The set of binomial likelihood functions for all 7 items represents their relative position along a "mental line" (Riley 1976, Prado et al., 2012, Bonnefond et al., 2014, Jensen et al., 2015). The item positions reflect relative preference calculated from both correct and incorrect trials.

The likelihood functions allowed us to simulate evidence accumulation for each pair and its dependence on symbolic distance. For any pair of items, e.g. E and F, a random deviate, x, was drawn from the probability density, P(x|E,EF), for the correct item (E) in the pair (EF), where P(x|E,EF) = L(E|x,EF). Then the log of the likelihood ratio was computed, e.g. log[L(E|x,EF) / L(F|x,EF)]. The log of the likelihood ratio (LLR) for each sample corresponds to the momentary evidence in drift diffusion models. The probability density was repeatedly sampled and the LLRs were summed until they reached one of two thresholds (boundaries) that represented the amount of evidence accumulated prior to correct and incorrect choices. The repeated sampling and summing of momentary evidence simulates a drift-diffusion process. The number of samples needed to reach threshold is an estimate of reaction time. The proportion of trials in which the correct boundary is reached first is an estimate of performance accuracy (percent correct).

Generalized Drift Diffusion Modeling (GDDM)

Behavioral data were fit to a Generalized DDM using the PyDDM toolbox (Shinn et al., 2020). We used several different approaches to understand the effect of each variable on the dynamics

of the GDDM model in relation to serial learning. In all approaches, the data to be fit were reaction time and accuracy segregated by symbolic distance or trial block. The two conditions (symbolic distance or trial block) were fit separately. Parameters were estimated for each NHP subject individually.

Standard model. In the first, "standard," approach, a minimal version of the GDDM was implemented with four parameters: drift rate (rate of evidence accumulation), boundary level (amount of accumulated evidence required for a decision), noise (instantaneous variability in drift rate), and non-decision time. The time evolution of the decision variable, x (accumulated evidence), was given by:

Eqn 1. $$dx = [drift * condition] * dt + dW$$

"*Drift*" is the drift rate parameter. "*Condition*" is either symbolic distance or trial block, depending on how the data are segregated, and is coded as an integer; [1…,6] for symbolic distance or [1,…,8] for trial block. "*W*" is the momentary noise represented as a Weiner process with standard deviation = 1. We fixed the noise parameters to $mu = 0$ and $sigma = 1.0$ to avoid redundancy with the boundary height parameter as scaling noise variability is equivalent to scaling the boundary height. The time resolution (dt) was 10 msec and the overall duration was 1.0 sec.

Table 1. PyDDM parameter ranges.						
Range	**Drift**	**Boundary**	**Noise**	**Leak (if present)**	**Tau**	**Non-decision**
Standard model						
min	0.0	0.1	1.0	0	-	50 msec
max	10.0	10.0	1.0	100	-	400 msec
Collapsing bound model						
min	0.0	0.3	1.0	-5.0	0.01	50 msec
max	5.0	3.0	1.0	5.0	10	400 msec

For this approach, the fitting algorithm searched for a single set of parameters (drift rate, bound, and non-decision time) that fit the data within each condition. All parameters except for noise were allowed to vary freely within a wide range (**Table 1**). The drift rate was the product of the scalar drift parameter and the condition index, and therefore varied linearly with condition. The boundary was constant both in time, and across conditions. The same parameter ranges were used for all NHP and all conditions.

Leaky integration. In the standard approach, there was no loss of evidence over time. However, evidence integration can be imperfect. Thus, in a second, "leaky integration" approach, imperfect integration was modeled by adding a leak term to the standard model:

Eqn 2. $$dx = [drift * condition + leak * x] * dt + dW$$

The *leak* is a constant that multiplies accumulated evidence (x) at each moment. If the leak term is negative, the evidence tends to decay exponentially toward zero. If the leak is positive, it acts like a gain factor that causes the evidence to increase exponentially (unstable integration).

Collapsing bound. In both the standard and leaky integration approaches, the decision boundary remained constant while evidence was being accumulated. To understand the effect

of allowing the boundary to change over time, a "collapsing bound" approach was tested. In this approach, the bound collapsed exponentially, thus introducing a sixth parameter, *tau*, the collapse rate. In this approach, the equation for momentary evidence was the same as **Eqn 1**, but the time-dependence of the boundary was given by:

Eqn 3. $$bound = B_0 * exp(-t/tau)$$

Where B_0 is the initial boundary level at $t = 0$. This model was tested both without a leak parameter (using **Eqn 1** for evidence accumulation) and with a leak (using **Eqn 2**).

Simulations. To illustrate how each model parameter affected dynamics and performance, simulations of the standard, leaky integration, and collapsing bound models are shown in **Fig 2**. The left panel illustrates the time evolution of accumulated evidence for a range of drift rates (signal strengths). The center panel shows the reaction time histograms for correct and error trials. The right panel shows accuracy and reaction time. Non-decision time was set to zero in all simulations.

For the standard model (**Fig 2**, top row), low signal strength was associated with chance performance accuracy and long reaction times. The addition of a positive leak term (**Fig 2**, middle row) restricted the range of reaction times by speeding up evidence accumulation, but also reduced accuracy for moderate signal strengths. The collapsing bound model (**Fig 2**, bottom row) performed similarly to the leaky integration model but was more effective in reducing long reaction times and produced more symmetric reaction time distributions.

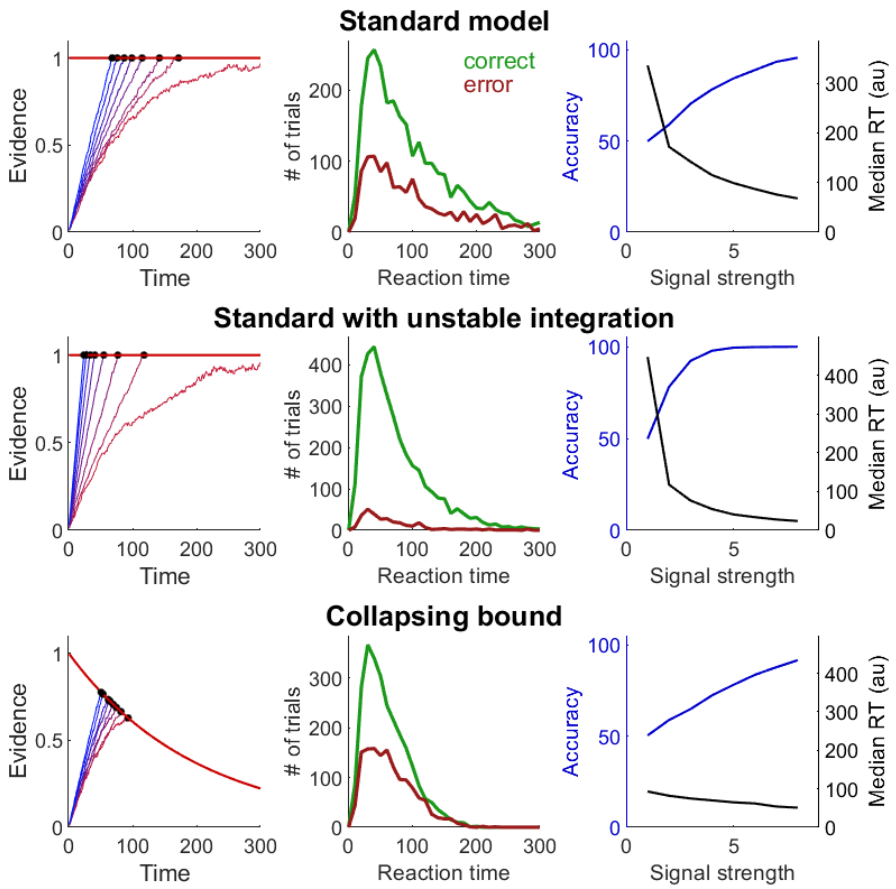

Figure 2. DDM Simulations. Top: Standard model with constant bound and no leak. Left: red traces represent weaker signal strengths, blue is stronger signal. Center: reaction time histograms for correct (green) and error (red) trials for a single, moderate signal strength. Right: Accuracy (percent correct) and reaction time as a function of signal strength. **Middle:** Standard model with leak. **Bottom:** Collapsing bound model.

Results

Behavioral Performance

Performance was quantified in terms of accuracy (percent correct) and response time (time from stimulus onset until choice saccade onset). A total of 273 behavioral sessions were recorded (F: 85, H: 98, L: 90 sessions). Each session started with a novel set of pictures that were arranged into a list by assigning a rank order to each picture (**Fig. 1A**). The ranking was initially unknown to the subject but was learned over the course of the session.

