



Chained study and the discovery of relational structure

Douglas B. Markant¹

Accepted: 5 June 2021 / Published online: 15 July 2021
© The Psychonomic Society, Inc. 2021

Abstract

Prior knowledge of relational structure allows people to quickly make sense of and respond to new experiences. When awareness of such structure is not necessary to support learning, however, it is unclear when and why individuals “spontaneously discover” an underlying relational schema. The present study examines the determinants of such discovery in discrimination-based transitive inference (TI), whereby people learn about a hierarchy of interrelated premises and are tested on their ability to draw inferences that bridge studied relations. Experiencing “chained” sequences of overlapping premises during training was predicted to facilitate the discovery of relational structure. Among individuals without prior knowledge of the hierarchy, chaining improved relational learning and was most likely to result in explicit awareness of the underlying relations between items. Observation of chained training sequences was also more effective than the self-generation of training sequences. These findings add to growing evidence that the temporal dynamics of training, including successive presentation of overlapping associations, are key to understanding spontaneous relational discovery during learning.

Keywords Transitive inference · Relational learning · Relational discovery

Introduction

Relational knowledge guides learning and generalization in novel circumstances, as when familiar schemas allow a learner to rapidly encode and reason about new events (Gilboa & Marlatte, 2017; Halford, Wilson, & Phillips, 2010). A fundamental question in cognitive science is how such abstract relational knowledge emerges from experience with fragments of a larger conceptual structure. Understanding this process is especially crucial in light of plentiful evidence that, without explicit instruction or salient relational cues, people often fail to recognize relational structure that connects observed events (Gick & Holyoak, 1983; Goldwater & Gentner, 2015, 2018).

The present study examines how the dynamics of learning affect whether people discover a common relational

structure across a set of interrelated experiences. I focus on the problem of discrimination-based transitive inference (TI), a domain in which there is clear evidence for such spontaneous discovery but a poor understanding of its causes. In transitive inference, a set of items are organized in a linear hierarchy (e.g., $A < B < C < D$). Participants learn a set of *premises* which correspond to discriminations between adjacent items in the hierarchy, where the higher ranked item in any pair is reinforced (A–B+, B–C+, C–D+, etc., with a + following the item that is reinforced in each premise). After learning to choose the reinforced item in each premise, people are then tested on their ability to make transitive inferences when faced with novel, non-adjacent pairs that were never experienced during training (e.g., A–C+).

Previous studies of discrimination-based TI have found that even without the aid of explicit instruction or relational cues, some people spontaneously discover the hierarchical organization of the items during training (Lazareva, 2012; Vasconcelos, 2008). In some cases, this explicit awareness of the hierarchy has been linked to improved relational inference at the time of test (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005), suggesting that the discovery of a familiar relational schema may support rapid encoding and integration of individual premises into a unified representation of the hierarchy. However, little is known about the factors which

A preliminary report of these findings was presented at the 2020 Annual Meeting of the Cognitive Science Society (see <http://psyarxiv.com/cp8bm/>).

✉ Douglas B. Markant
dmarkant@uncg.edu

¹ Department of Psychological Science, University of North Carolina at Charlotte, Colvard South Building, UNC Charlotte, 9201 University City Blvd., Charlotte, NC 28223, USA

support the discovery of the latent hierarchy and any ensuing benefits for relational learning in this task.

This study builds on prior work showing that the dynamics of training—specifically, the manner in which sequences of training examples are generated—can influence whether people identify abstract rules or relations when learning concepts (Birnbaum, Kornell, Bjork, & Bjork, 2013; Doumas, Hummel, & Sandhofer, 2008; Gentner, 2010; Markant & Gureckis, 2014b). In a similar vein, a main goal of the present work is to identify training conditions which help people discover latent relational structure and thereby make sense of seemingly disparate or contradicting experiences during study. Following on a recent study in a related task (Markant, 2020), I examine whether two kinds of training facilitate spontaneous discovery in discrimination-based TI: 1) experiencing “chained” sequences of overlapping premises from trial to trial, and 2) active control over the selection of premises for study.

Spontaneous discovery in transitive inference

Transitive inference has long been seen as a hallmark of logical reasoning (for a review, see Vasconcelos 2008). In early studies of TI the hierarchical organization of the items was readily apparent due to salient properties of the stimuli or the meaning of the relations themselves. Spatial relations (e.g., positions in a linear array) or shared physical features (e.g., relative lengths) naturally imply the property of transitivity. For instance, given the premises *Bob is taller than Dina* and *Dina is taller than Mark*, people easily infer that *Bob is taller than Mark*. These standard variants of TI are associated with explicit reasoning strategies such as logical deduction (Clark, 1969) or the integration of premises into a unified mental representation of the hierarchy (De Soto, London, & Handel, 1965; Hummel & Holyoak, 2001). These strategies involve awareness of the hierarchy or the reasoning process itself, as evinced by their reliance on working memory (Libben & Titone, 2008; Vandierendonck & De vooght, 1997) and sensitivity to relational complexity (Clark, 1969; Waltz et al., 2004).

