A novel rat spatial transitive inference paradigm for investigating memory schema and deliberation

Authors: Blake S. Porter, Catherine Shi, Evgeniia Kozlova, and Shantanu P. Jadhav

Affiliation: Neuroscience Program, Department of Psychology, and Volen National Center for Complex Systems, Brandeis University, Waltham, MA, 02453, USA

Correspondence: blakeporterneuro@gmail.com

Conflicts of Interest: Authors report no conflicts of interest

Funding Sources: This work was supported by a Smith Foundation Odyssey Award and NIH grants R01MH112661 and R01MH120228 to S.P.J. B.S.P was supported by the NIH NINDS T32 (NS0072-92).

Abstract

Inferential reasoning is a vital cognitive ability that enables animals to navigate novel situations by leveraging existing relational knowledge of memory schema. Transitive inference (TI), a classic paradigm for studying this process, tests the ability of subjects to infer relationships within a value hierarchy (e.g., A > B > C > D > E) after being trained only on adjacent pairs (e.g., A-B, B-C). While the ability to perform transitive inference is well-conserved across species, the neural mechanisms supporting it, particularly the role of the hippocampal-prefrontal circuit, remain underexplored due to limitations in existing rodent TI paradigms. These paradigms often rely on manual trial setups with pot-digging choice behavior using odors that makes it difficult to pinpoint the moment of inferential deliberation. To address these challenges, we developed a novel automated spatial TI task for rats using a radial maze, which leverages their natural foraging behaviors and includes a dedicated deliberation zone. We found that most rats successfully learned the premise pairs and successfully inferred, with performance comparable to studies using traditional odor-based tasks. We also investigated vicarious trial and error (VTE), a behavioral correlate of deliberation during choice trajectories. We found that VTE behavior was elevated during the initial stages of learning with access to new maze arms, when the first novel premise pair was introduced, and for incorrect trials, corresponding to hypothesized association of VTE events with uncertainty. Further, rats also had elevated VTE behavior on inference testing day for the inference test pair (B-D) compared to the other inner premise pairs (B-C & C-D) on their first testing day, suggesting the need to deliberate during inferential reasoning. Thus, VTEs may reflect deliberation in addition to uncertainty. Our findings demonstrate the feasibility of a spatial TI task that is well-suited for neurophysiological investigation and provides new insights into the behavioral correlates of inferential reasoning and deliberation.

Introduction

In an ever-changing world, animals will encounter novel situations where they must rely on existing knowledge to make the best decision. Inferential reasoning is the cognitive ability to draw upon existing knowledge that shares common features with a novel scenario to make an informed decision. Memory schemas, or relational networks of abstract and generalized information, likely aid inferential reasoning by enabling the retrieval of information that matches elements of a novel situation (Varga et al., 2022). Leveraging the relational network of schemas enables animals to infer potential action plans and their outcomes without having to explicitly experience them first (Tenenbaum et al., 2011). Thus, inferential reasoning not only saves time and energy but it can reduce risk by avoiding inferred harm.

Transitive inference (TI) is a hallmark inference paradigm that tests a subject's ability to infer the relationships of elements in a value hierarchy (e.g., A>B>C>D>E; Vasconcelos, 2008). Subjects are only trained on adjacent items of the set (e.g., A-B, B-C, etc.; also known as premise pairs) while testing consists of novel combinations (e.g., B-D). Thus, to successfully infer, subjects must form a memory schema of the complete value hierarchy from the premise pairs. Transitive inference reasoning appears to be a highly conserved cognitive function as it has been demonstrated in insects (Tibbetts et al., 2019), fish (Grosenick et al., 2007), rodents (Davis, 1992; Dusek & Eichenbaum, 1997), multiple bird species (Lazareva et al., 2004; Steirn et al., 1995; von Fersen et al., 1991), non-human primates (NHP; Brunamonti et al., 2016; Gillan, 1981; Mcgonigle & Chalmers, 1977; for Vasconcelos, 2008), and humans (Bryant & Trabasso, 1971; Preston et al., 2004). However, studying the neural mechanisms underlying inferential reasoning and memory schemas can be challenging in rodent models due to the commonly used TI paradigms having several disadvantages for investigating such questions.

Transitivity is an important property of relationships between stimuli, actions, and information more generally. Utilizing inferential reasoning on transitivity allows animals to understand the world around them without having to experience every possible relationship (Eichenbaum, 1999, 2000; Tenenbaum et al., 2011). An ethological example of the benefits of transitive inference is social hierarchies and rank (Becker et al., 2021). By observing peers fighting and considering their own rank, territorial beta fish utilize transitive inference to determine their social rank and abilities to fight other peers (Grosenick et al., 2007). Being able to infer their rank can help them avoid obvious defeats and take on easier rivals. However, we know little about the underlying physiological mechanisms that support transitive inference. Recent computational modeling work using artificial neural networks has demonstrated that transitive inference can occur in many model architectures (Frank et al., 2003; Kay et al., 2024;

Kumaran & McClelland, 2012; Miconi & Kay, 2025; Wu & B Levy, 2001; Wu & Levy, 1998), including hippocampal models (Frank et al., 2003; Kumaran & McClelland, 2012; Wu & B Levy, 2001; Wu & Levy, 1998). Human fMRI studies (Acuna et al., 2002; Heckers et al., 2004; Schlichting & Preston, 2016; Zalesak & Heckers, 2009; Zeithamova et al., 2012) and human (Koscik & Tranel, 2012; Wing et al., 2021) and animal lesion studies (DeVito, Kanter, et al., 2010; DeVito, Lykken, et al., 2010; Dusek & Eichenbaum, 1997) have demonstrated the necessity of the hippocampus and prefrontal cortices in transitive inference. Recent electrophysiological studies with NHP have provided some key insights into how prefrontal cortical circuits aid in inferential reasoning (Brunamonti et al., 2016; Ramawat et al., 2022, 2023). However, rodent electrophysiological data on these circuits during training and inference testing is sparse. One reason for that may be that existing TI paradigms for rodents are not amenable to the constraints of electrophysiological recordings. To address these issues, we have developed a novel spatial transitive inference task for rodents.

To date, most rodent transitive inference studies have relied on paradigms requiring rodents to discriminate between pots filled with scented sand where each odor is an item in the value hierarchy (e.g., A is paprika, B is coffee, etc.; (André et al., 2012; Davis, 1992; DeVito, Kanter, et al., 2010; DeVito, Lykken, et al., 2010; Dusek & Eichenbaum, 1997; Van der Jeugd et al., 2009; Van Elzakker et al., 2003). This setup makes it challenging to pinpoint when exactly rodents are deliberating and potentially utilizing inferential reasoning. Furthermore, many of the previous TI paradigms required the experimenter to manually set up each trial. We sought to develop an automated spatial version of the TI task to leverage rats' natural abilities to forage for food and to be amenable to physiological investigation. We also sought a paradigm that has a delineated area of the maze where rats can deliberate choices before execution, facilitating the elucidation of the behavioral and neural correlates of inferential reasoning.

A promising approach to quantify deliberative behaviors is vicarious trial and error (VTE; for review: Redish, 2016). Vicarious trial and error events were first described by Gentry and Muenzinger (Muenzinger, 1956) and popularized by Tolman (Tolman, 1939, 1948) as a behavioral readout of deliberation where rats vicariously test specific action plans and their outcomes before making a choice. This theory has been bolstered by evidence from neural recordings showing hippocampal-prefrontal oscillatory interactions (Amemiya & Redish, 2018; Rosenblum et al., 2025; Stout et al., 2022) and hippocampal theta sequences sweeping ahead of animals to multiple potential routes during VTEs (Johnson & Redish, 2007; Papale et al., 2016). Yet the exact role of VTEs and their prevalence remains poorly understood. VTEs can be most prevalent during learning while rats deliberate during periods of uncertainty. As a task becomes habitual, VTEs decline. However, VTEs may also persist in especially

challenging tasks even as the rats become proficient, possibly reflecting the ongoing need for deliberation. Recent findings have further challenged the deliberative nature of VTEs. Miles et. al. (2024) used a flexible, spatial decision making task to show that VTEs may also more generally be tied to periods of uncertainty, regardless of how proficient the animal is at the task. Thus, development of a spatial transitive inference task with a deliberation zone allows investigation of the role of VTEs in uncertainty and deliberation during both initial learning of premise pairs and during later tests of inferential reasoning, allowing valuable insights into the role VTEs play in decision making, and providing a foundation for neurophysiological investigation of the mechanisms underlying deliberation and inferential reasoning.

Methods

Subjects

All experimental methods were approved by the Brandeis University Institutional Care and Use Committee (#21001 and #24001-A) and conducted under the guidelines of the US National Institutes of Health. A total of 14 adult Long-Evans rats (1 female, 13 males; weighing between 350 – 536 grams) were used. Five of the rats were from a TH-Cre line (Witten et al., 2011) and used in a previous experiment (Ding et al., 2025) before the transitive inference task. All other rats were obtained from Charles River Laboratories or bred in-house.

Initially, animals were dual housed unless implanted (implanted animals were single housed after surgery) and given ad libitum access to food and water on a 12-hour/12-hour light/dark cycle (7am-7pm light). Before training began, rats were food-restricted to no less than 85% of their free feed weight. The rats were habituated to the experimenters and evaporated milk (Nestlé S.A.) before being trained to shuttle back and forth on a linear track. The linear track was in the same room as the main experimental radial maze and constructed of the same maze parts. While shuttling, pneumatic barriers were randomly triggered on the radial maze to habituate the rats to the sound. When rats began to shuttle on the linear track, a lick requirement to dispense the reward was introduced. Rats were trained to lick three times before a reward was dispensed. This lick requirement ensured that 1) there was a deliberate behavioral readout of their choice beyond just traveling to a reward well, 2) the reward was not predispensed so rats could not navigate by smell or the sound of the pump to the correct choice arm, and 3) it avoided rats accidentally dispensing a reward (e.g., by stepping on the lick sensor). Once readily shuttling (~ 5 rewards/min), non-implanted animals and the already implanted TH-Cre animals moved onto premise pair training. The other implanted rat underwent surgery and received a custom 64 tetrode microdrive targeting mPFC and HPC bilaterally with a similar surgical protocol to Ding et al. (2025). Electrophysiological data are not presented in this paper. Once the animal was

recovered, they were retested on the linear track to habituate to the weight of the implant. Afterwards, their training commenced in the same steps as non-implanted and TH-Cre animals.