Reaction times (RT) as a function of trial number for all sessions completed by NHP H are shown in **Fig. 3A**. The three phases of the task are color-coded (single target trials in gray, adjacent pairs in blue, all pairs in magenta). There was an abrupt shift in the RT distribution (**Fig. 3B**) at the transition from single stimuli (median RT 192 msec) to stimulus pairs (median RT 270 to 275 msec).

Adjacent-pair (training phase) and all-pair (testing phase) trials were grouped into blocks of 42 consecutive trials to obtain robust estimates of RT (**Fig. 3A**). The shape and

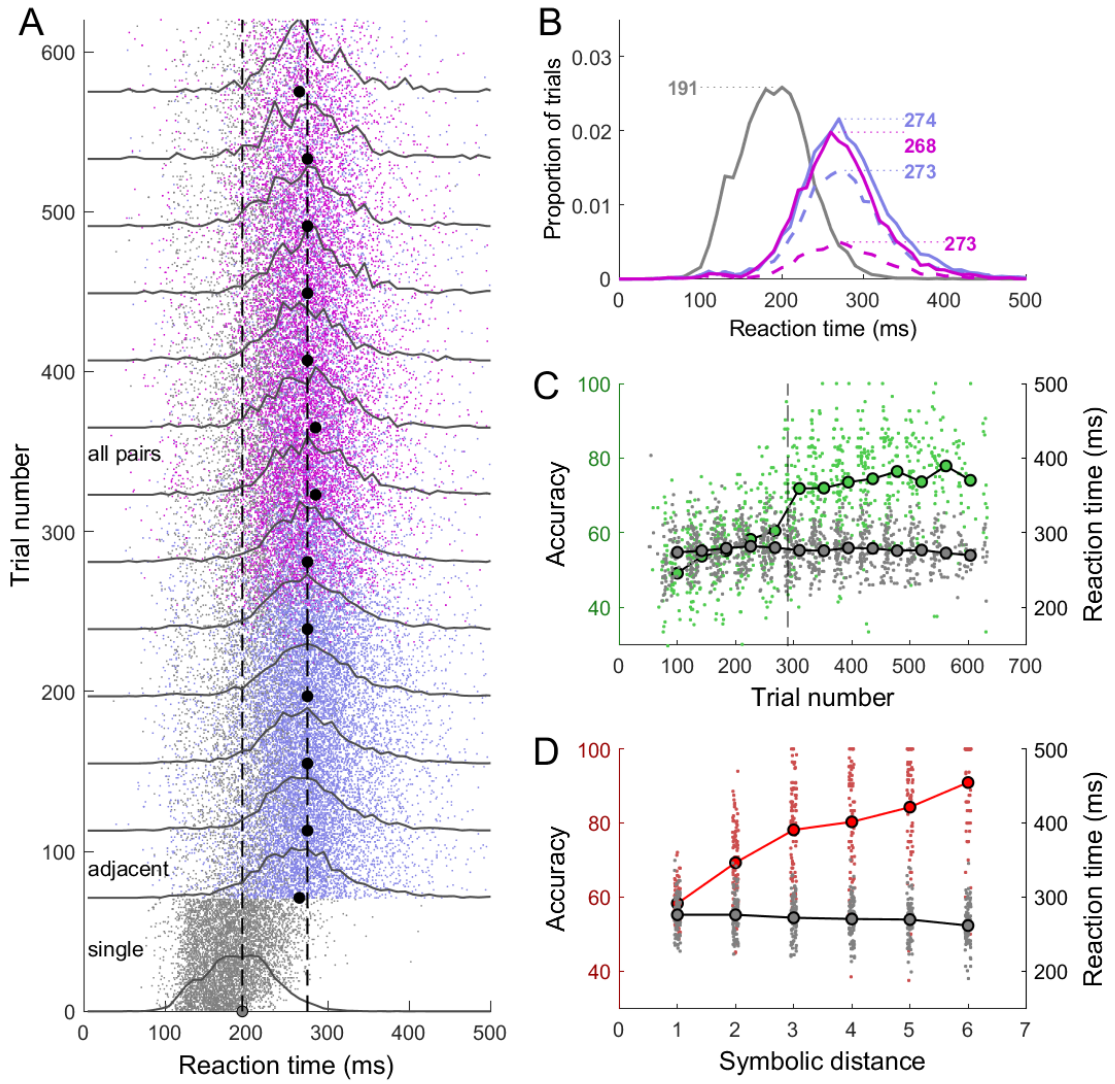

Figure 3. **Behavioral Performance on TI transfer task. A.** Reaction time vs trial number for NHP H. The colors represent the three task phases (blue – presentation of single items, cyan – adjacent pairs, red – all pairs). The gray dot on the x-axis and the thin dashed line is the median RT for single target trials. The black dots are the median RTs for choice trials binned in 50 trial segments. The thick dashed line is the median over all choice trials. **B.** Reaction time histograms. Colors as in A. For choice trials, solid lines are correct responses and dashed lines are incorrect. **C.** Accuracy (green) and reaction time (gray/black) during learning. Dashed vertical line is transition from training to testing phase. **D.** Accuracy (red) and reaction time (grey/black) as a function of symbolic distance.

median value of the RT distributions remained constant across trial blocks. The RT distributions were similar in shape for correct (**Fig 3B** solid lines) and incorrect (dashed lines)

trials, and for adjacent (**Fig. 3B,** blue lines) and non-adjacent pairs (magenta lines).

The effect of learning on accuracy and RT for this NHP is shown in **Fig. 3C**. During adjacent pair training, logistic regression showed a significant effect of trial number on accuracy for all NHPs (F: $p = 0.03$, df = 15,206; H: $p < 0.0001$, df = 16,012; L: $p < 0.0001$, df = 6658). Logistic regression coefficients for accuracy versus trial number were similar for all NHP (0.096 for NHP F, 0.07 for NHP H, 0.11 for NHP L). The effect of trial number on RT, as assessed by linear regression, was also significant ($p < 0.0001$ for all NHPs). However, the slopes of the regression lines were shallow, amounting to an increase of 5 to 7 msec for every 100 trials.

During all-pairs testing, the relationship between trial number and accuracy remained significant ($p < 0.0001$ for all NHP) and the RT trend turned slightly negative (-0.5 msec per 100 trials) and was significant for 2 NHP ($p < 0.002$). However, the relationship was inconsistent - negative for one NHP (H: decrease of 1.6 ms/400 trials), and positive for the other (L: increase of 4 ms/400 trials).

Symbolic distance (SD, the difference in rank between the two stimuli in each pair) had a strong effect on accuracy (**Fig. 3D**). Logistic regression coefficients for SD were consistent across NHP (0.35 for NHP F, 0.28 NHP H, 0.22 for NHP L). For RT, there was a small but significant reduction with SD for both NHPs (F: slope = -1.5 msec, $p < 0.0001$; H: -2.9 msec, $p < 0.0001$; L: -0.33 msec, $p = 0.02$). The results overall indicate a strong effect of trial number and SD on performance accuracy, but little effect on response time.

Accuracy and Reaction Time at Transfer

The transition from adjacent pair training to all-pairs testing was not signaled to the NHP while they were performing the task. Rather, non-adjacent pairs were simply added in to the mix of trials. The start of the testing phase is defined as the occurrence of the first non-adjacent pair after adjacent-only training. Performance accuracy showed an abrupt effect of symbolic