More recent research has demonstrated that explicit awareness of the hierarchy may not be necessary for transitive inference (for reviews see Lazareva 2012; Vasconcelos 2008). These studies have relied on *discrimination-based* TI tasks which lack any overt cues to the underlying hierarchical organization of the items (Delius & Siemann, 1998; Dusek & Eichenbaum, 1997; Frank et al., 2005, 2006; Greene et al. 2001; Leo & Greene, 2008; Libben & Titone, 2008; McGonigle & Chalmers, 1977; Smith & Squire, 2005). Each premise is instead presented as a choice between two items that are adjacent in the latent hierarchy, only one of which is reinforced because it is ranked higher (A–B+, B–C+, C–D+, etc., with a + following the item that

is reinforced in each premise). After learning through trial and error to select the reinforced item for each premise, people (and many non-human animals; see Vasconcelos, 2008) are able to make transitive inferences for novel, non-adjacent pairs, even if they remain unaware of the underlying hierarchy. These findings have lent support to implicit accounts of TI based on associative or reinforcement learning which do not rely on explicit logical reasoning.

Although these findings suggest explicit awareness may not be necessary for TI, it is typically associated with both faster learning of premises and more accurate inference. Performance in discrimination-based TI improves when participants are directly informed about the hierarchy prior to training (Greene, Spellman, Levy, Dusek, & Eichenbaum, 2001; Libben & Titone, 2008), are first exposed to a familiar rank-ordered example (i.e., playing cards, Moses et al., 2010), or when the choice feedback serves as a cue to an item’s rank in the hierarchy (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Siemann & Delius, 1996). Interestingly, even in the absence of such direct instruction or hints about the relational structure, several studies have shown that some participants spontaneously discover the hierarchy during training. This “serendipitous” awareness has similarly been linked to faster learning and more accurate inference (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005), but the evidence for this relationship is mixed (Delius & Siemann, 1998; Frank, Rudy, Levy, & O’Reilly, 2005; Greene et al., 2001). Little is known about why some participants detect the underlying relational structure, as well as the circumstances in which that discovery enhances relational inference. Discrimination-based TI is thus an ideal setting to examine how explicit relational knowledge spontaneously emerges during the course of trial-and-error learning.

Structuring training to promote spontaneous discovery

Past research has shown that the order of examples during training affects whether people learn abstract representations, including in category learning (Birnbaum et al., 2013; Carvalho & Goldstone, 2015; Elio & Anderson, 1984; Sana, Yan, & Kim, 2017) and relational learning (Don, Goldwater, Greenaway, Hutchings, & Livesey, 2020; Gentner, 2010; Goldwater, Don, Krusche, & Livesey, 2018). Collectively, this work shows that comparison across training examples drives the discovery of abstract relational features (Doumas et al., 2008; Gentner, 2010). Because people are more likely to compare examples that follow one another, juxtaposing related examples can draw attention to shared features and promote the discovery of an abstract rule or analogical mapping (Goldwater et al., 2018).

The same process may play a role in relational discovery in discrimination-based TI. A key challenge in this task is that non-endpoint items are involved in premises with contradicting contingencies (e.g., B is always reinforced when presented with A, but is never reinforced when presented with C). In addition, these items are typically reinforced at similar overall rates across training, further obscuring the latent hierarchy. To make sense of these contradictions, the learner must recognize a common pattern across overlapping sets of premises (e.g., B is “better” than A, but “worse” than C; C is “better” than B, but “worse” than D) which is consistent with a rank-ordered organization of the items. People may therefore be more likely to discover the hierarchy when they experience “chained” sequences of overlapping premises in successive trials (e.g., A–B+, followed by B–C+, followed by C–D+, etc.), making the relational commonality more salient.

Chained sequences improve inference accuracy in standard TI tasks in which participants are already aware of the hierarchy (Andrews, 2010; Markant, 2020; Waltz et al., 2004). One reason for this benefit is that experiencing chained premises allows a learner to integrate premises into a unified mental representation of the hierarchy, an effortful process that depends on holding information about overlapping premises in mind. Although past work has shown that presentation order can affect inference in discrimination-based TI (Lazareva, Gazes, Elkins, & Hampton, 2020; Wynne, 1995), no existing studies have directly examined the effects of chained study on inference or explicit awareness of the hierarchy. However, the possibility that chaining would provoke discovery of the hierarchy has led some researchers to minimize chaining when studying implicit learning in TI (Ellenbogen et al., 2007; Frank et al., 2005, 2006), while others have speculated that variation in the amount of chaining across studies has led to conflicting evidence for spontaneous discovery and its relationship to performance (Lazareva & Wasserman, 2010; Libben & Titone, 2008).