Premise pair training

An 8-arm radial maze, with automated rule implementation, reward delivery, and barrier control, was used for the experiment (Porter et al., 2025). Each arm had unique cues such as hot glue dots, black foam, or vinyl stripes. Colorful shapes were adhered to the walls of the experiment room. Arms were separated by walls so that rats could not jump from one arm to another. Arm access was controlled by 12” pneumatic barriers situated between the walls and between the maze center and radial arms. At the end of each arm was a reward well that would dispense an evaporated milk reward (~ 0.1mL) after rats licked three times. One arm was designated as the Home arm and was the same arm across all rats. For each rat, five arms were randomly assigned to the five items of the value hierarchy (i.e., A,B,C,D,E). The two arms 45° from Home were not used. Any arm assignments that were assessed to have a simple egocentric strategy that could be used to successfully solve the task (e.g., always go to the left- or right-most arm) were not used and new assignments were generated until satisfactory. A rat’s arm assignments stayed the same throughout all of training and testing.

Our transitive inference trial structure was unique from previous versions of rodent TI. Rats always began at Home where they received a reward. Upon licking at Home, the barriers were set for the next trial. Rats would enter the center of the radial maze and, given the two open arms, choose the correct one based on the latent value hierarchy (A>B>C>D>E). If rats chose correctly (e.g., arm A for A-B), they would be rewarded at the reward well at the end of the correct arm, and Home would be blocked off. Rats would then need to go to the second arm (e.g., B) for a second reward, the only choice available, which would then lower the Home barrier and raise the first arm barrier. If the animal chose incorrectly at the beginning of the premise-pair trial (e.g., chose arm B for A-B), rats would receive no reward at the incorrect arm and the barrier to the correct arm (e.g., A) would be raised so it could not be visited. Thus, for incorrect trials, there was no reward and rats had to return Home to initiate the next trial. Thus, for correct trials, rats receive rewards at the end of both arms comprising the premise pair, with the critical choice being the first arm of the pair.

Once readily shuttling on the linear track, all the rats began training on the transitive inference task’s premise pairs. Rats received no pre-exposure or pre-training to the 8-arm maze. Rats were trained once per day for five to seven days per week. Training sessions lasted at least fifteen minutes but were stopped at one hour. After fifteen minutes, the session was ended when the rat did not attempt a trial for two

minutes. All rats developed a habit of sitting and waiting at the Home arm to end the session.

The premise pair training (A>B>C>D>E) had three phases (**Figure 1A**). Rats progressed to the next phase when they met the criterion of performing 75% or more correct for all present premise pairs for two days in a row. This criterion was chosen to be similar to previous rodent TI tasks (DeVito, Kanter, et al., 2010; Dusek & Eichenbaum, 1997; Silverman et al., 2015; Van der Jeugd et al., 2009; Van Elzakker et al., 2003). For the first phase of training, rats were trained on the premise pairs A-B and B-C. The second phase added C-D, and the third phase added D-E, with all pairs included in each phase (e.g., the second phase had A-B, B-C, C-D trials, and the third phase had all four premise pair trials). By default, premise pairs were randomly intermixed. However, to avoid random long streaks and to facilitate learning, if one premise pair had three fewer rewards than any of the other premise pairs, that lower-reward premise pair was presented until its reward count caught up to the others. Thus, if rats were struggling to get certain premise pairs correct, those pairs would be presented more often. On occasion, if a rat was severely struggling with certain premise pairs, a training day would begin with fixed blocks (5 – 10 trials long) of only those premise pairs before the intermixing of all premise pairs. These block sessions were counted in the total training time but were never used for determining the criterion. Only fully intermixed sessions were used to determine if a rat was ready for the next training phase or testing. Once a rat reached the final criterion of performing 75% or better on all four premise pairs for two days in a row, inferencing testing commenced the following day.

Transitive inference test day

Inference testing was conducted for just one day. While rats may be capable of learning the new pairs within a single day, this is unlikely given the training times of premise pairs (**Figure 2**). But repeated exposure to the test pairs over days would likely result in them learning them through trial and error rather than inferring their choices. On test day, rats were presented with the four premise pairs (A-B, B-C, C-D, and D-E) as well as an inference test pair (B-D) and a novelty control pair (A-E). The reinforcement schedule was the same as training for all pairs, i.e., correct B-D and A-E choices were rewarded along with all premise pairs (Vasconcelos, 2008). Pair presentation was controlled by randomly generating blocks of six trials where each item-pair was presented once in the block. This ensured a roughly equal number of trials across pairs. The first two blocks were only the four premise pairs; all subsequent blocks were six pairs.

Behavioral analysis

All data was collected with SpikeGadgets' environmental control unit (ECU) and their software. An overhead camera (Mako G-158C, Allied Vision Inc.) synced to SpikeGadgets recorded video at 30 frames per second. Rats' position was tracked with DeepLabCut (DLC; Nath et al., 2019). For non-implanted animals, their nose was tracked for position. For implanted animals, the light-emitting diodes on their implant were tracked for position. To ensure accurate tracking, any DLC tracked position with a likelihood less than 0.98 was removed and a linear interpolation was applied. A boxcar filter of 5 frames (0.17 seconds) was then applied to remove jitter that could impact the Idphi metric.

Training day selection

Rats had a highly variable number of training days. In order to compare performance (**Figure 2B**) and VTEs (**Figure 4D & E**) across training, we divided the premise pair training times into thirds. The first day was selected as the first day a premise pair was presented. The criteria day is the day before rats move onto the next training phase (for A-B, B-C, and C-D) or the day before test day (for D-E). The midpoint day was determined by the day between the first and the criterion day, rounded up to a day if needed. Some rats did not have a midpoint day due to meeting the criterion in just two days. In these rare cases, the midpoint day was represented as a "not a number" (NaN). However, if a day had fewer than five trials, it was not used, and the day search was iterated until a day with more than five trials was found. This was primarily done to avoid the earliest training days for some rats who were slow to habituate to the maze and attempted very few trials.

VTE and zIdPhi

In order to quantify VTEs, we used the same integrated absolute change in angular velocity of the head (Idphi) metric developed by Papale et al. (2012) on choice trajectories from Home to the first arm of choice. For every recording session, a region of interest was manually drawn for the center area of the maze as well as for every arm. The center zone comprised the area inside all the arm barriers. The arm zones were drawn at the barrier to the length of the arm. A choice trajectory was determined from the frame the rat (nose or LED) entered the center zone to the frame they entered an arm that was not Home. If rats entered the center and then re-entered the Home arm, the trial was not analyzed. If rats entered a choice arm but did not lick at the reward well and went to the other choice arm, only the initial trajectory from Home was analyzed. Trajectories longer than two seconds were also excluded so that behaviors such as grooming or excessive exploration were not analyzed.

Choice arms were all at different angles from the Home arm. Two arms were at ±90°, two arms were at ±135°, and one arm was 180° from Home. To account for this, Idphi values were z-scored within an arm. This ensured the relatively low baseline Idphi values for the 180° arm / relatively high values for the 90° arms did not bias the mean or standard deviation for other arms. Z-scoring was done independently for each rat using all their trajectories from both training and test to determine zIdPhi values for analysis.

Statistical testing

Parametric tests were used unless otherwise stated. Error bars show the standard error of the mean. Significance was tested with an alpha of 0.05. Bootstrapping for testing correct versus incorrect trials was run 100 times using MATLAB's datasample function with replacement.

Results

We used a five-item transitive inference set (A>B>C>D>E; **Figure 1A**) where arms of an eight-arm radial maze represented items of the set (**Figure 1B**). Arm availability was automated based on trial type and controlled by pneumatic barriers (Porter et al., 2025; **Figure 1D**). Specifically, for a given trial type, only the arms associated with the items of the pair were open (e.g., A-B, only the arms for A and B were open). All other arms had an opaque barrier blocking them off. One arm served as Home where rats needed to return to initiate the next trial. For each rat, item identity was randomly assigned to arms. However, configurations were rejected if they were solvable by an obvious egocentric strategy (e.g., "always go left", which can be seen in **Figure 1B** for illustrative purposes only). A rat's configuration was constant for the entire experiment.

Our transitive inference paradigm was altered from more established paradigms to better investigate memory schema. Primarily, we 1) altered how premise pairs are introduced during training (**Figure 1A**) and 2) how the element of lesser value in a pair is treated (**Figure 1C**). During training, previous studies tend to introduce all premise pairs at once with blocks of trials for each pair. Block sizes are reduced over the course of training until the paradigm is switched to an intermixed, pseudorandom presentation of all premise pairs. We chose to iteratively add premise pairs once rats were performing well on the existing pairs, and by default, presenting pairs intermixed (**Figure 1A**). This design allows us to ask questions around schema accommodation and assimilation as new pairs are added and the existing task schema needs to be altered. Training began with intermixed presentations of A-B and B-C trial types. We used a criterion of 75% correct for all present pairs for two days in a row to be considered proficient. Once the criterion was reached, the next premise pair was added

to the existing premise pairs until all four premise pairs were present (i.e., A-B, B-C, C-D, and D-E).

A second alteration was the treatment of the second item in the set. For correct trials, rats visited the second item in the premise pairs for “free” (e.g., a correct A-B trial consisted of rats getting a reward at A, then another reward at B; **Figure 1C & D**). We chose to do this in order to avoid the end item (i.e., E) never being rewarded and, during periods of high performance, never being visited. Incorrect trials proceeded as usual in that if rats made the wrong choice (e.g., B for an A-B trial), rats received no reward and needed to return Home to start a new trial (arm A would be blocked off once B was chosen).

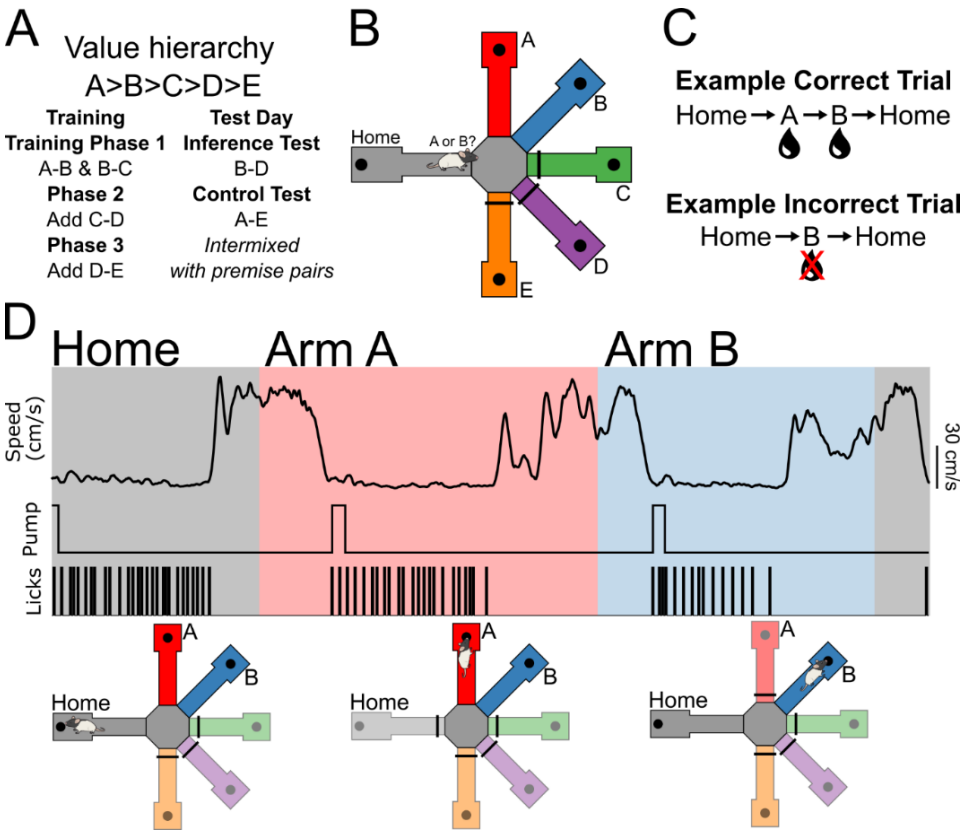

Figure 1: Spatial transitive inference paradigm. A) Paradigm overview of training and testing schedules. **B)** Maze schematic depicting an example A-B trial where rats exit home and must decide on the correct arm to visit. Other arms/items are blocked off by barriers (black bars). Black circles represent reward wells. **C)** Example correct and incorrect trial structure for A-B trial. **D)** Top: Timeline of one example A-B trial showing speed, reward Pump, and licks at reward wells. Animals start at Home to initiate a new trial. Rats then travel to the center and decide on which arm to choose. In this trial, the rat is correct and consumes a reward at arm A before traveling to arm B for a second reward and then returning Home. **Bottom:** schematic state of the barriers and arm availability during the trial. Greyed-out arms are blocked by barriers (black bars). *Note: Sequential arm order is shown for illustrative purposes. No rat had arms in sequential order.*