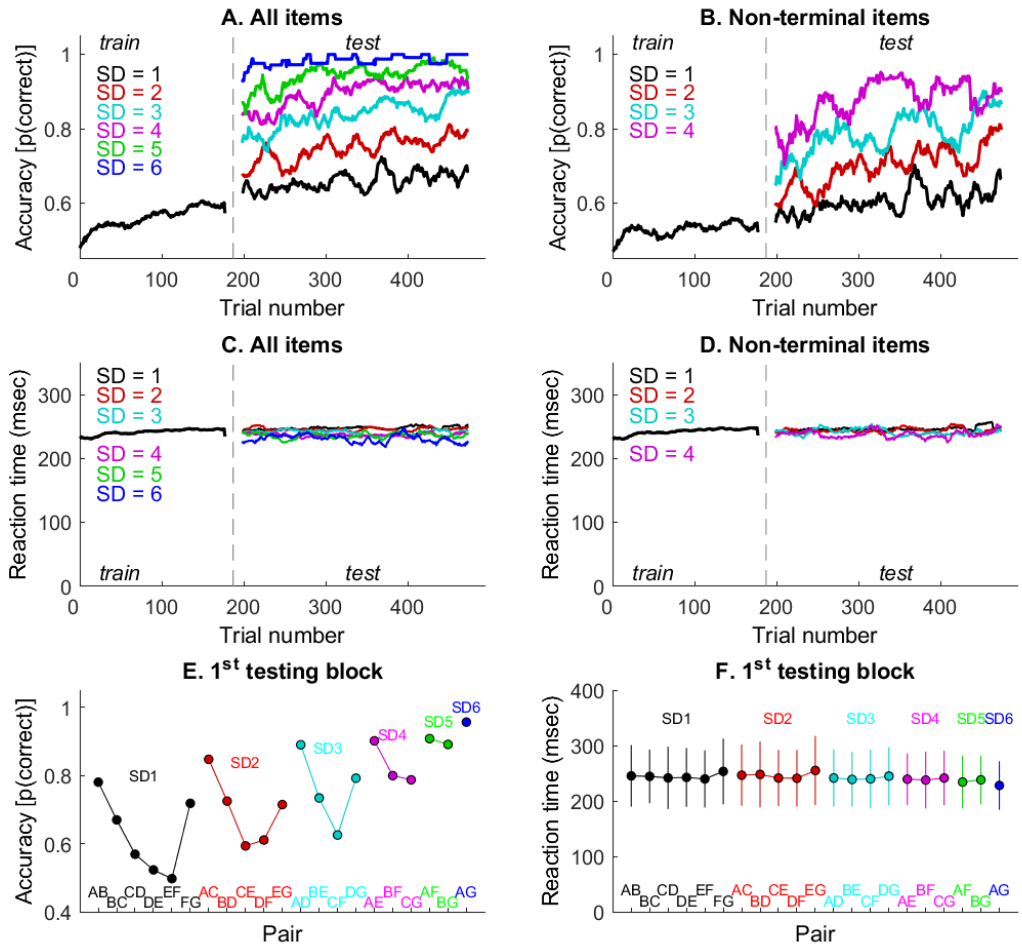

Figure 4. Accuracy and reaction time at transfer. A. Accuracy vs. trial during training and testing averaged over all NHP and sessions. **B.** Accuracy for pairs with non-terminal items. **C.** Reaction time during training and test. **D.** Reaction time for critical pairs. **E.** Accuracy for each pair during the first block of testing. **F.** Reaction time for each pair during the first block of testing.

distance from the first testing trial onward. **Fig. 4A** shows training and testing accuracy averaged over all sessions for the 3 NHP. After ~200 training trials, accuracy was close to 60% correct. At the start of all-pairs testing, accuracy for adjacent pairs (**Fig. 4A** black lines follows the same trajectory as training. However, at the beginning of testing, accuracy for non-adjacent pairs (**Fig. 4A** colored lines) was superior to adjacent pairs for each symbolic distance even though all non-adjacent pairings were novel. We have documented this remarkable jump in performance at transfer with other NHP and other testing systems (Jensen et al., 2015). The same pattern of results was obtained even when terminal pairs (those containing items A or G) were excluded (**Fig. 4B**). By contrast, reaction times showed minimal variation across training and testing (**Fig. 4C,D**) To examine performance at transfer more closely, the accuracy (**Fig 4E**) and reaction time (**Fig. 4F**) for each item pair during the first block (42 trials) of testing are shown. During a block of 42 trials, each pair was shown twice; once in each spatial configuration.

The test pairs, BD, CE, DF, BE, CF, and BF are ”critical” in the sense that they did not occur during training nor do they contain terminal items. Above chance performance on these pairs, as well as the overall effect of symbolic distance, is considered evidence that NHP relied to some extent on transitive relations in making their choices. For example, during training, B and F were chosen almost equally often from the pairs BC and FG (in fact F was correctly chosen from FG, and therefore rewarded, more often than B was correctly chosen from BC). Yet, during the first block of testing, B was chosen over F in roughly 80% of the BF trials.

Speed-accuracy trade-off

Although there was not a clear relationship between accuracy and choice reaction time when data were combined over sessions (**Figs 3C** & **3D**), a more fine-grained analysis revealed subtler trends. In this analysis, we calculated the residual accuracy and average reaction time in each block of 42 consecutive trials during the adjacent-pair training and all-pairs testing. "Residual" is here defined as the accuracy or average reaction time within each block of 42 trials normalized by subtracting the averages of each of those measures calculated over all choice trials for the entire session. By compensating for session-to-session variability, the residuals provided estimates of the change in accuracy per unit change in reaction time that could be combined across sessions.

Figs. 5A,B,C show an improvement in accuracy with increasing reaction times during the training phase of the task (first 200 choice trials of each session) for all three NHP. Each data point represents a single 42-trial block from one session. For every additional millisecond of reaction time, there was an average increase in performance ranging from 0.24 to 0.43 percent correct. The correlation between accuracy and RT was significant for each NHP. This is opposite to the typical pattern for perceptual decisions where increasing accuracy is associated

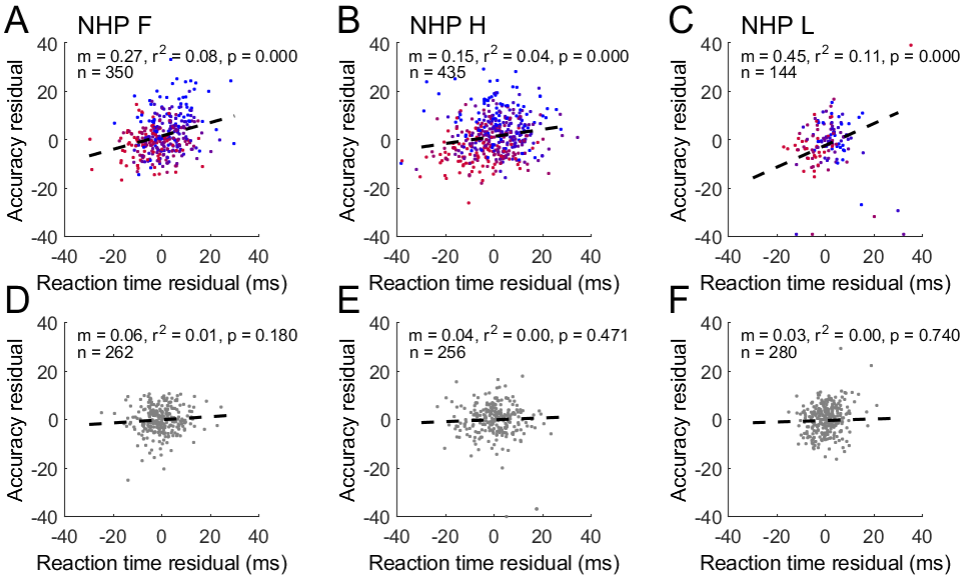

Figure 5. Speed-accuracy trade-off. A,B,C. Improvement in accuracy v. increase in RT during training phase for each NHP. 'm' is the linear regression slope, 'r^2' is the coefficient of determination, 'p' is the significance of the correlation (Pearson), and 'n' is the number of 50-trial blocks per session times the number of sessions. The color gradient indicates trial blocks that were earlier in training (red) vs. later (blue). **D,E,F.** All-pairs testing phase.

with shorter reaction times. During the subsequent testing phase (the first 200 trials after end of training)(**Fig. 5 D,E,F**), no significant relationship between accuracy and reaction time was found for any NHP.

During the all-pairs testing phase, we sorted trials into six groups by symbolic distance (SD), rather than by trial block. When this was done, a significant negative relationship was found between the change in accuracy and increasing reaction time (**Fig. 6 A,B,C**); trials with larger SD had higher accuracy and shorter RTs.

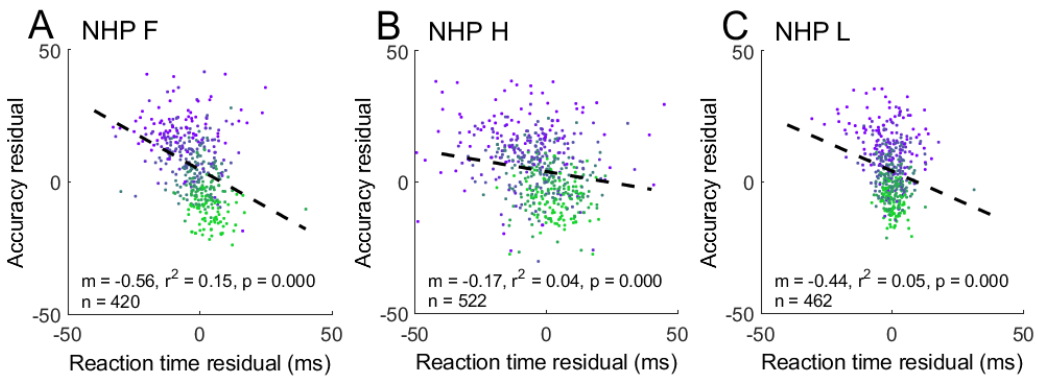

Figure 6. Speed-accuracy trade-off for trials grouped by symbolic distance. A,B,C. Decrement in accuracy v. increase in RT during testing phase for each NHP. 'm' is the linear regression slope, 'r^2' is the coefficient of determination, 'p' is the significance of the correlation (Pearson), and 'n' is the number of symbolic distances times the number of sessions. The color gradient indicates smaller (green) vs. larger (magenta) symbolic distances.