Does active control aid relational discovery?

A second factor which may impact the discovery of relational structure is the opportunity to control the order in which premises are experienced. Self-directed exploration has long been seen as central to the construction of conceptual knowledge from experience (Bruner, 1961; Phillips, 1995), and there is substantial evidence that active exploration improves memory for studied materials compared to passive observation (Markant, Ruggeri, Gureckis, & Xu, 2016; Murty, DuBrow, & Davachi, 2015; Ruggeri, Markant, Gureckis, & Xu, 2019; Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). Recent work has also shown that active control aids the discovery of abstract relationships in

many forms of conceptual learning (Gureckis & Markant, 2012), including function learning (Henriksson & Enkvist, 2016), category learning (Markant & Gureckis, 2014b), and causal structure learning (Sobel & Kushnir, 2006). In contrast to a predetermined training sequence, active control allows learners to tailor the selection of training examples according to their own uncertainty or hypotheses about the target concept.

In the domain of transitive inference, a recent study showed that active control enhances relational learning in a standard TI task in which participants were aware of the hierarchy (Markant, 2020). Participants were instructed to learn the ranks of individuals in a social hierarchy, with each premise encoding the relationship between an employee and their direct supervisor (e.g., Person A is supervised by Person B; Person B is supervised by Person C). Compared to a passive training condition in which premises were presented in a random order, active participants who controlled the selection of premises performed better on tests of transitive inference, an advantage that arose in part from their preference to chain premises during training. This study suggests that if learners have the appropriate relational schema in mind, active control gives them the opportunity to order premises in a way that is more effective for learning the correct hierarchy.

Active selection of premises has not been previously studied in the context of discrimination-based TI. As such, it is unknown whether active learners who are unaware of the hierarchy would have a similar preference to chain premises during study, or if they are more likely to discover the hierarchy compared to passive conditions in which the training sequence is predetermined. While there is strong evidence in favor of active control in relatively well-defined domains, there is continued debate over whether active learners can search for information effectively in the absence of a clear set of alternative hypotheses or prior knowledge about the domain (Mayer, 2004). In the context of discrimination-based TI, active control may only be beneficial when learners are informed about the nature of the underlying hierarchy. Learners who lack such prior knowledge may fail to generate training sequences that draw attention to the common structure across premises, thereby lowering their chance of discovering the hierarchy.

Overview of the current study

The present study examined the effects of chaining and learner control on the discovery of relational structure in discrimination-based TI. Participants played a card game (Fig. 1) in which they learned to choose cards to reveal hidden rewards. Each card’s rank in an underlying hierarchy determined whether it would be rewarded when paired with other cards. Training trials began with a *stage 1 choice*

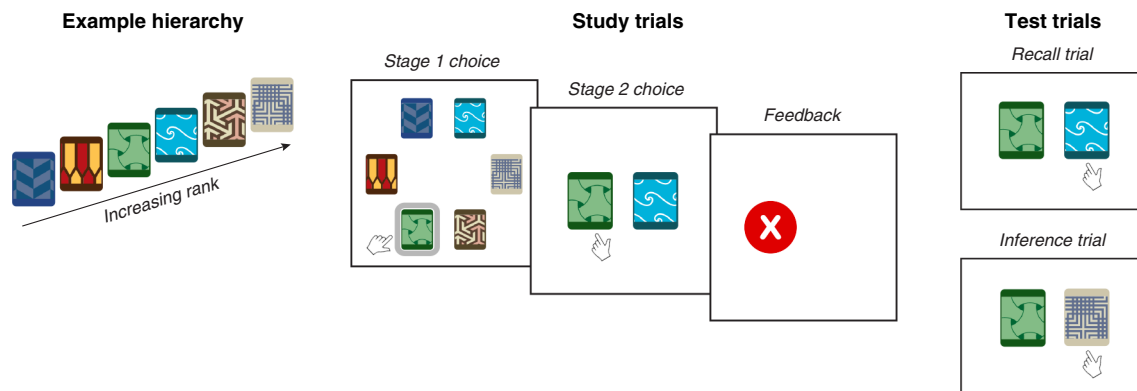


Fig. 1 Depiction of transitive inference task. **Left:** Six cards were arranged in a hierarchy that was unknown to participants. **Middle:** During the training phase, participants learned about premise pairs comprised of items that were adjacent in the hierarchy. In each study trial they selected one item to learn about (stage 1 choice). The selected item was then randomly paired with an adjacent item in the hierarchy