A total of 14 rats were trained on the transitive inference paradigm (**Figure 2A**). Of these 14, four failed to learn the premise pairs. Failure to learn was not operationalized, but was generally characterized by rats reaching more than 30 days of training with no consistent progress towards reaching the final criterion for all four premise pairs. These rats would often have stochastic day-to-day performance, going from performing greater than 75% for a pair on one day to below chance on the same pair the next day. All four of these rats did finish phase 1 of training (A-B, B-C) and thus understood the general task structure. However, adding the subsequent pairs seemingly proved too challenging. The data from the four rats who failed to learn are not analyzed in this paper.

Ten rats successfully learned all four premise pairs and performed at least 75% correct on each of the four pairs for two days in a row. Due to all rats having a different number of training days, we divided the days each premise pair into the first day it was presented, the criterion day (crit.) before the next phase (next training phase or test day), and a midpoint day (mid) between the two (**Figure 2B**). Considerable variability was seen across different rats on time to learn specific premise pairs, though no one pair was significantly different from the others ($F(3,36) = 1.13$, $p = 0.35$; **Figure 2C**). Most rats obtained 75% correct on A-B within the first day or two (mean 2.2 ± 1.6 days) as they essentially just needed to learn to go to A. B-C tended to have a longer time (3.2 ± 1.2 days) to achieve 75% for the first time as rats needed to overcome their avoidance of B from A-B in order to get it correct. Afterwards, rats picked up C-D (2.2 ± 1.6 days) and D-E (2.3 ± 1.4 days) relatively quickly. However, when a new premise pair was added, the performance on other premise pairs tended to temporarily fall, possibly reflecting the need to update the previous end item to now be a correct choice in the right context (e.g., C for B-C when C-D is added). There was no significant difference in the average total training days for each training phase ($F(2,27) = 2.1$, $p = 0.14$; **Figure 2D**). However, a great deal of variance can be seen across rats, especially for the final ABCDE phase. Some rats seemed to be able to quickly integrate D-E within a few days, while others needed upwards of 18 days to reach the criterion of 75% correct for all four pairs for 2 days in a row. On average, rats took 21.2 days (standard deviation, S.D. = 10.6) to complete all training, consisting of an average of 2,215 total trials (S.D. = 737). Despite the variability in training times, all rats who learned the task were performing around 90% correct on all four premise pairs on the final training day with no significant difference between pairs ($F(3,36) = 1.89$, $p = 0.15$; **Figure 2E**)

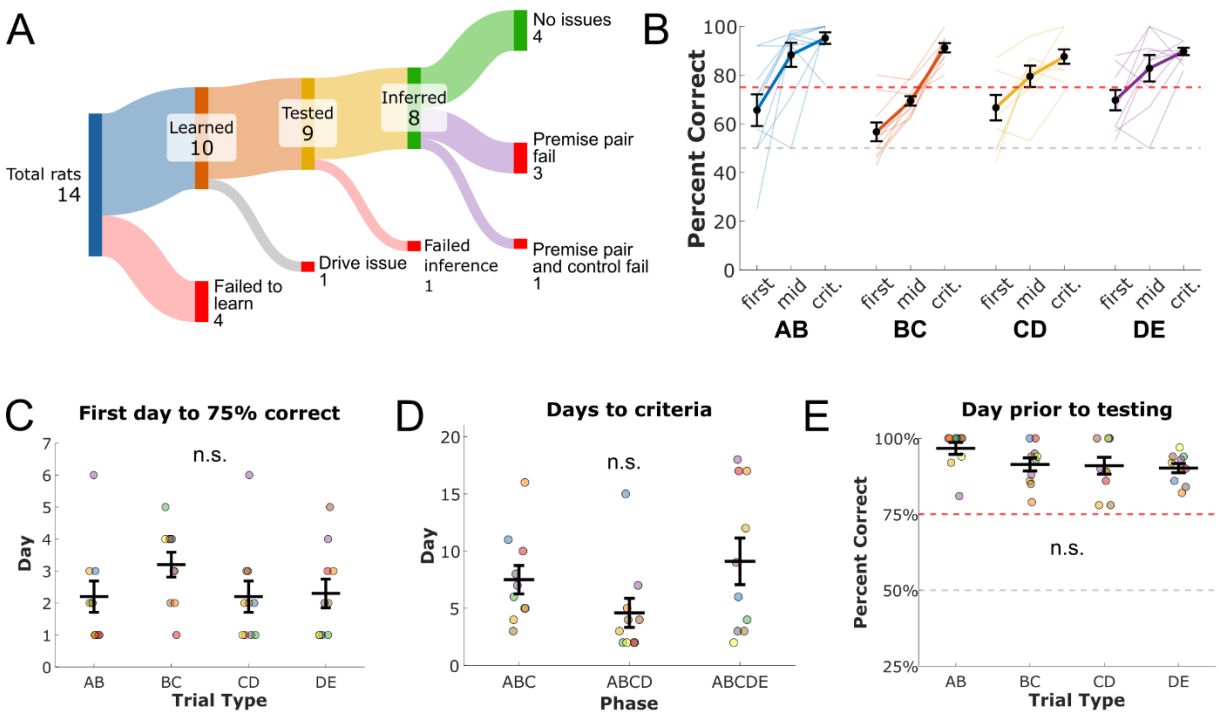

Figure 2: Transitive inference training. A) Sankey diagram of rat training and testing progression. **B)** Performance over training for premise pairs. The first training day, the final criteria day for a training phase, and a midpoint day were chosen. **C)** Number of days to first achieve one day of 75% correct for each trial type. **D)** Number of days to finish a training phase by reaching 75% correct for each trial type for two days in a row. **E)** Rats' performance on premise pairs on the training day before the test day. **C-D**, each color represents data from one rat. n.s., $p > 0.05$.

Once rats reached the criterion for all premise pairs, they were tested on their ability to infer a novel pair of items, B-D. We also tested rats' ability to perform on A-E as a novelty control. Overall, there was no significant difference in performance across trial types on test day ($F(5,48) = 0.83$, $p = 0.54$; **Figure 3A**). Eight out of the nine rats tested successfully inferred B-D, as shown by their average performance on B-D being around 85% correct. The rat who failed the inference test did worse than chance (~33% correct), possibly indicating a strong bias to arm D when faced with uncertainty. Curiously, of the eight animals that successfully inferred, four of them had premise pair performance below 75% and one animal performed A-E at 50%. The marked drop in premise pair performance, especially for A-B, was unexpected as the animals were performing very well on these trial types in the days leading up to testing (**Figure 2E**). It may be the case that the additional testing pairs disrupt the existing task schema and alter performance. Next, we investigated the time course of B-D performance on the first 10 trials (**Figure 3B**). Using a four trial window with a two trial overlap, we find rats on average perform ~ 75% initially and achieve near 100% performance for trials 7-10.

A hallmark of memory schemas is that new information can be scaffolded onto existing information so long as they share some common features. To investigate this,

we compared how well rats performed on the first day they encountered C-D and D-E during training and their performance on test day, the first time they saw B-D and A-E. We excluded A-B and B-C since their first day was the first day the rats were on the maze and likely no memory schema had formed yet. We found that rats performed significantly differently on the first day a trial type was presented ($F(3,28) = 7.1$, $p = 0.0011$; **Figure 3C**). Rats did significantly better on the test pairs (B-D & A-E) compared to the premise pairs (C-D & D-E) on the first day they are presented (all p's < 0.05). Overall, rats' high performance on the test pairs may indicate utilizing the memory schema of the value hierarchy as opposed to trial and error guessing when new arms are added.

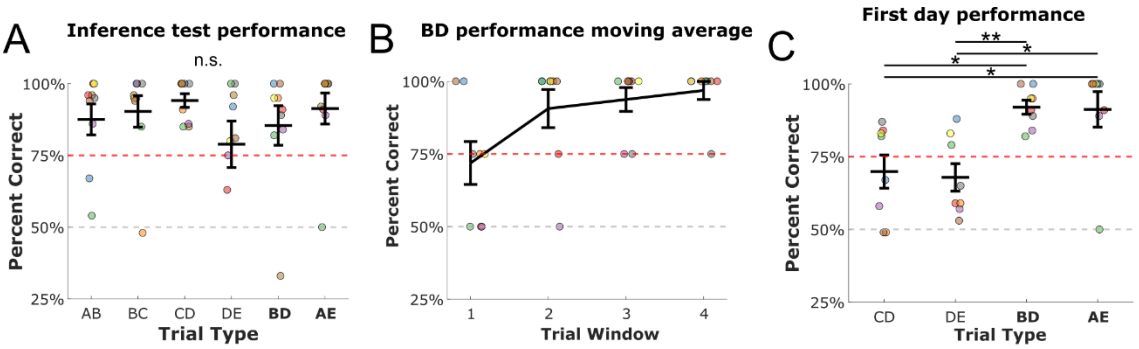

Figure 3: Inference testing. A) All rats' performance on testing day shows rats performed well on all trial types with no significant difference across trial types. However, one of the eight rats failed to infer B-D, and four rats had issues with other presented pairs. **B)** Performance of rats who successfully inferred on the first 10 B-D trials. A 4 trial window was used with a 2 trial overlap. **C)** Performance of rats who successfully inferred on the first day they were presented with a trial type. B-D and A-E have the same performance as seen in **A**. Rats do significantly better on test pairs than premise pairs when first introduced. Each color represents data from one rat. Error bars show SEM. n.s., not significant. *, $p < 0.05$; **; $p < 0.01$.