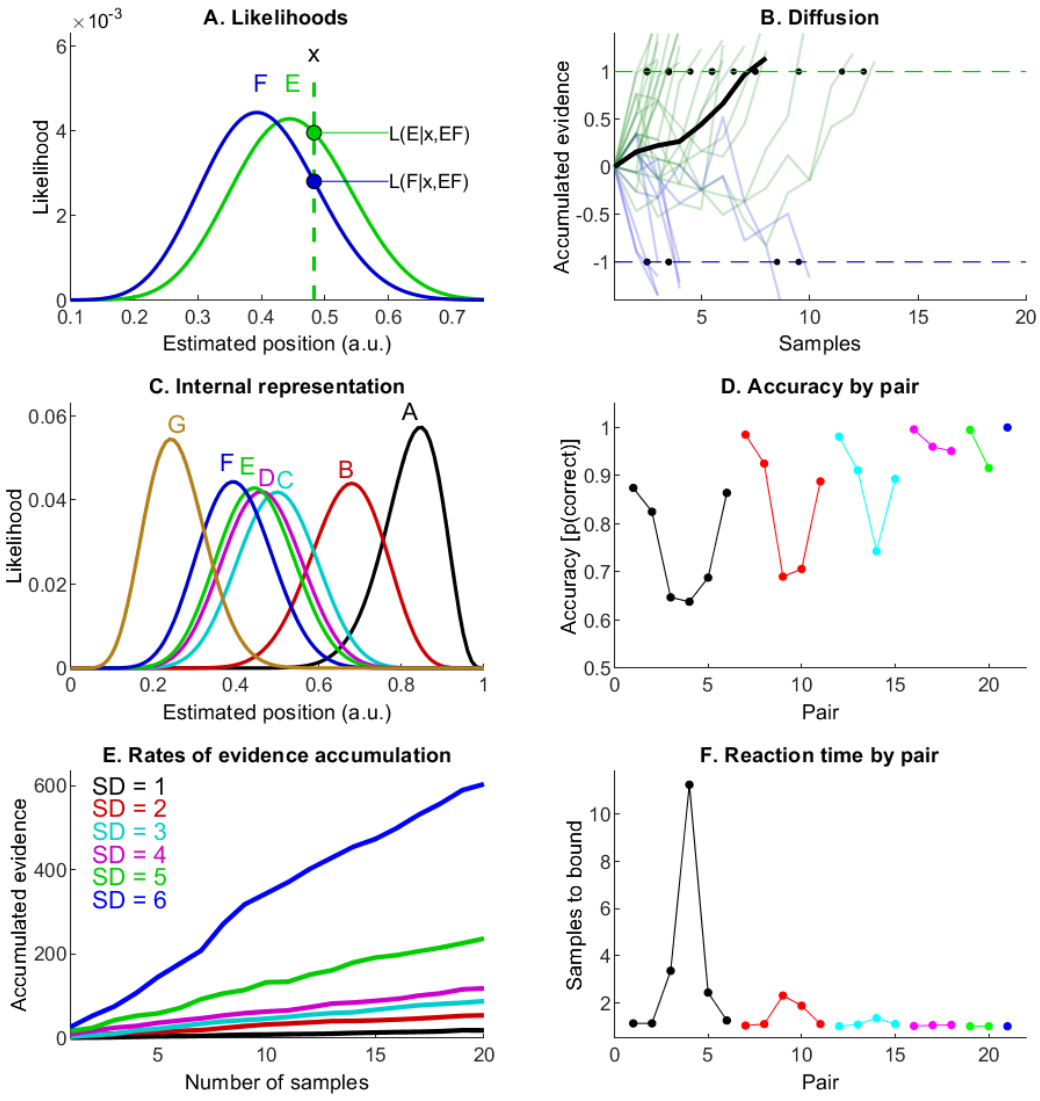

Figure 7. Sequential probability ratio test (SPRT). A. Binomial likelihood functions for items E and F. **B.** Accumulation of evidence for the pair EF based on repeated sampling of relative likelihood. Upper bound (green dashed line) corresponds choices of E, lower bound (blue) corresponds to choosing F. Solid black line is the average accumulated evidence. **C.** Binomial likelihood functions for all items. **D.** Accuracy for each pair predicted by SPRT. **E.** Evidence accumulation for different symbolic distances. **F.** Reaction time (samples to bound) predicted by SPRT.

354

355 ***Rate of evidence accumulation estimated from sequential sampling.***

356 An optimal method for making binary decisions is to repeatedly sample the relative likelihood

357 of the two alternatives (Wald & Wolfowitz 1948) until the sum of the log-likelihood ratios
358 reaches a predetermined value (bound). We applied this method to determine the optimal rate
359 of evidence accumulation (i.e. diffusion) in the TI task. Because the NHPs made a binary
360 choice on each trial, it is reasonable to model their choices using a set of binomial likelihood
361 distributions, one for each item (Jensen et al., 2015). The parameters of those distributions
362 (success rate, p, and number of trials, n) correspond to the number of times each item was
363 chosen relative to the number of times it was presented (choice probability). This is
364 demonstrated for the pair EF in **Fig. 7A** which shows the likelihood functions estimated from
365 the pattern of choices made at the end of the training phase of the task averaged over all 3
366 NHP.

367 The likelihood functions for each item were used to simulate a sequential sampling
368 decision process (**Fig. 7B**) with constant bounds. The rate of evidence accumulation and the
369 momentary noise were determined by the means and widths of the likelihood functions for E
370 and F, combined with random draws from the likelihood function for E. The upper and lower
371 bounds were set to approximate the observed probability of choosing E from the pair EF.

372 The likelihoods for all 7 items, estimated from choice probabilities at the end of testing,
373 are shown in **Fig. 7C**. We propose that this set of functions forms an internal representation of
374 the ordering of the TI list and the degree of uncertainty associated with the relative position of
375 each item (Jensen et al., 2015). Simulations were run to estimate the rate of evidence
376 accumulation for each pair during the testing phase. Using a fixed bound, we could estimate
377 the accuracy and reaction time (number of samples to bound) for each pair. The pattern of

accuracy vs. pair (**Fig. 7D**) was qualitatively similar to the behavioral data (**Fig. 4C**), thus validating the approach, at least in part.

Evidence accumulation rates were combined over pairs to estimate the rate of evidence accumulation as a function of symbolic distance (**Fig. 7D**). Evidence accumulation was approximately linear with the number of samples, reflecting the fact that, on average, each sample provided the same amount of momentary evidence regardless of when it was drawn. The rate of accumulation (drift rate) increased with symbolic distance.

It is obvious that there is no constant boundary level that could be reached after a fixed number of samples for all symbolic distances. This is confirmed by plotting the number of samples (a proxy for reaction time) needed to reach a bound that was fixed for all pairs (**Fig. 7F**). Reaction times produced by this version of the SPRT were inversely related to accuracy, a result that does not match the behavioral data (**Fig. 4D**).

In general, the accuracy and reaction times obtained with a simple but theoretically well-motivated sequential sampling model are at odds with the experimental data and suggest that a fuller treatment is needed. We therefore turned to the PyDDM package to test more elaborate models.

PyDDM Results

Choice reaction times during the TI transfer task were generally under 300 msec even though subjects were allowed up to 1500 msec after stimulus onset to respond by making a saccadic eye movement to the target or distractor. Choice RTs averaged only 35 msec longer than simple RTs for NHP F (215 vs. 180 msec median RT), 85 msec for NHP H (275 vs. 190 msec), and 30 msec for NHP L (211 vs. 181 msec).