in which an item was selected for study. The selected item was then paired with an item immediately adjacent in the hierarchy, at which point participants made a *stage 2 choice* and received feedback about whether the chosen item was hiding a reward. Stage 2 choices correspond to the typical structure of discrimination-based TI tasks, while the novel stage 1 choices furnish the opportunity to control the order of premises during training. Upon reaching a learning criterion, participants were given a standard forced-choice test which evaluated both recall of studied premise pairs and their ability to make transitive inferences. In addition, they completed two tests of their explicit relational knowledge: 1) A ranking test in which they attempted to rank items and reported their confidence in their chosen order, and 2) a post-task questionnaire which assessed their awareness of the hierarchy.

Participants were randomly assigned to one of three training conditions. There were two passive training conditions in which the training sequences were predetermined, such that participants were forced to select particular items in the stage 1 choices. In the Passive-Frequency condition, items were presented in a pseudo-random order such that all items were selected equally often by the end of training. In the Passive-Adjacent condition, items were selected that were adjacent to the previously chosen item in the hierarchy, thereby creating chained sequences of overlapping premises from trial to trial. Lastly, in the Active training condition participants were free to select any of the items in the hierarchy for study.

As noted above, previous work suggests that benefits from chained study (whether the result of passive observation or active selection) might only emerge when learners are aware of the hierarchy. In anticipation of this possibility, participants' prior knowledge of the hierarchy

and participants chose one item from the pair (stage 2 choice) to reveal whether a reward was hidden beneath it. **Right:** In each test trial, participants were asked to predict which of two cards was hiding a reward. *Recall trials* involved premise pairs that were directly experienced during training, whereas *inference trials* involved novel, non-adjacent pairs

was also manipulated. Participants in the Informed condition were told from the outset about the hierarchical nature of the items, whereas Non-informed participants were simply instructed to learn to pick the correct item in each pairing through trial and error. The Informed condition therefore provides a benchmark for both explicit awareness and the effects of training condition when directly instructed about the hierarchy.

Based on the results of Markant (2020), Passive-Adjacent and Active training were expected to improve inference and ranking performance (compared to the Passive-Frequency condition) when participants were informed about the hierarchy (Informed condition). The central question of the study concerns the effects of training condition among participants who do not have the benefit of that prior knowledge (Non-informed condition). If chained study facilitates the discovery of the hierarchy, Passive-Adjacent training should lead to both higher inference accuracy and greater explicit knowledge as assessed on the ranking test and awareness questionnaire. A similar advantage was expected among Non-informed, Active participants, although this may depend on the extent to which those learners choose to generate chained sequences during training. Finally, if spontaneous discovery allows people to make more productive use of the training period to learn the correct hierarchy, then explicit awareness should be positively related to accuracy on the inference and ranking tests.

Method

The procedures described below were approved by the Institutional Review Board at UNC Charlotte (IRB #18-0558).

Participants

Two-hundred and fifty-two people were recruited from Amazon Mechanical Turk. The sample size was chosen based on a target of approximately 30 participants in each condition, with the possibility that up to 20% of participants may be excluded due to failures to pass attention checks or to reach the learning criterion. Twenty-four individuals were excluded because they failed attention check questions (see Appendix) and a further 23 individuals were excluded because they failed an instruction comprehension question. This left $N = 205$ participants for the analysis (age $M = 36.66$ years, $SD = 10.25$, ranging from 22 to 71 years; 37% female, 43% male, 20% no sex indicated). Participants received a base payment of \$1 and a bonus of up to \$3 ($M = \$2.24$, $SD = 0.55$) based on their performance in the task, which took an average of 21.39 minutes ($SD = 8.02$) to complete.

Materials and procedure

Participants learned about a 6-item hierarchy made up of cards with unique graphical patterns (Fig. 1, left). The task was described as a card game in which the goal was to “learn to pick the right card that is hiding a reward.” Cards were randomly assigned to each rank in the hierarchy for each participant. The rank of each item determined whether it should be selected to find the reward, such that the higher-ranked item in any given pair was always reinforced. The stimuli were designed to avoid any perceptual features which might serve as a cue to a card’s rank in the hierarchy.

The experiment was based on a 2×3 factorial design with instructional condition (Informed or Non-informed) and training condition (Passive-Frequency, Passive-Adjacent, or Active) as between-subjects factors.

Instructional manipulation

The instructional manipulation determined whether participants were told about the hierarchical organization of the cards and occurred at the beginning of the task. Participants in the Non-informed condition saw the following text:

Each card may or may not be rewarded when paired with other cards. Your performance in the game will depend on whether you can learn the correct choice for each pairing of cards.