VTEs potentially provide a valuable behavioral readout of deliberation. Previous studies generally proposed that VTEs remain low before rats understand the layout of the maze (i.e., no schema), peak while rats are figuring out the rules (i.e., schema forming and being tested), and diminish as the task becomes habitual (i.e., reliable schema being used; Redish, 2016). However, sufficiently hard tasks may show relatively elevated levels of VTEs even after task proficiency is reached. Our task is unique among many existing paradigms since arms D and E are added during training, necessitating schema updating. Furthermore, our rats were naïve to the maze at the start of the experiment. We were interested in the progression of VTEs over training and how adding new premise pairs altered VTEs.

We measured VTEs using the zIdPhi metric of integrated absolute value of the angular head velocity (Papale et al., 2012). However, since all choice arms are in

different locations at different angles from the Home arm, trajectories were z-scored independently for each arm. By doing this, we avoid mixing the relatively low Idphi values from the trials to the arm 180 degrees from Home with relatively higher Idphi values for trials to the arms 90 degrees from Home. Any choice trajectory lasting longer than 2 seconds was removed to avoid including any trials where behaviors like grooming or excessive exploration occurred. After removing these trials, the mean choice time was 0.75 seconds (± 0.24 S.D.). Three example trajectories from the same session and the same trial type are shown in **Figure 4** (**A-C**). A ballistic, correct trajectory can be seen in **Figure 4A** while **Figure 4B & 4C** show VTE trials. We measured the zIdPhi values across training (same timepoints as **Figure 2B**), which peak early on (A-B and B-C day 1) and gradually diminish with proficiency. ZIdPhi values peak again when new arms and trial types are added. We tested whether zIdPhi values were different on the first day of each premise pair introduction. We found a significant difference in trial types for zIdPhi values on the first day of presentation ($F(3,31) = 4.71$, $p = 0.008$). Post-hoc multiple comparisons showed zIdPhi values were significantly higher on the first day of A-B and B-C compared to D-E ($p = 0.02$, $p = 0.02$, respectively). We also wanted to test whether VTEs increased for all trial types when a new arm was added. Thus, we focused on zIdPhi values for the first day of C-D and D-E (**Figure 4E**). On the first day of C-D, zIdPhi values are significantly higher than A-B and B-C zIdPhi values ($F(2,27) = 5.1$, $p = 0.013$). In contrast, on the first day of D-E, zIdPhi values are no different across the four trial types ($F(3,36) = 1.8$, $p = 0.16$). Taken together, the decrease in the zIdPhi of trials over training, especially by D-E's introduction, may indicate that by the time D-E is introduced, most rats have an established schema of the task that can readily assimilate new premise pairs, and D-E trials do not especially evoke VTEs.

An increased presence of VTEs has been reported to be associated with both correct and incorrect trials. The ambiguity may be due to how different tasks are performed and their difficulty levels. We tested zIdPhi values between correct and incorrect trials via bootstrapping by downsampling with replacement the correct trials down to match the number of incorrect trials. The first three days of training for all animals were not included, as rats were still learning the task in general and had high error rates. In our transitive inference task, we found that incorrect trials had significantly higher zIdPhi values than correct trials (100% of bootstrap permutation tests had p values < 0.05; mean p value < 0.0001; 95% confidence intervals for effect size: -0.24 to -0.20; **Figure 4F**). We also tested each rat individually to determine if their zIdPhi values were different across correct and incorrect trials. Eight out of the 10 rats had mean p-values less than 0.05. Thus, in our task, higher zIdPhi values are strongly associated with incorrect trials.

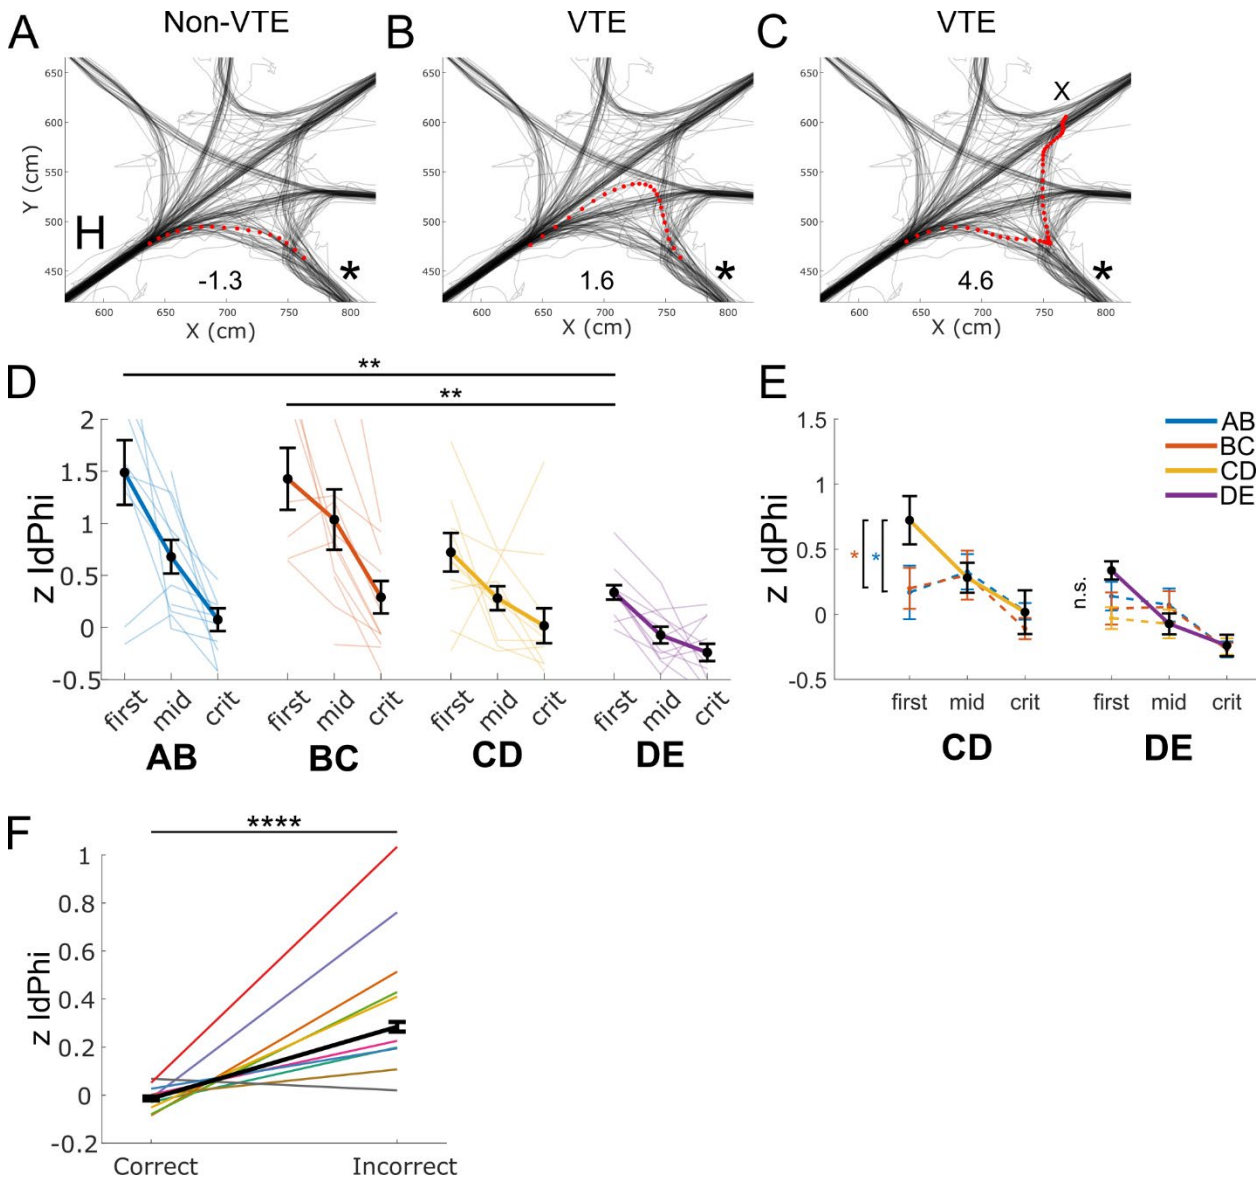

Figure 4: VTEs during premise pair training. **A-C)** Example choice trajectories from Home (H, bottom left) to the arm of choice. All examples are the same trial type where the correct arm is noted by the *. The inset value indicated zIdPhi. **A)** Smooth ballistic trajectory to the correct arm. **B)** Example VTE event where rat begins on path to wrong arm but reorients to correct arm. **C)** Example VTE event where rat stops at the correct arm to reorient and choose the incorrect arm denoted by the X. **D)** Average zIdPhis across training for each trial type. The first training day, the final criteria day for a training phase, and a midpoint day were chosen. **E)** Same as D, but breaking out first day of CD and D-E. zIdPhi for all premise pairs for the days is plotted. **F)** zIdPhi values for correct vs. incorrect trials. Each colored line corresponds to the average values for each rat. Black marker and line indicate mean and SEM from all trials from all rats. Statistical testing employed bootstrapping to downsample the correct trials to the incorrect trials. The average p-value is shown; $p < 0.05$ for 100% of permutations. n.s., not significant. *, $p < 0.05$; **; $p < 0.01$; ***, $p < 0.001$; ****, $p < 0.0001$.

VTEs have long been regarded as a behavioral hallmark of deliberation, especially during challenging tasks that may rely on memory schema. We were interested in how VTE-like behavior occurred during inference testing when rats were

challenged with new pairs. We found the zIdPhi values were significantly different across trial types on test day (F(5,938), 5.51, $p < 0.001$; **Figure 5A**). Perhaps most critical is that zIdPhi values were significantly higher for B-D trials compared to B-C trials ($p = 0.012$) as well as C-D trials ($p < 0.0001$). B-D to B-C is of particular note as rats make the same choice (go to B), the new inference test configuration led to increased zIdPhi values, possibly as a result of needing time to perform inferential reasoning on the memory schema of the value hierarchy. As with training, we also found that zIdPhi values were significantly higher for incorrect trials than correct trials on test day (100% of bootstrap permutation tests had p values < 0.05; mean p value < 0.0001; 95% confidence intervals for effect size: -0.26 to -0.18; **Figure 5B**). Individually, five out of the ten rats had mean p-values less than 0.05. Thus, when rats are less ballistic in their routes, they are more likely to get the trial wrong.