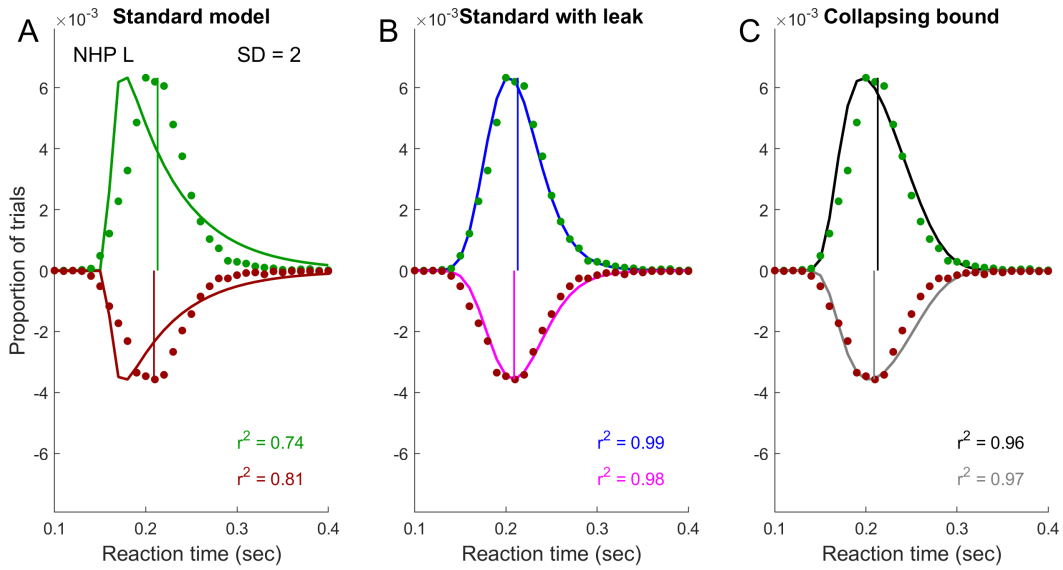

Figure 8. Generalized Drift Diffusion Model fits. Data from all sessions for NHP L with symbolic distance (SD) equal to 2. Green dots are reaction time distributions for correct trials, red dots for error trials. Vertical red and green lines indicate the median RT for each distribution. Green and red curves are fits for correct and error trials, respectively, using the standard model. The correlation between the fits and the data are shown in red and green text. Blue and magenta curves and text are fits obtained with the leaky integration model. Black and grey curves and text are fits obtained with the collapsing bound model.

RTs remained nearly constant despite learning-related improvements in accuracy. RTs were also constant across symbolic distance despite large differences in accuracy. Drift diffusion models (DDMs) with constant bounds and variable drift rate typically show increases in accuracy that are accompanied by shorter reaction times. As the current results did not show this pattern, it raises the question of whether DDMs can fit the behavior exhibited by NHP during the transitive

Table 2. PyDDM fitted parameters for each model, NHP, and condition. For "condition," SD is symbolic distance, TB is trial block. Model 1 is the standard model with constant bound and no leak. Model 2 is the standard model with a leak term. Model 3 is the collapsing bound model with no leak. Model 4 is the collapsing bound with leak. BIC is Bayesian information criterion. VAC is variance accounted for (R^2). Bolded lines indicate models with the highest VAC.

NHP - Condition	Model	Leaky Integration	Collapsing Bound (Tau)	Drift	Bound	Noise	Non-decision	BIC	VAC
NHP F - SD	1	-	-	0.68	0.30	1.0	152 msec	-73232	0.78
NHP F - SD	2	43.1	-	0.99	9.0	1.0	105 msec	-85582	0.96
NHP F - SD	3	-	0.15	0.55	0.58	1.0	122 msec	-83621	0.93
NHP F - SD	4	-4.9	0.14	0.50	0.58	1.0	118 msec	-83706	0.93
NHP H - SD	1	-	-	0.61	0.35	1.0	184 msec	-39867	0.75
NHP H - SD	2	29.8	-	0.83	8.6	1.0	116 msec	-52508	0.96
NHP H - SD	3	-	0.15	0.50	0.88	1.0	126 msec	-48531	0.91
NHP H - SD	4	-4.7	0.15	0.45	0.92	1.0	110 msec	-48737	0.92
NHP L - SD	1	-	-	0.57	0.30	1.0	154 msec	-103034	0.73
NHP L - SD	2	69.4	-	0.93	8.2	1.0	133 msec	-120970	0.97
NHP L - SD	3	-	0.15	0.47	0.46	1.0	137 msec	-118035	0.92
NHP L - SD	4	-3.9	0.10	0.47	0.47	1.0	140 msec	-118598	0.93
NHP F - TB	1	-	-	0.29	0.30	1.0	151 msec	-54357	0.79
NHP F - TB	2	44.5	-	0.41	9.9	1.0	100 msec	-64415	0.96
NHP F - TB	3	-	0.15	0.25	0.65	1.0	116 msec	-57016	0.95
NHP F - TB	4	-4.9	0.14	0.50	0.58	1.0	118 msec	-83706	0.95
NHP H - TB	1	-	-	0.27	0.35	1.0	187 msec	-34058	0.74
NHP H - TB	2	28.6	-	0.36	9.5	1.0	105 msec	-45301	0.96
NHP H - TB	3	-	0.15	0.23	0.99	1.0	125 msec	-38285	0.91
NHP H - TB	4	-4.8	0.15	0.21	0.92	1.0	107 msec	-38316	0.91
NHP L - TB	1	-	-	0.28	0.25	1.0	156 msec	-76581	0.73
NHP L - TB	2	71.2	-	0.46	8.2	1.0	133 msec	-90833	0.96
NHP L - TB	3	-	0.10	0.22	0.86	1.0	116 msec	-80666	0.93
NHP L - TB	4	-3.7	0.10	0.24	0.54	1.0	132 msec	-80857	0.93

inference task and thus shed light on how decisions are made during TI learning.

To test this, we first fit a generalized DDM with four parameters (drift rate, noise, boundary, and non-decision time, see Methods) to the current behavioral data using the PyDDM package (Shinn et al. 2020). An example of the PyDDM fits for data segregated by

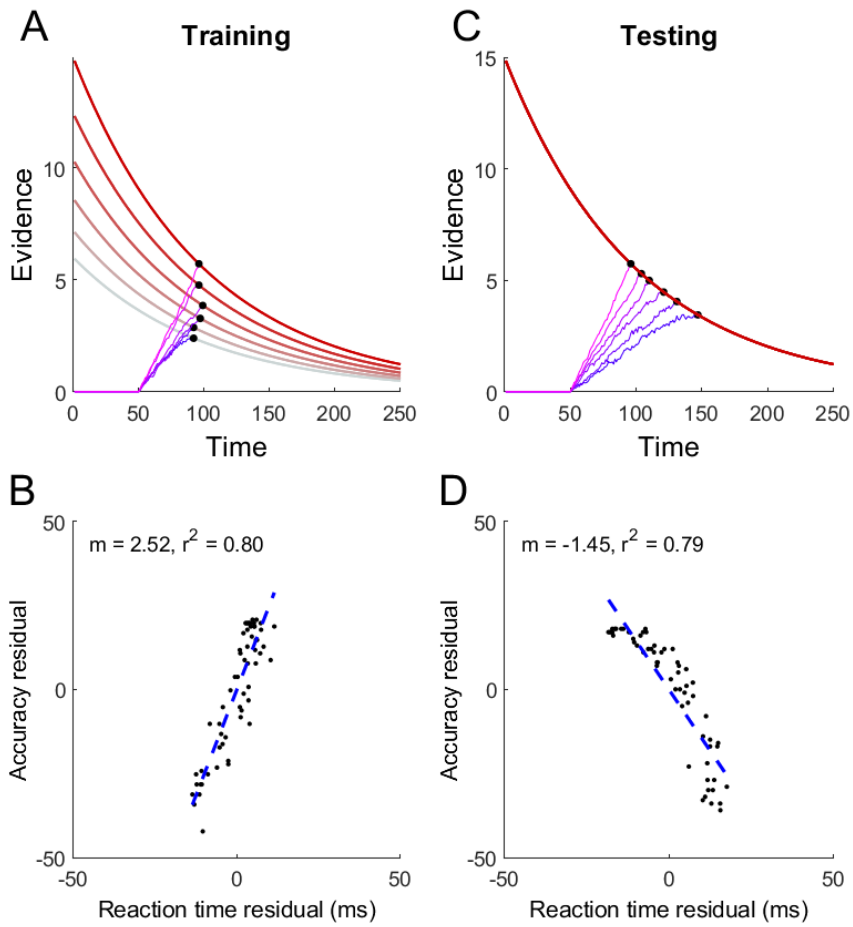

Figure 9. Variable collapsing bound model. A. Initial boundary position increases with signal strength during training. **B.** Positive correlation between residual reaction time and accuracy. **C.** Initial boundary position stabilizes after training. **D.** Negative relationship between reaction time and accuracy during testing.

outcome (green = correct, red = error) for NHP L is shown in **Fig. 8** (the fits for all NHP are shown in supplemental **Figs. S1** and **S2**). The fits for the standard model are shown as the green and red curves. These fits were relatively poor. This was confirmed by the variance

accounted for (VAC, measured by r^2), which was generally less than 80%. The fitted parameter values are shown in **Table 2**.