In contrast, participants in the Informed condition were told there was an underlying hierarchy and given a familiar example:

Each card has a rank that determines whether it will be rewarded over other cards. For example, the

top-ranked card will always be rewarded regardless of what other card it is paired with (just as an ace is ranked higher than all other playing cards), while the bottom-ranked card is never rewarded. Your performance in the game will depend on whether you can learn the correct ranking of the six cards.

A comprehension quiz followed the instructions in which participants had to identify a valid statement about their condition (Non-informed condition: “Each card may or may not be rewarded when paired with other cards”; Informed condition: “Each card’s rank determines when it will be rewarded over other cards.”). There were no further differences between the Informed and Non-informed conditions subsequent to the comprehension quiz.

Training phase

The training phase included up to 10 blocks, with each block comprised of 12 study trials followed by 10 recall trials. Each study trial began when the participant clicked on a circle in the center of the display, causing the six items to be displayed in a ring (Fig. 1, middle). Items were randomly assigned locations in the ring at the beginning of the task and occupied the same locations throughout training. Participants then selected an item for study (the *stage 1 choice*) according to their training condition:

- **Passive-Frequency training:** A predetermined item was highlighted and participants were instructed to select it by clicking on it. The selected item was randomly sampled from the set of items that had been studied the least often up to that point. As a result, this condition produced sequences in which all items were selected with equal frequency by the end of training and repeated selections of any given item tended to be spaced apart.
- **Passive-Adjacent training:** A predetermined item was highlighted and participants were instructed to select it by clicking on it. The selected item was sampled from the set of items that were adjacent in the hierarchy to whichever item had been selected on the previous trial (excluding the first trial, for which the item was randomly selected from the full set). Whichever adjacent item had been studied the least often was chosen; if the adjacent items had been studied an equal number of times then one was chosen at random.
- **Active training:** Participants were free to select any of the six items on every study trial.

Following the stage 1 choice, the selected item was randomly paired with the item either immediately subordinate

or superordinate in the hierarchy to form a premise.¹ Left-right positions of the two items were randomized. Participants in all conditions then chose one of the two items (the *stage 2 choice*) and received feedback about whether it was hiding a reward (a green dollar sign when the higher-ranking item was chosen and a red X when the lower-ranking item was chosen). Feedback was displayed for 1 s, after which the task proceeded to the next trial.

After 12 study trials, participants completed 10 recall trials (two trials per premise) which tested their memory of the studied premises (Fig. 1, right). On each recall trial a premise pair appeared in the center of the display and participants were instructed to select the card that was hiding the reward. No feedback was provided until the end of the block, at which point participants were told the proportion of correct responses. The training phase ended either after 10 blocks or when participants reached a criterion of 100% correct responses in a block, indicating that they chose the higher-ranking card for every premise pair twice.

Test phase

The test phase was comprised of 45 trials, with three repetitions of every possible pairing of items from the 6-item hierarchy. Recall trials involved premise pairs that were directly experienced during the study phase (5 unique pairs), whereas inference trials involved novel pairings of non-adjacent items (10 unique pairs).

In each test trial, two items from the hierarchy appeared side-by-side in the center of the display (Fig. 1, right). As with the recall trials during the training phase, participants were instructed to select the item from each pair that was hiding a reward. No feedback was presented until the end of the study, at which point participants were informed about their overall accuracy.

Awareness questionnaire

Immediately following the test phase, participants responded to a set of questions intended to assess their explicit awareness of the hierarchy (see [Appendix](#) for questions

and response coding). The questions were adapted from a questionnaire used by Kumaran and Ludwig (2013) and Moses, Villate, Binns, Davidson, and Ryan (2008). The *awareness score* was the proportion of questions (out of three) in which a participant's response indicated awareness of the hierarchical organization of the items or the ability to use logical reasoning to draw inferences during the test. Four participants failed to respond to one or more of the awareness questions and were excluded from analyses involving awareness scores.

Ranking elicitation

After the awareness questionnaire, participants were asked to create a linear ranking of the six items in the hierarchy according to "how likely rewards are when you choose them, ranging from low likelihood of rewards on the left to high likelihood of rewards on the right." The six cards were displayed in a random order and the position of each item could be changed by clicking on arrow buttons. Participants were self-paced and could make any number of changes before recording their response. Ranking accuracy was the proportion of items that were in the correct position. Following the ranking elicitation, participants also rated their confidence that they ranked the items correctly on a scale from 1 (not at all confident) to 5 (completely confident).