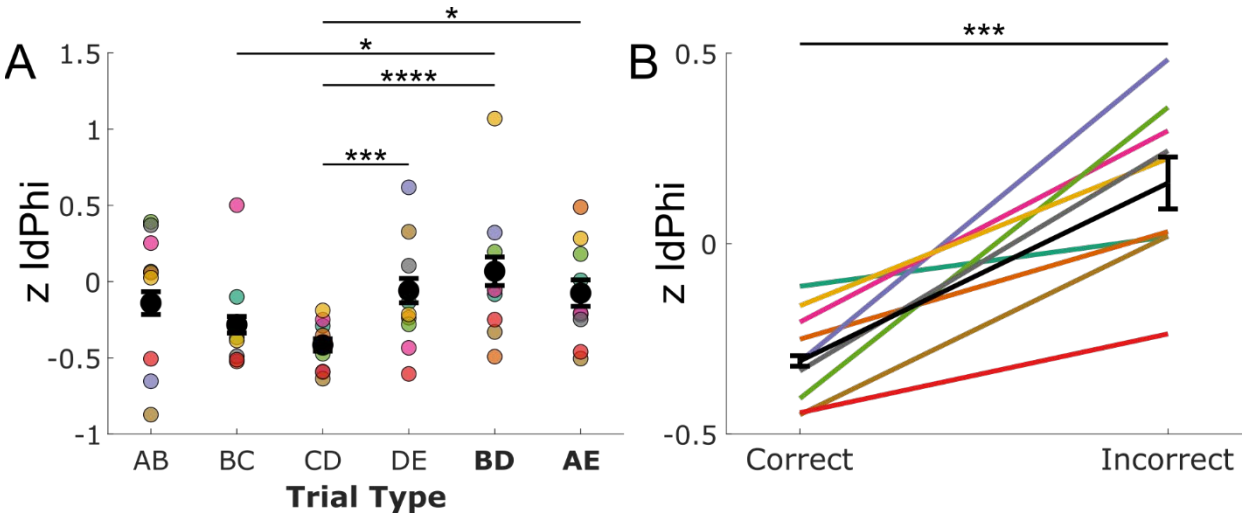

Figure 5: VTEs on inference test day. A) zIdPhi values across trial types on testing day. Each point is the average zIdPhi for each rat. Black circles represent the mean with SEM from all rats and all trials. **B)** zIdPhi values for correct vs. incorrect trials. Each colored line corresponds to the average values for each rat. The black line indicates the mean and SEM from all trials from all rats. Statistical testing employed bootstrapping to downsample the correct trials to the incorrect trials. The average p-value is shown; $p < 0.05$ for 100% of permutations. *, $p < 0.05$; **; $p < 0.01$; ***, $p < 0.001$; ****, $p < 0.0001$.

Discussion

We demonstrate the feasibility of a spatial transitive inference task for rats utilizing an automated radial maze. A majority of the animals trained on the task were able to learn the premise pairs and reach high levels of performance, and only one of nine trained animals failed the inference test. We investigated deliberation-like behavior by analyzing vicarious trial and error events, quantified using zIdPhi values, during training and testing. We found a marked drop in zIdPhi values over training as rats

became proficient. Adding a new premise pair early on resulted in high zIdPhi values for that pair, but adding a second new premise pair did not increase its zIdPhi values beyond the previous premise pairs. On test day we found zIdPhi values to be elevated for the test pairs (B-D and A-E). We also found a robust increase in zIdPhi values for incorrect trials across both training and testing. Overall, we demonstrate that rats can utilize inferential reasoning to successfully decide on the most appropriate actions in novel situations.

The most common rodent TI paradigm utilizes odors as the stimuli mapped to the items of the value hierarchy. The odors are mixed into sand and placed in pots where rats have to dig for a reward. While these studies were critical to establishing rodents' abilities to perform transitive inference (Davis, 1992; Roberts & Phelps, 1994; Van Elzakker et al., 2003) and the brain regions necessary for inferential reasoning (DeVito, Kanter, et al., 2010; Dusek & Eichenbaum, 1997; Van der Jeugd et al., 2009), there are many drawbacks. Manual baiting and setup of every trial leads to lower throughput and could introduce variability to experimental conditions (e.g., depth of reward, pot positions). It is also ambiguous to determine the precise time period when rats may be perceiving the scents and making an inference judgment. Our version of the TI paradigm addresses these issues by using an automated maze and having each arm represent an item in the value hierarchy. It's important to note that while our study uses space (the distinct arms) as the items, the spatial layout is not inherently related to the value hierarchy, as was the case in some rodent TI tasks (Roberts & Phelps, 1994). The radial maze also inherently has a decision area in the center of the maze where rats can perceive the open arms and make their decision. Furthermore, a rat's choice of an arm is easier to quantify than the traditional time digging metric. Our new task, coupled with precise position tracking and electrophysiological recordings, will give valuable insights into the neuronal mechanisms of transitive inference such as the role of theta and replay sequences in schema formation and decision making. Such studies would also allow the comparison of rodent neural activity to non-human primates and humans during inferential reasoning to investigate shared and distinct mechanisms, such as common neural manifolds (Ramawat et al., 2023).

Despite the differences of our spatial TI task and training schedule from the traditional scented-pot version, the rats' behavior and performance are remarkably similar to the existing literature. Across studies, premise pair acquisition failure rates are similar, with about 20-25% of the rats trained failing to learn all premise pairs on five-item sets (Davis, 1992; Dusek & Eichenbaum, 1997; Silverman et al., 2015; Van Elzakker et al., 2003). In our task, rats generally end on higher performance levels (~90%) compared to the odor tasks (~80%) despite our criterion of only 75%. Surprisingly, rats in this spatial task do not appear to have the typical "V" shaped performance curve. Generally, rodent and NHP (Brunamonti et al., 2016) subjects'

average performance is highest for the end premise pairs (e.g., A-B, D-E) and lowest for the “inner pairs” (e.g., B-C, C-D). This is likely due to the end pairs having a simple rule to get correct, i.e., always pick A and never pick E. The lack of V-shaped performance on our task could be attributed to a reward provided on arm E for correct D-E trials. In contrast, in traditional paradigms, a reward is never encountered on the end item. Thus, our rats may have a higher propensity to choose arm E on D-E trials as rewards have been encountered there. We also observed that some rats had reduced performance on original premise pairs on testing day, despite successful B-D inference. We hypothesize that the availability of a new trial type in B-D leads to uncertainty that affected performance on original premise pairs, possibly caused by schema assimilation and accommodation.

In addition, it appears that training durations are remarkably similar across rodent TI paradigms. In our paradigm of intermixed premise pairs, rats took an average of 21 days to complete training. Intermixing trials is more common in NHP (Gillan, 1981) studies compared to rodent ones, with previous reports that rats cannot learn with intermixed training (Donald, 1987). Previous studies using blocks of trials that transition into intermixed report training time in days were similar at 26 (DeVito, Kanter, et al., 2010) and 18 (DeVito, Lykken, et al., 2010). Interestingly, a touchscreen version of the task also took 22 days of training (Silverman et al., 2015). However, our rats required an order of magnitude more trials, averaging 2,215 (737 S.D.), compared to block design tasks. The training regime of most other TI tasks is the presentation of blocks (usually starting around 10 trials long) of each premise pair until the criterion is reached and the block size is reduced until the final phase of intermixing. Where reported, rats required 300-400 total trials to learn the odor-based version of the task (Dusek & Eichenbaum, 1997; Van der Jeugd et al., 2009). Based on the training methods, we estimate the studies reporting training days had similar trial counts of a few hundred (DeVito, Kanter, et al., 2010; DeVito, Lykken, et al., 2010). Thus, while total training days may be similar, it’s likely more efficient in terms of total trials to train rats with blocks of specific premise pairs before transitioning to intermixed presentation. Furthermore, these differences raise important questions on the most efficient way to build memory schemas.

To our knowledge, there is only one other rodent spatial transitive inference task that was developed for mice (Abdou et al., 2024). In their version of the task, a corridor has five distinct (color, geometry, and texture) “rooms” flanking each side. Each context was assigned an item on the hierarchy (A>B>C>D>E) and mice made their choice between two rooms by staying in one of them for 10 seconds. However, direct comparisons to this paradigm are challenging. Their training schedule consisted of only two premise pairs a day with no intermixing of all four premise pairs in one training session. Furthermore, mice in their task failed the first inference test and only

successfully inferred after sleeping and being retested. As a result, we hesitate to make direct comparisons to their valuable study.

Unlike previous rodent TI paradigms, our radial maze version has a clear deliberation area that can be used to investigate potential reasoning behaviors such as vicarious trial and error events. In addition, by serially adding premise pairs, we have distinct time points where the task gets more challenging. Traditionally, VTEs have been ascribed to reflect deliberation as rodents vicariously sample potential choices (Redish, 2016). Furthermore, it's been proposed that VTE rates peak during learning and diminish as choices become automated. Neural recordings in spatial decision-making tasks have demonstrated the relevance of the hippocampal-prefrontal circuit in deliberative behaviors reflected in VTEs (Amemiya & Redish, 2018; Hasz & Redish, 2020, 2020; Johnson & Redish, 2007; Miles et al., 2024; Papale et al., 2016; Schmidt et al., 2019; Stout et al., 2022). However, recent studies have posited that VTE behaviors are heterogeneous and may not always reflect deliberation (Miles et al., 2024). Miles reported high VTE during periods of learning that coincided with correct trials, possibly reflecting deliberation. However, a separate time of high VTEs was found when rats were not learning and getting trials wrong, possibly reflecting uncertainty. We find that high zIdPhi values, a quantification of VTEs, are highest early in training and diminish as rats become proficient. Adding in the first additional premise pair, C-D, results in high zIdPhi values. However, this elevation is not seen when D-E is added, potentially reflecting that adding E evokes less uncertainty than adding D. The lower zIdPhi values for D-E may also reflect a more efficient schema assimilation process for D-E, as the rats have more experience with the task compared to when C-D is added. On test day, we find that zIdPhi values are higher for B-D trials compared to B-C, despite rats doing equally well on both. Thus, rats may have higher ZIdPhi values while inferentially reasoning on B-D trials and having lower ZIdPhi values for the well-trained B-C trials. Overall, these findings support the notion that VTEs occur during periods of learning and reflect the rats deliberating their possible choices.

The immediate utility of VTEs in decision making is also unclear. If VTEs primarily reflect deliberation, one might expect that after a VTE, a rodent will have a higher probability of making a correct choice. However, in support of Miles' (2024) data and others (Schmidt et al., 2013; Stout et al., 2022), we also find that ZIpdhi values are consistently higher on incorrect trials than on correct trials both during training and testing. This may further indicate that during schema-based tasks, VTEs can be associated with periods of uncertainty when rats are more likely to make errors. Overall, we further bolster the notion that VTEs can occur under different experimental conditions and cognitive states and further investigations into their cause and purpose are needed, including neurophysiological investigation during behavioral periods of uncertainty and deliberation, which may reveal differences..