To estimate the effect of within-session learning on model fits, the data were grouped into 8 blocks of 42 consecutive trials spanning the first 400 trials of paired presentations. During these trials, accuracy improved from chance levels to roughly 80% correct. The variance accounted for (r^2) was generally $< 80\%$ for each trial block condition for all NHP (**Fig S2**). These results indicate that a DDM with constant bounds and no leak term provided a poor fit to data in which accuracy improves across conditions while RTs remain relatively constant. We next explore the role of each model parameter by testing variants of the standard DDM.

Effect of Drift Rate. The drift term captures whether the rate of evidence accumulation is sensitive to the stimulus conditions, as indexed by symbolic distance, or stage of learning, as indexed by trial block. The fitted parameters indicate an increased rate of evidence accumulation (drift rate) as a function of both SD and trial block (TB). This is expected as drift rate has a strong effect on accuracy. The drift rate estimates were similar across NHP. However, there was a factor of 2 difference between conditions, most likely because performance accuracy reached higher asymptotic levels as a function of symbolic distance (**Fig. 3D**) than it did for trial block (**Fig. 3C**).

Effect of Leak. The leak term modulates the rate of evidence accumulation. It can be thought of as a gain control that dampens or enhances the effect of drift rate and has a more powerful effect as more evidence is accumulated. Adding a leak term to the standard model improved the fits, as shown by the blue and magenta curves (**Fig. 8**). These fits were

substantially better than the standard non-leaky model, resulting in VACs in the range of 88% to 97%. The Bayesian Information Criterion (BIC) scores were also improved (**Table 2**).

The leak term can either slow or accelerate the rate of evidence accumulation depending on its sign. For the standard model with a leak, the fitted leak term was positive, indicating unstable integration. The magnitude of the leak term was quite large such that it played a dominant role in the model's performance, minimizing the effect of drift rate and boundary.

Effect of Boundary. The initial boundary level sets the amount of accumulated evidence needed to render a decision, whereas the collapse rate determines if that amount of evidence is stable over time. The collapsing bound model with no leak provided fits that were as good as the standard (constant bound) model with leaky integration, as measured by VAC and BIC. Adding a leak term to the collapsing bound model further improved the fits but only by a small amount (**Table 2**).

When a leak term was added to the collapsing bound model, the fitted values were negative, indicating true leaky integration (loss of evidence over time). This contrasts with the function of the leak parameter in the constant bound model, which was positive, indicative of unstable integration.

Effect of Non-decision Time. All NHP had non-decision times that were similar across conditions. However, there were differences among the NHP. Non-decision times for NHP F and L were on average 30-32 msec shorter than NHP H. It is of interest that the simple RTs for NHP F and L had the same average (181 msec), while H had the longest simple RTs at 192 msec suggesting that simple RTs may be related to non-decision time.

455 Speed-accuracy trade-off. The modeling approaches considered so far have been
456 evaluated on the basis of their ability to fit average accuracy and reaction time data. We have
457 not yet considered the relationship between residual accuracy and reaction time (i.e. the speed-
458 accuracy trade-off) shown in **Figs. 4** and **5**. The striking feature of these trade-offs is that they
459 have positive slope as a function of trial block during training and negative slope as a function
460 of symbolic distance during testing.

461 The standard DDM tends to produce a negative relationship between accuracy and
462 reaction time as indicated in **Fig. 2.** One way to turn this into a positive relationship is to
463 allow the boundary to vary. If the boundary height increases with signal strength, then
464 accuracy and reaction time can increase together. The relationship between boundary position
465 and signal strength that leads to increasing reaction times can be calculated analytically
466 (Bogacz et al 2006, 2008). These calculations reveal that accuracy is related to the product of
467 drift rate and boundary position, while reaction time is more closely related to their ratio.

468 **Fig. 9A** illustrates a "variable collapsing bound model" in which an increasing initial
469 boundary position leads to reaction times that increase with signal strength (**Fig. 9B**). When
470 the initial boundary position is fixed (**Fig. 9C**), reaction time decreases with signal strength
471 leading to a negative relationship between accuracy and reaction time (**Fig. 9D**).

472 In the collapsing bound model, initial boundary position was constant across trial blocks.
473 We tested a version of the collapsing bound model in which initial boundary position was free to
474 vary in each trial block. In this case, the fitted initial boundary position tended to increase
475 during training for all 3 NHP (**Fig. 10A**). This was true whether there was a leak term

included in the model or not. When each symbolic distance was fit separately (**Fig. 10B**), the initial bound position was constant across symbolic distance.

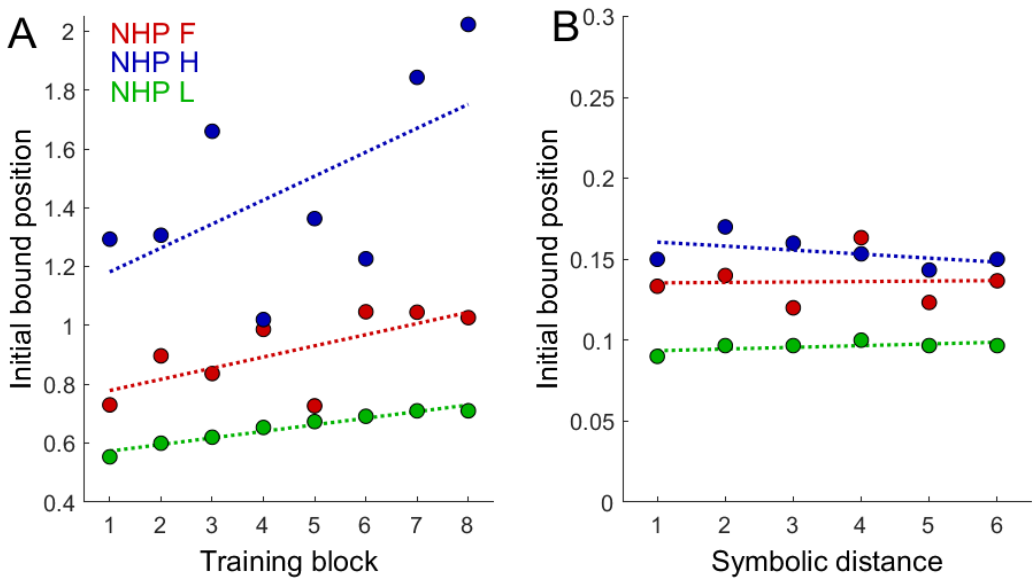

Figure 10. Variable collapsing bound fitted to training and testing data. A. Initial boundary position fit independently for each training block for all 3 NHP. Solid symbols are parameters for a model with a leak term. Dashed lines are linear fits to these data. **B.** Initial boundary position fit to each symbolic distance. Same conventions as **A**.

Discussion

The transitive inference (TI) paradigm assesses a subject's ability to infer the serial order of a set of stimuli based solely on pairwise comparisons, without explicit spatial or temporal cues. In this task, correct responses are rewarded for selecting the item of lower rank from a pair of images. When presented with a novel set of items, nonhuman primates (NHPs) initially performed at chance, but accuracy improved over the course of ~200-300 training trials. Accuracy then increased dramatically at the transition to all-pairs testing and showed a

symbolic distance effect for non-terminal pairs, suggesting that subjects had acquired an internal representation of the list's ordinal structure. How this internal structure and other latent variables contribute to decision making in this task was the subject of the current study.

To elucidate the mechanisms underlying TI decisions, we examined both accuracy and reaction time (RT) through the lens of evidence accumulation models. We found that choice RTs were remarkably fast—generally under 300 milliseconds and only marginally longer (by 30–85 ms) than simple reaction times—suggesting that decisions were made rapidly even though there was little time pressure. Despite these rapid responses, performance improved reliably with training, implying that learning and emergence of an internal representation were supported even within this brief time window. During learning, gains in accuracy were associated with longer reaction times, a pattern opposite to that typically found for perceptual decisions.