Results

Overview of analyses

The main analyses focused on the effects of instructional condition (Non-informed or Informed) and training condition (Passive-Frequency, Passive-Adjacent, or Active) on each dependent measure (number of blocks to criterion, recall accuracy, inference accuracy, ranking accuracy, ranking confidence, and awareness). Additional analyses are then presented which examine how test performance was related to post-task awareness and the makeup of the training experience.

Unless stated otherwise, regression models included two between-subjects factors (instructional condition and training condition) and their interaction as predictors. Omnibus tests were conducted using analysis of variance (for linear regression models) or analysis of deviance (for logistic regression models). In addition, planned pairwise contrasts were performed within each instructional and training condition using the *multcomp R* package with adjustments for multiple tests (Hothorn, Bretz, & Westfall, 2008). All analyses were conducted in *R* (R Core Team, 2018).

¹Random pairing was necessary to ensure that stage 1 choices were not predictive of reinforcement in stage 2 (e.g., if the selected item was always paired with the immediately superordinate item, participants could learn a simple rule to choose the new option in stage 2 without learning the premises themselves). A consequence of this design is that some trials in the Passive-Adjacent condition did not involve chaining, since selecting an adjacent item in stage 1 sometimes led to the repetition of the same premise from the previous trial or a non-overlapping premise in stage 2. Thus, the Passive-Adjacent condition maximizes the frequency of chaining (overlapping premises in successive trials) under the constraint of random pairing in stage 2.

Training

$N = 20$ participants failed to reach the training criterion (100% correct in a block of 10 recall trials) and were excluded from further analysis. More participants failed to reach the criterion in the Non-informed condition ($N=14$) than the Informed condition ($N=6$), a marginally significant difference ($\chi^2(1) = 3.72, p = 0.05$). Among participants who reached the training criterion, there was also a main effect of instructional condition on the number of blocks to criterion ($F(1, 179) = 10.23, MSE = 4.18, p = .002, \hat{\eta}_G^2 = .054$), with Informed participants requiring fewer blocks ($M = 3.21, SD = 2.12$) than Non-informed participants ($M = 4.16, SD = 2.02$). In contrast, there were no significant differences between training conditions in either the number of participants who failed to reach the criterion or the number of blocks to criterion. Thus, prior knowledge of the hierarchy led to more efficient acquisition of the premises, but the type of training sequence did not impact participants' ability to learn the studied premises in the training period.

Test accuracy

Participants completed a standard forced-choice test of their ability to choose the higher ranked item in each possible pairing of items from the hierarchy, including studied premises (recall trials) and novel, non-adjacent pairs (inference trials). Test responses were scored according to whether participants correctly chose the superordinate item in each test pair (0 = incorrect, 1 = correct) and the proportion of correct responses was modeled using logistic regression.

Recall trials

Accuracy on recall trials involving studied premise pairs was generally high in the final test (Fig. 2A, left), but was higher in the Informed condition than the Non-informed condition ($\chi^2(1) = 8.12, p = 0.004$). There was no effect of training condition ($\chi^2(2) = 4.16, p = 0.12$) and no interaction ($\chi^2(2) = 0.01, p = 0.99$), indicating that all participants could remember the correct responses for the studied premises regardless of training type.

Inference trials - Endpoint

Performance on inference trials was analyzed separately based on whether the test pair included an *endpoint* item (either the lowest- or highest-ranking item in the hierarchy). Endpoint and non-endpoint trials are typically separated in discrimination-based TI because they may involve qualitatively different strategies (Dusek & Eichenbaum,

1997; Smith & Squire, 2005). In particular, because endpoint items are associated with constant reinforcement histories (e.g., the lowest-ranking item is never reinforced), learners can respond accurately based on a single endpoint item without any inferential reasoning or comparison with the other item in the pair.

On inference trials involving an endpoint (Fig. 2A, middle) there was a main effect of instructional condition on accuracy ($\chi^2(1) = 47.97, p < .001$), with accuracy higher among Informed participants than Non-informed in each training condition (Passive-Frequency: $OR = 1.95 [1.34, 2.82], z = 4.88, p < .001$; Passive-Adjacent: $OR = 2.12 [1.31, 3.41], z = 4.27, p < .001$; Active: $OR = 1.47 [1.02, 2.13], z = 2.83, p = 0.04$). In addition, there was a main effect of training condition ($\chi^2(2) = 63.15, p < .001$), but no interaction ($\chi^2(2) = 3.36, p = 0.19$). Within the Non-informed condition, inference accuracy was higher in the Passive-Adjacent condition than both the Passive-Frequency ($OR = 2.13 [1.46, 3.12], z = 5.41, p < .001$) and Active conditions ($OR = 1.70 [1.16, 2.50], z = 3.73, p = 0.002$), whereas accuracy did not differ between the Active and Passive-Frequency conditions ($OR = 1.25 [0.88, 1.79], z = 1.74, p = 0.42$). Similarly, within the Informed condition, Passive-Adjacent inference accuracy was higher than both the Passive-Frequency ($OR = 2.32 [1.45, 3.71], z = 4.87, p < .001$) and Active conditions ($OR = 2.44 [1.54, 3.89], z = 5.23, p < .001$) but there was no difference between Active and Passive-Frequency conditions ($OR = 0.95 [0.65, 1.39], z = -0.37, p = 1.00$).