To conclude, we've demonstrated the feasibility of a spatial transitive inference task for rats that results in similar inference test success to the common odor-based TI paradigm. Our version of the task, with sequential training and deliberation zone, is more amenable to physiological recordings which are needed to investigate the neural mechanisms underlying inferential reasoning and memory schema. Computational model investigations into the potential mechanisms of solving transitive inference have demonstrated that numerous artificial neural networks can successfully solve TI tasks (Kay et al., 2024; Miconi & Kay, 2025), including hippocampal models (Frank et al., 2003; Kumaran & McClelland, 2012; Wu & B Levy, 2001; Wu & Levy, 1998). We expect that this spatial version of the classic transitive inference task, coupled with physiological recordings from the hippocampus and prefrontal cortex, can aid in elucidating the mechanisms supporting memory schema and inferential reasoning.

References

Abdou, K., Nomoto, M., Aly, M. H., Ibrahim, A. Z., Choko, K., Okubo-Suzuki, R., Muramatsu, S., & Inokuchi, K. (2024). Prefrontal coding of learned and inferred knowledge during REM and NREM sleep. *Nature Communications*, *15*(1), 4566. https://doi.org/10.1038/s41467-024-48816-x

Acuna, B. D., Eliassen, J. C., Donoghue, J. P., & Sanes, J. N. (2002). Frontal and parietal lobe activation during transitive inference in humans. *Cerebral Cortex (New York, N.Y.: 1991)*, *12*(12), 1312–1321. https://doi.org/10.1093/cercor/12.12.1312

Amemiya, S., & Redish, A. D. (2018). Hippocampal Theta-Gamma Coupling Reflects State-Dependent Information Processing in Decision Making. *Cell Reports*, *22*(12), 3328–3338. https://doi.org/10.1016/j.celrep.2018.02.091

André, J. M., Cordero, K. A., & Gould, T. J. (2012). Comparison of the performance of DBA/2 and C57BL/6 mice in transitive inference and foreground and background contextual fear conditioning. *Behavioral Neuroscience*, *126*(2), 249–257. https://doi.org/10.1037/a0027048

Becker, D. V., Unkelbach, C., & Fiedler, K. (2021). Inference in Social Cognition. In *Oxford Research Encyclopedia of Psychology*. https://doi.org/10.1093/acrefore/9780190236557.013.834

Brunamonti, E., Mione, V., Di Bello, F., Pani, P., Genovesio, A., & Ferraina, S. (2016). Neuronal Modulation in the Prefrontal Cortex in a Transitive Inference Task: Evidence of Neuronal Correlates of Mental Schema Management. *The Journal of Neuroscience*, *36*(4), 1223–1236. https://doi.org/10.1523/JNEUROSCI.1473-15.2016

Bryant, P. E., & Trabasso, T. (1971). Transitive Inferences and Memory in Young Children. *Nature*, *232*(5311), Article 5311. https://doi.org/10.1038/232456a0

Davis, H. (1992). Transitive inference in rats (Rattus norvegicus). *Journal of Comparative Psychology*, *106*(4), 342–349. https://doi.org/10.1037/0735-7036.106.4.342

DeVito, L. M., Kanter, B. R., & Eichenbaum, H. (2010). The hippocampus contributes to memory expression during transitive inference in mice. *Hippocampus*, *20*(1), 208–217. https://doi.org/10.1002/hipo.20610

DeVito, L. M., Lykken, C., Kanter, B. R., & Eichenbaum, H. (2010). Prefrontal cortex: Role in acquisition of overlapping associations and transitive inference. *Learning & Memory*, *17*(3), 161–167. https://doi.org/10.1101/lm.1685710

Ding, M., Tomsick, P. L., Young, R. A., & Jadhav, S. P. (2025). Ventral tegmental area dopamine neural activity switches simultaneously with rule representations in the medial prefrontal cortex and hippocampus. *Journal of Neuroscience*. https://doi.org/10.1523/JNEUROSCI.1670-24.2025

Donald, M. A. (1987). *Logical limitations in the rat: A failure to show transitive reasoning.* [Unpublished]. University of Guelph.

Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*, *94*(13), 7109–7114. https://doi.org/10.1073/pnas.94.13.7109

Eichenbaum, H. (1999). The hippocampus and mechanisms of declarative memory. *Behavioural Brain Research*, *103*(2), 123–133. https://doi.org/10.1016/S0166-4328(99)00044-3

Eichenbaum, H. (2000). A cortical–hippocampal system for declarative memory. *Nature Reviews Neuroscience*, *1*(1), 41–50. https://doi.org/10.1038/35036213

Frank, M. J., Rudy, J. W., & O'Reilly, R. C. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus. II. A computational analysis. *Hippocampus*, *13*(3), 341–354. https://doi.org/10.1002/hipo.10084

Gillan, D. J. (1981). Reasoning in the chimpanzee: II. Transitive inference. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*(2), 150–164. https://doi.org/10.1037/0097-7403.7.2.150

Grosenick, L., Clement, T. S., & Fernald, R. D. (2007). Fish can infer social rank by observation alone. *Nature*, *445*(7126), 429–432. https://doi.org/10.1038/nature05511

Hasz, B. M., & Redish, A. D. (2020). Spatial encoding in dorsomedial prefrontal cortex and hippocampus is related during deliberation. *Hippocampus*, *30*(11), 1194–1208. https://doi.org/10.1002/hipo.23250

Heckers, S., Zalesak, M., Weiss, A. P., Ditman, T., & Titone, D. (2004). Hippocampal activation during transitive inference in humans. *Hippocampus*, *14*(2), 153–162. https://doi.org/10.1002/hipo.10189

Johnson, A., & Redish, A. D. (2007). Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. *Journal of Neuroscience*, *27*(45), 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007

Kay, K., Biderman, N., Khajeh, R., Beiran, M., Cueva, C. J., Shohamy, D., Jensen, G., Wei, X.-X., Ferrera, V. P., & Abbott, L. F. (2024). Emergent neural dynamics and geometry for generalization in a transitive inference task. *PLOS Computational Biology*, *20*(4), e1011954. https://doi.org/10.1371/journal.pcbi.1011954

Koscik, T. R., & Tranel, D. (2012). The Human Ventromedial Prefrontal Cortex Is Critical for Transitive Inference. *Journal of Cognitive Neuroscience*, *24*(5), 1191–1204. https://doi.org/10.1162/jocn_a_00203

Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, *119*(3), 573–616. https://doi.org/10.1037/a0028681

Lazareva, O. F., Smirnova, A. A., Bagozkaja, M. S., Zorina, Z. A., Rayevsky, V. V., & Wasserman, E. A. (2004). Transitive responding in hooded crows requires linearly ordered stimuli. *Journal of the Experimental Analysis of Behavior*, *82*(1), 1–19. https://doi.org/10.1901/jeab.2004.82-1

Mcgonigle, B. O., & Chalmers, M. (1977). Are monkeys logical? *Nature*, *267*(5613), 694–696. https://doi.org/10.1038/267694a0

Miconi, T., & Kay, K. (2025). Neural mechanisms of relational learning and fast knowledge reassembly in plastic neural networks. *Nature Neuroscience*, *28*(2), 406–414. https://doi.org/10.1038/s41593-024-01852-8

Miles, J. T., Mullins, G. L., & Mizumori, S. J. Y. (2024). Flexible decision-making is related to strategy learning, vicarious trial and error, and medial prefrontal rhythms during spatial set-shifting. *Learning & Memory*, *31*(7), a053911. https://doi.org/10.1101/lm.053911.123

Muenzinger, K. F. (1956). On the origin and early use of the term vicarious trial and error (VTE). *Psychological Bulletin*, *53*(6), 493–494. https://doi.org/10.1037/h0044135

Nath, T., Mathis, A., Chen, A. C., Patel, A., Bethge, M., & Mathis, M. W. (2019). Using DeepLabCut for 3D markerless pose estimation across species and behaviors. *Nature Protocols*, *14*(7), 2152–2176. https://doi.org/10.1038/s41596-019-0176-0

Papale, A. E., Stott, J. J., Powell, N. J., Regier, P. S., & Redish, A. D. (2012). Interactions between deliberation and delay-discounting in rats. *Cognitive, Affective, & Behavioral Neuroscience*, *12*(3), 513–526. https://doi.org/10.3758/s13415-012-0097-7

Papale, A. E., Zielinski, M. C., Frank, L. M., Jadhav, S. P., & Redish, A. D. (2016). Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making. *Neuron*, *92*(5), 975–982. https://doi.org/10.1016/j.neuron.2016.10.028

Porter, B. S., Olson, J. M., Leppla, C. A., Duvelle, É., Bladon, J. H., Meer, M. A. A. van der, & Jadhav, S. P. (2025). Adapt-A-Maze: An Open-Source Adaptable and Automated Rodent Behavior Maze System. *eNeuro*, *12*(7). https://doi.org/10.1523/ENEURO.0138-25.2025

Preston, A. R., Shrager, Y., Dudukovic, N. M., & Gabrieli, J. D. E. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus*, *14*(2), 148–152. https://doi.org/10.1002/hipo.20009

Ramawat, S., Marc, I. B., Ceccarelli, F., Ferrucci, L., Bardella, G., Ferraina, S., Pani, P., & Brunamonti, E. (2023). The transitive inference task to study the neuronal correlates of memory-driven decision making: A monkey neurophysiology perspective. *Neuroscience & Biobehavioral Reviews*, 105258. https://doi.org/10.1016/j.neubiorev.2023.105258

Ramawat, S., Mione, V., Di Bello, F., Bardella, G., Genovesio, A., Pani, P., Ferraina, S., & Brunamonti, E. (2022). Different Contribution of the Monkey Prefrontal and Premotor Dorsal Cortex in Decision Making During a Transitive Inference Task. *Neuroscience*, *485*, 147–162. https://doi.org/10.1016/j.neuroscience.2022.01.013

Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, *17*(3), 147–159. https://doi.org/10.1038/nrn.2015.30

Roberts, W. A., & Phelps, M. T. (1994). Transitive Inference in Rats: A Test of the Spatial Coding Hypothesis. *Psychological Science*, *5*(6), 368–374. https://doi.org/10.1111/j.1467-9280.1994.tb00287.x

Rosenblum, H. L., Kim, S., Stout, J. J., Klintsova, A. Y., & Griffin, A. L. (2025). Choice Behaviors and Prefrontal–Hippocampal Coupling Are Disrupted in a Rat Model of Fetal Alcohol Spectrum Disorders. *Journal of Neuroscience*, *45*(10). https://doi.org/10.1523/JNEUROSCI.1241-24.2025

Schlichting, M. L., & Preston, A. R. (2016). Hippocampal–medial prefrontal circuit supports memory updating during learning and post-encoding rest. *Neurobiology of Learning and Memory*, *134*, 91–106. https://doi.org/10.1016/j.nlm.2015.11.005

Schmidt, B., Duin, A. A., & Redish, A. D. (2019). Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. *Journal of Neurophysiology*, *121*(6), 1981–2000. https://doi.org/10.1152/jn.00793.2018