As expected from previous studies, symbolic distance (SD)—the rank difference between paired stimuli—was a strong predictor of accuracy: larger SDs led to higher performance. However, the impact of symbolic distance on RT was minimal, consistent with prior work (Munoz et al., 2020). This dissociation between accuracy and RT contrasts with findings from perceptual decision-making paradigms, where RTs are typically longer (500–1000 ms) and inversely related to decision difficulty or accuracy (e.g., Roitman & Shadlen, 2002; Ratcliff et al., 2003; Kim & Shadlen, 1999; Kiani et al., 2014; Jun et al., 2021). In those studies, RT-accuracy trade-offs are commonly interpreted as signatures of gradual evidence accumulation, as formalized in sequential sampling models like the drift diffusion model (DDM).

We initially explored whether a simple sequential sampling process could explain the

pattern of results observed in the TI task. Simulations based on binomial likelihoods yielded an evidence accumulation rate highly sensitive to symbolic distance (**Fig. 7D**). However, these models failed to capture the relative invariance of RT across SD conditions in the behavioral data (**Figs. 3D** and **4D**), motivating the need for a more flexible and biologically plausible framework.

The DDM provides a compelling theoretical framework for modeling decision dynamics. It allows for flexible manipulation of core parameters—including drift rate, decision boundary, and integration properties—resulting in distinct operational regimes. A constant bound model (**Fig. 2**, top row), which assumes a fixed decision threshold over time, has proven effective across many perceptual decision-making studies (Voskuilen et al., 2016). However, under conditions of low evidence (i.e., zero drift rate), this model predicts unreasonably long RTs unaccompanied by improvements in accuracy—a situation we refer to as the "zero-drift-rate catastrophe."

To circumvent this limitation, we explored models incorporating time-dependent (collapsing) boundaries (**Fig. 2,** bottom row). These models reduce the decision threshold over time, accelerating responses in low-drift scenarios and yielding more symmetrical RT distributions—features more consistent with the observed behavioral data.

Another extension involved introducing leaky integration, where the accumulated evidence is affected by a decay (or amplification) factor. With a positive leak (unstable integration), accumulation becomes more rapid and boundary-crossing more likely, again minimizing long RTs (**Fig. 2**, middle row). However, fitting the current data using a constant-

bound DDM required large positive leak values, suggesting that drift rate played a diminished role relative to the gain term—raising questions about the model's plausibility.

When leak was added to the collapsing bound model, the best-fitting leak parameter was negative, indicating true leaky integration. This variant slightly improved model fit over the collapsing bound alone, though both models yielded excellent correspondence with behavioral data. Importantly, the collapsing bound model did not require extreme parameter values, making it a more parsimonious and biologically plausible account of TI decision-making.

Examining the empirical relationship between accuracy and RT (**Figs. 4** and **5**) revealed a weak but significant positive correlation during the early training phase. Within the DDM framework, this relationship reflects the interplay of drift rate and boundary height. Increases in drift rate reduce RT and improve accuracy, whereas increases in decision boundary prolong RT but also improve accuracy. The observed pattern suggests that, during learning, drift rate may have increased alongside modest adjustments in boundary height, enabling more accurate decisions at the cost of slightly delayed responses.

This observation led us to propose a hybrid "variable collapsing bound" model (**Fig. 9**). In this framework, the decision boundary gradually increases during early training (Fig. 9A), producing a positive RT–accuracy correlation (**Fig. 9B**). Once learning stabilizes, the boundary collapses over time within each trial but remains fixed across trials (**Fig. 9C**), resulting in the more familiar negative RT–accuracy relationship seen in symbolic distance-sorted data (**Fig. 9D**; cf. **Fig. 5**).

Together, these findings provide new insights into the cognitive and computational

processes underlying TI learning in nonhuman primates. By extending drift-diffusion modeling to a domain driven by internally inferred structure rather than direct sensory evidence, we highlight the adaptability of decision mechanisms to task demands. The success of the variable collapsing bound model suggests that decision policies are not fixed but evolve dynamically with learning, enabling fast yet accurate behavior even in abstract cognitive tasks.

This interpretation aligns with prior work suggesting that the utility of collapsing boundaries may depend on species and task context. While human data generally favors fixed-boundary models—with studies showing negligible practical advantages for collapsing variants (Voskuilen et al., 2016)—evidence from nonhuman primates is more mixed. Some paradigms support collapsing boundaries, indicating that decision strategies may be more flexible in nonhuman species or under specific task demands. Hawkins et al. (2015) further emphasize that urgency signals and boundary dynamics can be modulated by contextual factors such as time pressure and task structure. In this light, our findings contribute to a growing body of evidence suggesting that decision boundaries are not static but can adapt to optimize performance in complex, internally guided tasks.

Conclusion

The drift diffusion model (DDM) is traditionally employed in paradigms where accuracy and reaction time exhibit an inverse relationship. However, the flexibility of this modeling framework enables it to accommodate less conventional patterns of behavior, thereby offering insight into the interaction of latent cognitive variables across diverse decision-making contexts. In the

present study, we demonstrate that a generalized DDM can be effectively applied to a reasoning task—transitive inference—where reaction times remain relatively constant despite significant learning-related improvements in accuracy. Our results suggest that, during the acquisition phase, the starting point of an exponentially collapsing decision boundary can be gradually modulated to reflect the evolving rate of evidence accumulation. This dynamic adjustment allows for an increased accuracy with minimal cost in response time, effectively enhancing decision efficiency. Furthermore, task performance was modulated by internal variables, such as symbolic distance, underscoring the role of internally computed metrics in guiding behavior. The collapsing boundary model provides a mechanistic explanation for these observations, minimizing the impact on reaction time as accuracy improves. In contrast, under conditions of slow evidence accumulation, this framework reduces the total evidence required by collapsing the decision threshold over time, acknowledging the diminishing returns of prolonged integration. Collectively, these findings highlight the utility of variable-bound DDMs in capturing the dynamics of decision-making under conditions where evidence is internally generated and temporally constrained.

Data availability

The datasets generated and/or analyzed during the current study are available from the corresponding author on request.

Competing interests

The authors declare they have no competing interests.

Acknowledgements

This study was supported by NIH R01-MH111703 (VPF and HST).

Author Contributions

The study was designed by FM, GJ, YA, VPF, and HST. Data were collected by FM and YA. Data were analyzed by FM, GJ, MS, JDM and VPF. The manuscript was written by FM, GJ, MS, and VPF.

Supplemental Figures

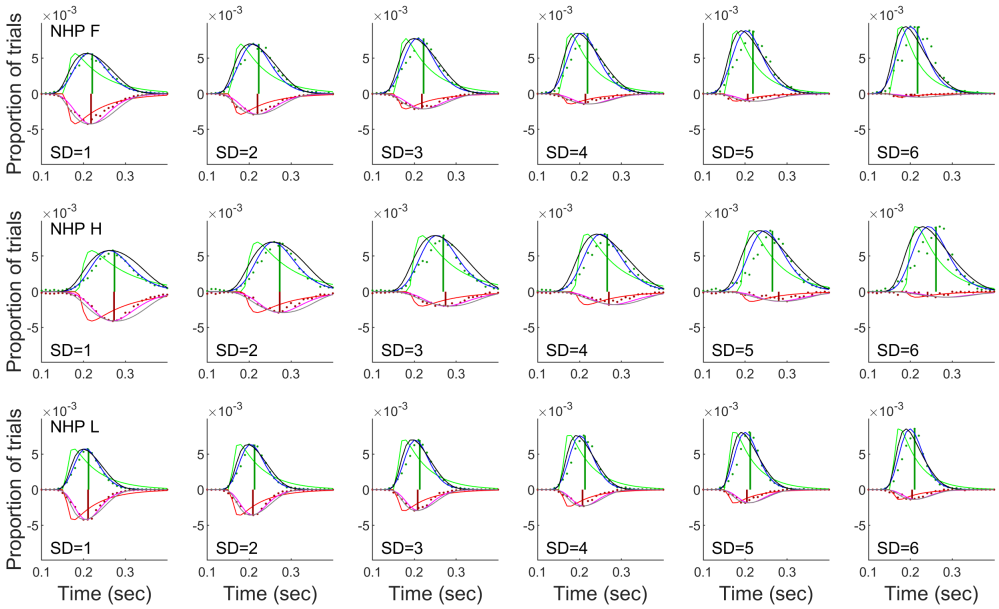

Figure S1. Drift Diffusion model fit for symbolic distance. Data from all sessions for all NHP segregated by trial block. Green dots are reaction time distributions for correct trials, red dots for error trials. Vertical lines indicate the median RT for each distribution. Green and red curves are PyDDM fits for the constant bound model. Blue and magenta curves are fits for the constant bound with leak model. Black and grey curves are fits for the collapsing bound model.