Inference trials - Non-endpoint

Compared to endpoint trials, non-endpoint inference trials provide a stronger test of relational learning because they require the integration of multiple premises to identify the higher ranked item. A similar pattern of results was obtained as for endpoint trials (Fig. 2A, right). There was a main effect of instructional condition ($\chi^2(1) = 59.76, p < .001$) such that accuracy was higher among Informed participants in each training condition (Passive-Frequency: $OR = 2.63 [1.61, 4.29], z = 5.35, p < .001$; Passive-Adjacent: $OR = 1.79 [1.08, 2.99], z = 3.10, p = 0.01$; Active: $OR = 2.35 [1.44, 3.85], z = 4.72, p < .001$). There was also a main effect of training condition ($\chi^2(2) = 18.25, p < .001$) but no interaction ($\chi^2(2) = 2.24, p = 0.33$). Within the Informed condition, there were no significant pairwise differences between training conditions. Within the Non-informed condition, non-endpoint inference accuracy in the Passive-Adjacent condition was significantly higher than the Passive-Frequency condition ($OR = 2.06 [1.27, 3.35], z = 4.07, p < .001$) and marginally higher than accuracy in the Active condition ($OR = 1.57 [0.97, 2.54], z = 2.56, p = 0.08$). Among Non-informed participants, Passive-Adjacent