Schmidt, B., Papale, A., Redish, A. D., & Markus, E. J. (2013). Conflict between place and response navigation strategies: Effects on vicarious trial and error (VTE)

behaviors. *Learning & Memory*, *20*(3), 130–138. https://doi.org/10.1101/lm.028753.112
Silverman, J. L., Gastrell, P. T., Karras, M. N., Solomon, M., & Crawley, J. N. (2015). Cognitive Abilities on Transitive Inference Using a Novel Touchscreen Technology for Mice. *Cerebral Cortex (New York, NY)*, *25*(5), 1133–1142. https://doi.org/10.1093/cercor/bht293
Steirn, J. N., Weaver, J. E., & Zentall, T. R. (1995). Transitive inference in pigeons: Simplified procedures and a test of value transfer theory. *Animal Learning & Behavior*, *23*(1), 76–82. https://doi.org/10.3758/BF03198018
Stout, J. J., Hallock, H. L., George, A. E., Adiraju, S. S., & Griffin, A. L. (2022). The ventral midline thalamus coordinates prefrontal–hippocampal neural synchrony during vicarious trial and error. *Scientific Reports*, *12*(1), 10940. https://doi.org/10.1038/s41598-022-14707-8
Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, *331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788
Tibbetts, E. A., Agudelo, J., Pandit, S., & Riojas, J. (2019). Transitive inference in Polistes paper wasps. *Biology Letters*, *15*(5), 20190015. https://doi.org/10.1098/rsbl.2019.0015
Tolman, E. C. (1939). Prediction of vicarious trial and error by means of the schematic sowbug. *Psychological Review*, *46*(4), 318–336. https://doi.org/10.1037/h0057054
Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208. https://doi.org/10.1037/h0061626
Van der Jeugd, A., Goddyn, H., Laeremans, A., Arckens, L., D'Hooge, R., & Verguts, T. (2009). Hippocampal involvement in the acquisition of relational associations, but not in the expression of a transitive inference task in mice. *Behavioral Neuroscience*, *123*, 109–114. https://doi.org/10.1037/a0013990
Van Elzakker, M., O'Reilly, R. C., & Rudy, J. W. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus. I. An empirical analysis. *Hippocampus*, *13*(3), 334–340. https://doi.org/10.1002/hipo.10083
Varga, N., Morton, N. W., & Preston, A. R. (2022). Schema, Inference, and Memory. In *Oxford Handbook of Human Memory*. Oxford University Press.
Vasconcelos, M. (2008). Transitive inference in non-human animals: An empirical and theoretical analysis. *Behavioural Processes*, *78*(3), 313–334. https://doi.org/10.1016/j.beproc.2008.02.017
von Fersen, L., Wynne, C. D., Delius, J. D., & Staddon, J. E. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *17*(3), 334–341. https://doi.org/10.1037/0097-7403.17.3.334
Wing, E. A., D'Angelo, M. C., Gilboa, A., & Ryan, J. D. (2021). The Role of the Ventromedial Prefrontal Cortex and Basal Forebrain in Relational Memory and Inference. *Journal of Cognitive Neuroscience*, *33*(9), 1976–1989. https://doi.org/10.1162/jocn_a_01722
Witten, I. B., Steinberg, E. E., Lee, S. Y., Davidson, T. J., Zalocusky, K. A., Brodsky, M., Yizhar, O., Cho, S. L., Gong, S., Ramakrishnan, C., Stuber, G. D., Tye, K. M., Janak, P. H., & Deisseroth, K. (2011). Recombinase-Driver Rat Lines: Tools,

Techniques, and Optogenetic Application to Dopamine-Mediated Reinforcement. *Neuron*, *72*(5), 721–733. https://doi.org/10.1016/j.neuron.2011.10.028

Wu, X., & B Levy, W. (2001). Simulating symbolic distance effects in the transitive inference problem. *Neurocomputing*, *38–40*, 1603–1610. https://doi.org/10.1016/S0925-2312(01)00512-4

Wu, X., & Levy, W. B. (1998). A Hippocampal-Like Neural Network Model Solves the Transitive Inference Problem. In J. M. Bower (Ed.), *Computational Neuroscience: Trends in Research, 1998* (pp. 567–572). Springer US. https://doi.org/10.1007/978-1-4615-4831-7_94

Zalesak, M., & Heckers, S. (2009). The role of the hippocampus in transitive inference. *Psychiatry Research*, *172*(1), 24–30. https://doi.org/10.1016/j.pscychresns.2008.09.008

Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, *75*(1), 168–179. https://doi.org/10.1016/j.neuron.2012.05.010

A

Value hierarchy

A>B>C>D>E

Training	Test Day
Training Phase 1	**Inference Test**
A-B & B-C	B-D
Phase 2	**Control Test**
Add C-D	A-E
Phase 3	*Intermixed*
Add D-E	*with premise pairs*

B

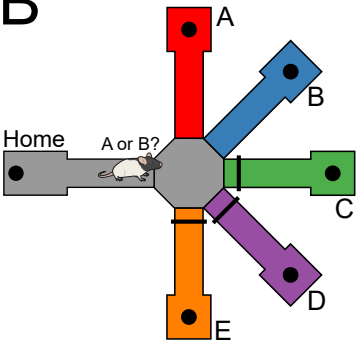

C

Example Correct Trial

Home → A → B → Home

Example Incorrect Trial

Home → B → Home

D

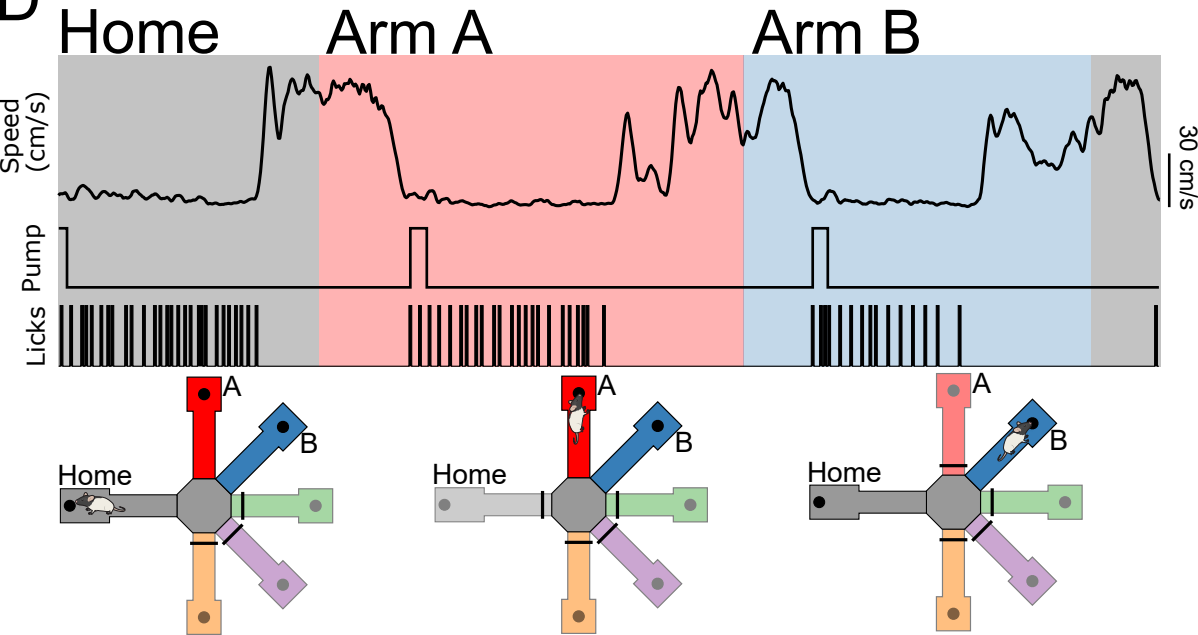

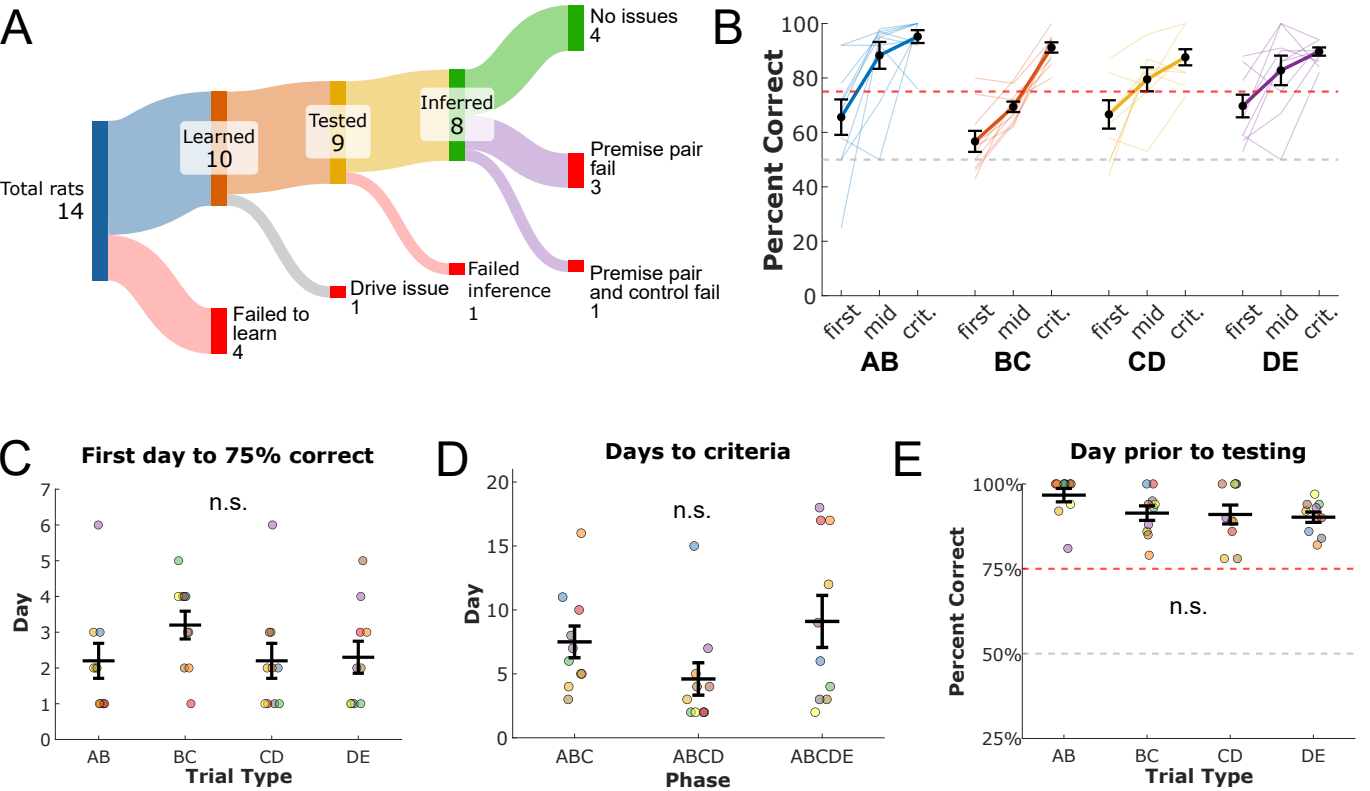
A
Total rats
14
Learned
10
Tested
9
Inferred
8
No issues
4
Premise pair
fail
3
Premise pair
and control fail
1
Failed
inference
1
Drive issue
1
Failed to
learn
4
B
Percent Correct
100
80
60
40
20
0
first
mid
crit.
AB
BC
CD
DE
C
First day to 75% correct
n.s.
Day
7
6
5
4
3
2
1
0
AB
BC
CD
DE
Trial Type
D
Days to criteria
n.s.
Day
20
15
10
5
0
ABC
ABCD
ABCDE
Phase
E
Day prior to testing
Percent Correct
100%
75%
50%
25%
n.s.
AB
BC
CD
DE
Trial Type