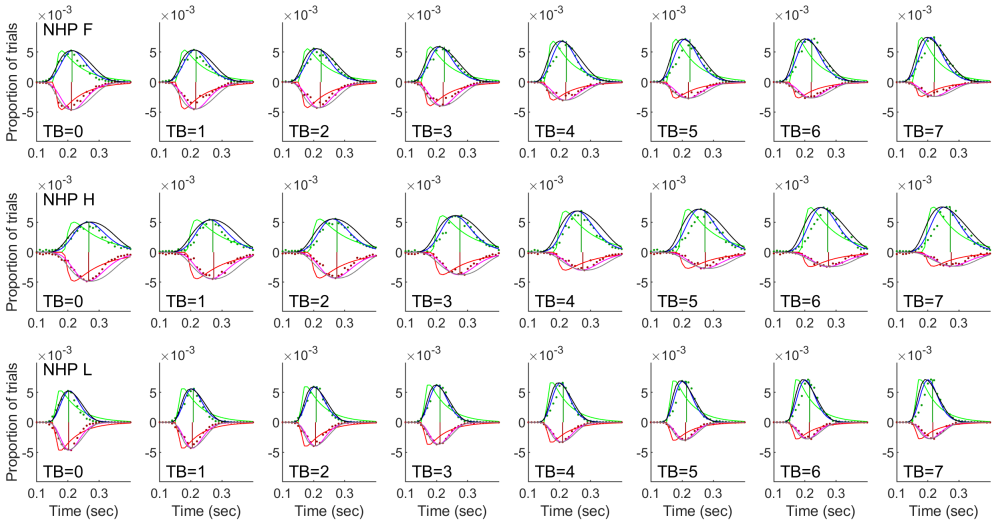

Figure S2. Drift Diffusion model fit for trial block. Data from all sessions for all NHP segregated by trial block. Green dots are reaction time distributions for correct trials, red dots for error trials. Vertical lines indicate the median RT for each distribution. Green and red curves are PyDDM fits for the constant bound model. Blue and magenta curves are fits for the constant bound with leak model. Black and grey curves are fits for the collapsing bound model.

References

1. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol Rev. 2006 Oct;113(4):700-65. doi: 10.1037/0033-295X.113.4.700. PMID: 17014301.
2. Bogacz R, Wagenmakers E-J, Forstmann BU, Nieuwenhuis S (2008) The neural basis of the speed-accuracy trade-off. Trends in Neurosciences. 33(1):10-16
3. Bonnefond M, Castelain T, Cheylus A, Van der Henst J-B (2014). Reasoning from transitive premises: An EEG study. Brain and Cognition 90, 100-108. https://doi.org/10.1016/j.bandc.2014.06.010
4. Churchland AK, Kiani R, Shadlen MN. Decision-making with multiple alternatives. Nat Neurosci.

2008 Jun;11(6):693-702. doi: 10.1038/nn.2123. Epub 2008 May 18. Erratum in: Nat Neurosci. 2008 Jul;11(7):851. PMID: 18488024; PMCID: PMC2453226.

5. Cohen JY, Crowder EA, Heitz RP, Subraveti CR, Thompson KG, Woodman GF, Schall JD. Cooperation and competition among frontal eye field neurons during visual target selection. J Neurosci. 2010 Mar 3;30(9):3227-38. doi: 10.1523/JNEUROSCI.4600-09.2010. PMID: 20203182; PMCID: PMC2844339.

6. Ding L, Gold JI. Caudate encodes multiple computations for perceptual decisions. J Neurosci. 2010 Nov 24;30(47):15747-59. doi: 10.1523/JNEUROSCI.2894-10.2010. PMID: 21106814; PMCID: PMC3005761.

7. Hawkins, G. E., Forstmann, B. U., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, 35(6), 2476-2484. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

8. Jensen G, Muñoz F, Alkan Y, Ferrera VP, Terrace HS. Implicit value updating explains transitive inference performance: the betasort model. PLOS Computational Biology, 11, e1004523. (2015).

9. Judge SJ, Richmond BJ, Chu FC. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Research*, *20*, 535-538. (1980).

10. Jun EJ, Bautista AR, Nunez MD, Allen DC, Tak JH, Alvarez E, Basso MA. Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. Nat Neurosci. 2021 Aug;24(8):1121-1131. doi: 10.1038/s41593-021-00878-6. Epub 2021 Jun 28. PMID: 34183869; PMCID: PMC8338902.

11. Kiani R, Cueva CJ, Reppas JB, Newsome WT. Dynamics of neural population responses in

prefrontal cortex indicate changes of mind on single trials. Curr Biol. 2014 Jul 7;24(13):1542-7. doi: 10.1016/j.cub.2014.05.049. Epub 2014 Jun 19. PMID: 24954050; PMCID: PMC4191655.

12. Kim JN, Shadlen MN. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci. 1999 Feb;2(2):176-85. doi: 10.1038/5739. PMID: 10195203.
13. Munoz, F., Jensen, G., Kennedy, B.C. *et al.* Learned Representation of Implied Serial Order in Posterior Parietal Cortex. *Sci Rep* **10**, 9386 (2020). https://doi.org/10.1038/s41598-020-65838-9
14. Murray, J., Bernacchia, A., Freedman, D., Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, Wang X-J. A hierarchy of intrinsic timescales across primate cortex. Nat Neurosci 17, 1661–1663 (2014). https://doi.org/10.1038/nn.3862
15. Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*(2), 59–108. https://doi.org/10.1037/0033-295X.85.2.59
16. Ratcliff R, Cherian A, Segraves M. A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. J Neurophysiol. 2003 Sep;90(3):1392-407. doi: 10.1152/jn.01049.2002. Epub 2003 May 21. PMID: 12761282.
17. Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci. 2002 Nov 1;22(21):9475-89. doi: 10.1523/JNEUROSCI.22-21-09475.2002. PMID: 12417672; PMCID: PMC6758024.
18. Shankar S, Massoglia DP, Zhu D, Costello MG, Stanford TR, Salinas E. Tracking the temporal evolution of a perceptual judgment using a compelled-response task. J Neurosci. 2011 Jun 8;31(23):8406-21. doi: 10.1523/JNEUROSCI.1419-11.2011. PMID: 21653845; PMCID: PMC3134312.
19. Maxwell Shinn, Norman H Lam, John D Murray (2020) A flexible framework for simulating and fitting generalized drift-diffusion models eLife 9:e56938

20. Middlebrooks PG, Schall JD. Response inhibition during perceptual decision making in humans and macaques. Atten Percept Psychophys. 2014 Feb;76(2):353-66. doi: 10.3758/s13414-013-0599-6. PMID: 24306985; PMCID: PMC4141461.

21. Prado J, Mutreja R, Booth JR. Fractionating the neural substrates of transitive reasoning: task-dependent contributions of spatial and verbal representations. Cereb Cortex. 2013 Mar;23(3):499-507. doi: 10.1093/cercor/bhr389. Epub 2012 Jan 23. PMID: 22275478; PMCID: PMC3563338.

22. Riley, C. A. (1976). The representation of comparative relations and the transitive inference task. *Journal of Experimental Child Psychology, 22*(1), 1–22. https://doi.org/10.1016/0022-0965(76)90085-0

23. Soltani A, Murray JD, Seo H, Lee D. Timescales of Cognition in the Brain. Curr Opin Behav Sci. 2021 Oct;41:30-37. doi: 10.1016/j.cobeha.2021.03.003. Epub 2021 Mar 31. PMID: 34026949; PMCID: PMC8136243.

24. Stanford, T., Shankar, S., Massoglia, D. *et al.* Perceptual decision making in less than 30 milliseconds. *Nat Neurosci* **13**, 379–385 (2010). https://doi.org/10.1038/nn.2485

25. Wald A (1945) Sequential tests of statistical hypotheses. Ann Math Statist. 16(2):117-186. Doi: 10.l1214/aoms/1177731118

26. Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. The Annals of Mathematical Statistics. 19(3):326-339.

27. Voskuilen C, Ratcliff R, Smith PL. Comparing fixed and collapsing boundary versions of the diffusion model. J Math Psychol. 2016 Aug;73:59-79. doi: 10.1016/j.jmp.2016.04.008. Epub 2016 May 24. PMID: 28579640; PMCID: PMC5450920.

28. Yoo SBM, Hayden BY, Pearson JM. Continuous decisions. Philos Trans R Soc Lond B Biol Sci. 2021 Mar;376(1819):20190664. doi: 10.1098/rstb.2019.0664. Epub 2021 Jan 11. PMID: 33423634;

PMCID: PMC7815426.