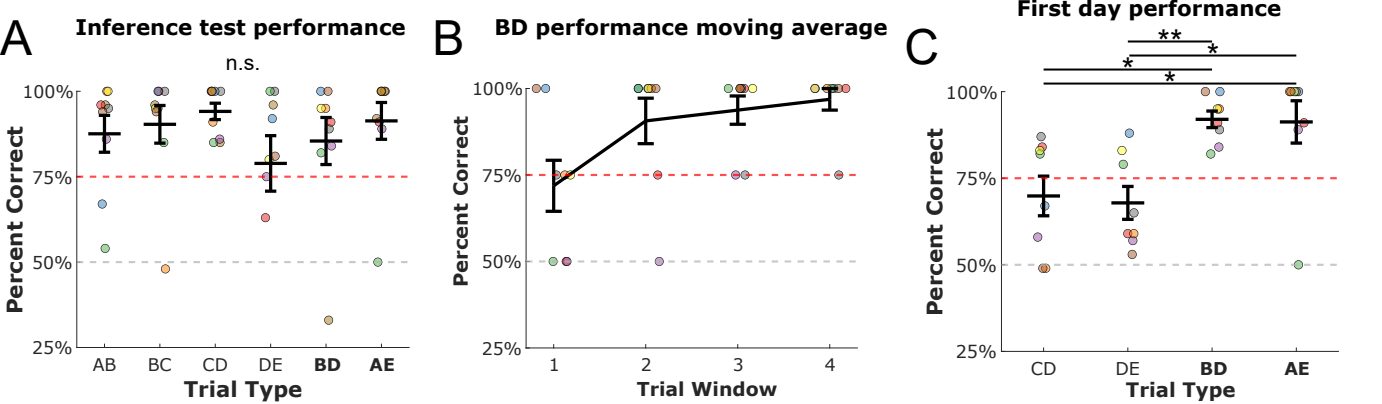
A
Inference test performance
n.s.
100%
75%
50%
25%
Percent Correct
AB
BC
CD
DE
BD
AE
Trial Type
B
BD performance moving average
100%
75%
50%
25%
Percent Correct
1
2
3
4
Trial Window
C
First day performance
**
*
*
*
100%
75%
50%
25%
Percent Correct
CD
DE
BD
AE
Trial Type

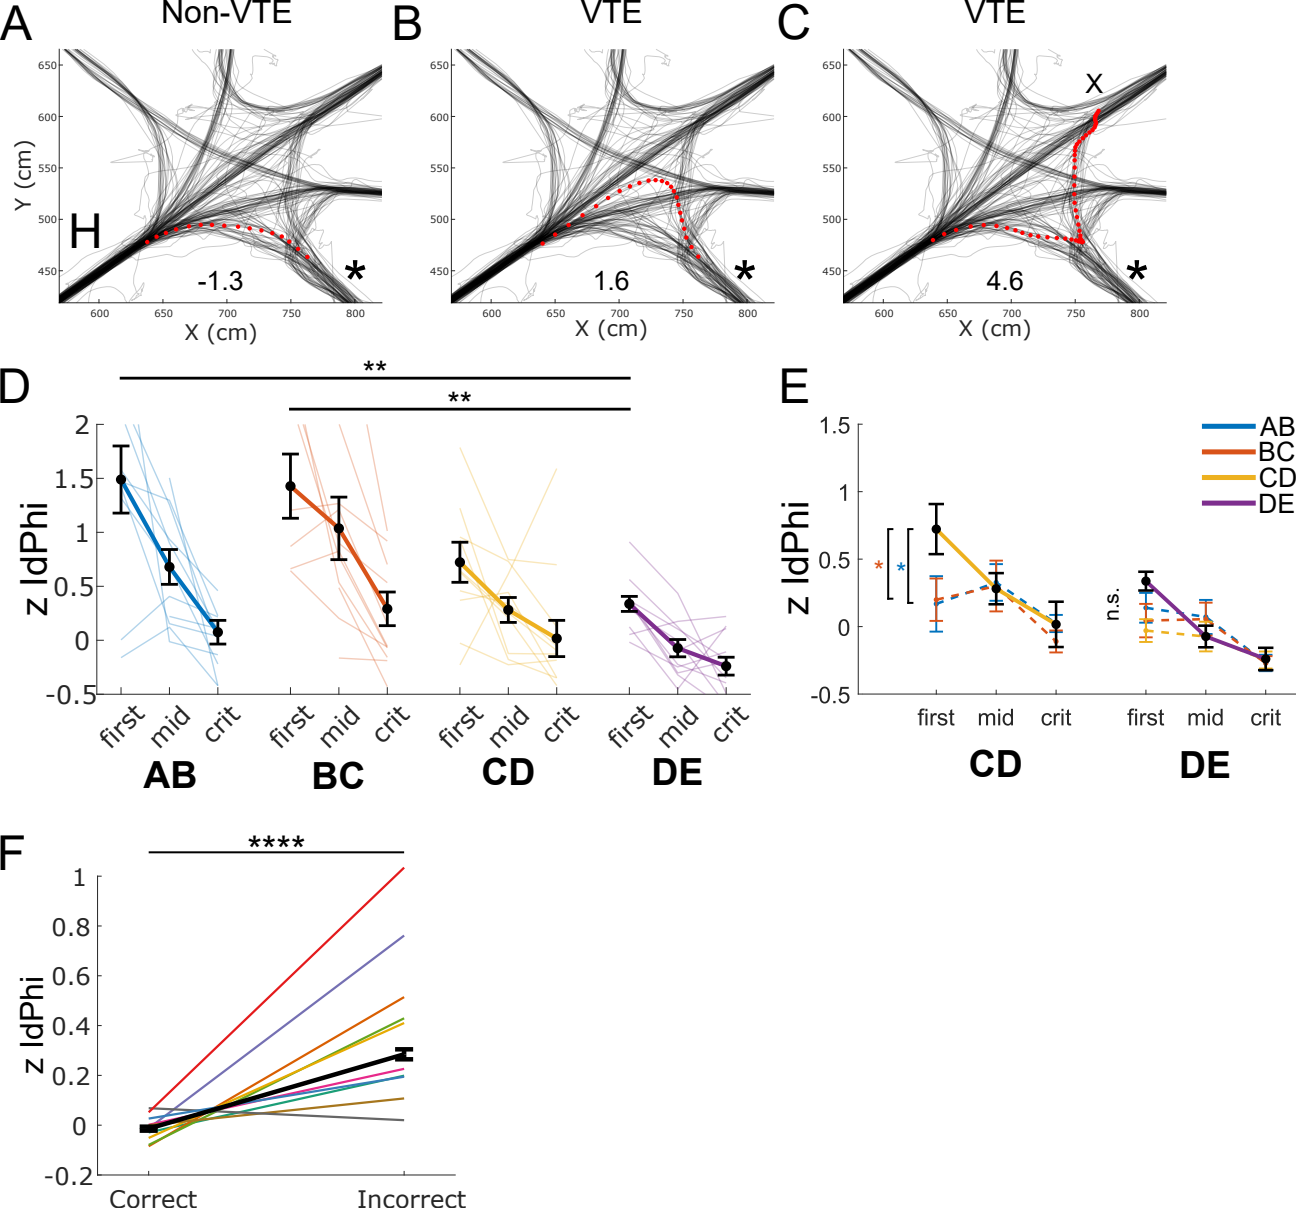
A
Non-VTE
B
VTE
C
VTE
Y (cm)
X (cm)
H
-1.3
1.6
4.6
X
*
D
**
**
z IdPhi
first
mid
crit
AB
BC
CD
DE
E
AB
BC
CD
DE
n.s.
F
****
z IdPhi
Correct
Incorrect

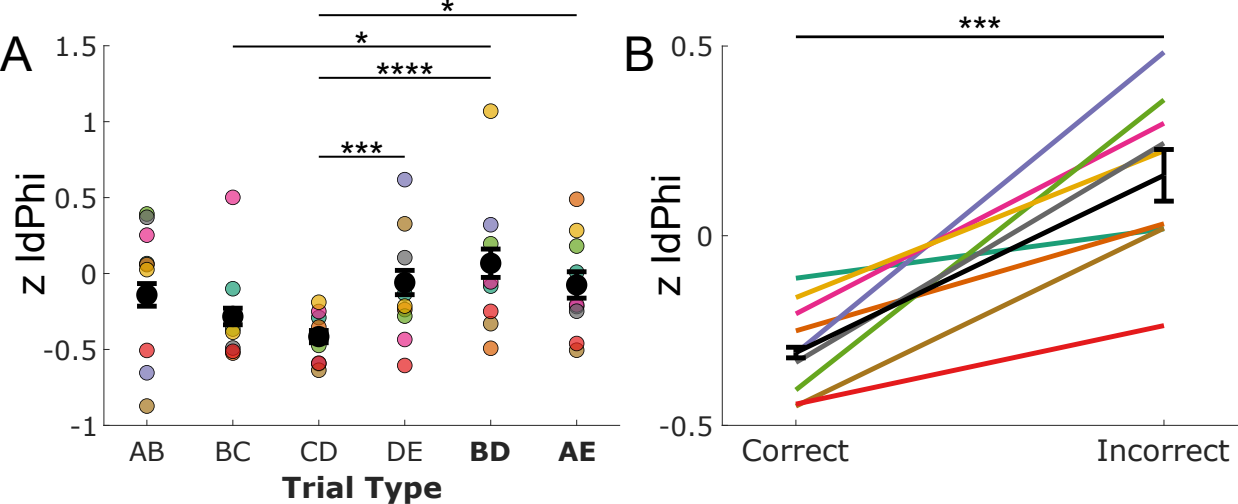
A
B
z ldPhi
1.5
1
0.5
0
-0.5
-1
*
*
****
***
AB
BC
CD
DE
BD
AE
Trial Type
z ldPhi
0.5
0
-0.5
***
Correct
Incorrect