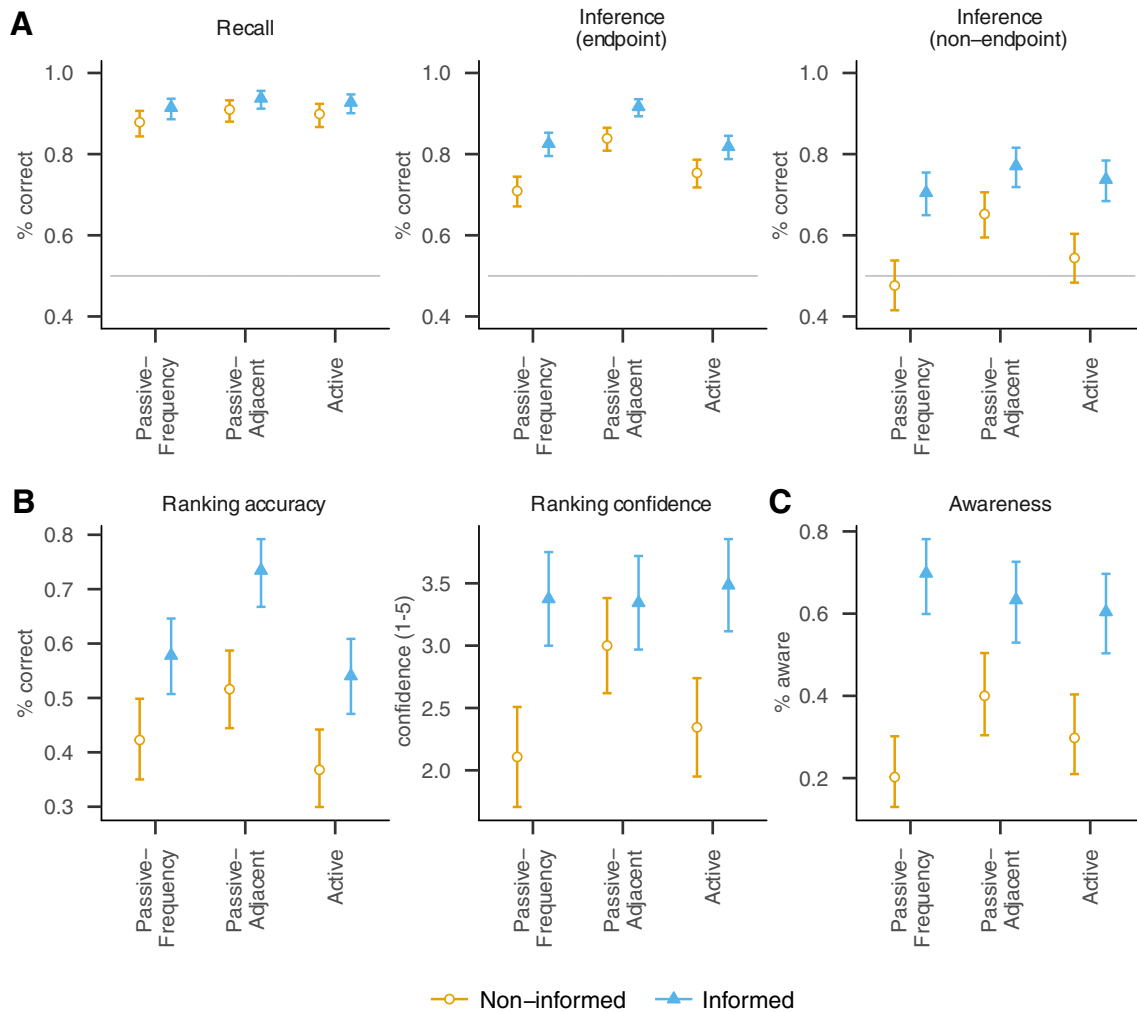


Fig. 2 Means and 95% confidence intervals for performance on the forced-choice test (A), ranking (B), and post-task awareness questionnaire (C). Horizontal lines indicate chance performance on the forced-choice test

training was the only condition that led to overall accuracy that was above chance (Fig. 2A, right).

In sum, despite there being no differences in memory for the studied premises between the training conditions, Passive-Adjacent training led to higher performance for both endpoint and non-endpoint inference. However, this advantage over other training types was most consistent when participants were not informed of the hierarchy beforehand.

Ranking accuracy and confidence

As discussed in the Introduction, successful inference on the forced choice test does not necessarily imply explicit knowledge of the hierarchy, as implicit, associative mechanisms can produce similar patterns of performance (Delius & Siemann, 1998; Dusek & Eichenbaum, 1997; Frank et al., 2005, 2006). Asking participants to rank the

items is a more direct test of explicit knowledge, and in particular, whether Non-informed participants were able to discover the hierarchical nature of the task.

Accuracy on the ranking test was defined as the proportion of six items that were assigned the correct rank (Fig. 2B, left). There were main effects of both instructional condition ($\chi^2(1) = 38.08, p < .001$) and training condition ($\chi^2(2) = 24.85, p < .001$), but no interaction ($\chi^2(2) = 1.23, p = 0.54$). Ranking accuracy was higher among Informed participants in each training condition (Passive-Frequency: $OR = 1.87 [1.05, 3.34], z = 2.93, p = 0.03$; Passive-Adjacent: $OR = 2.59 [1.43, 4.70], z = 4.34, p < .001$; Active: $OR = 2.02 [1.14, 3.59], z = 3.31, p = 0.008$). Within the Informed condition, Passive-Adjacent ranking accuracy was higher than both Passive-Frequency ($OR = 2.02 [1.11, 3.66], z = 3.20, p = 0.01$) and Active conditions ($OR = 2.35 [1.31, 4.23], z = 3.94, p < .001$), while in the Non-informed condition accuracy was higher following Passive-Adjacent

training than Active training ($OR = 1.83 [1.02, 3.29]$, $z = 2.82$, $p = 0.04$). There were no other pairwise differences.

A similar pattern was seen in confidence judgments of elicited rankings (Fig. 2B, right). There was no main effect of training condition ($F(2, 179) = 2.02$, $MSE = 1.17$, $p = .135$, $\hat{\eta}_G^2 = .022$), but there was a main effect of instructional condition ($F(1, 179) = 32.49$, $MSE = 1.17$, $p < .001$, $\hat{\eta}_G^2 = .154$) and a significant interaction ($F(2, 179) = 3.31$, $MSE = 1.17$, $p = .039$, $\hat{\eta}_G^2 = .036$). Within the Non-informed condition, Passive-Adjacent confidence was higher than Passive-Frequency confidence ($\beta = 0.89 [0.12, 1.67]$, $z = 3.16$, $p = 0.01$). No other pairwise comparisons were significant. As was seen for inference accuracy in the forced choice test, the Passive-Adjacent condition was associated with the best performance in terms of ranking accuracy and confidence, with the clearest advantage over other training conditions again emerging when participants were not informed about the hierarchy.

Post-task awareness

Responses to the post-task awareness questions were coded based on whether participants endorsed the statement that indicated explicit awareness of the hierarchy (0 = unaware, 1 = aware; see Appendix). The probability of making “aware” responses was modeled with logistic regression based on three responses for each participant. As expected, post-task awareness (Fig. 2C) was higher among Informed participants than Non-informed participants ($\chi^2(1) = 65.82$, $p < .001$) for all three training conditions (Passive-Frequency: $OR = 9.11 [3.51, 23.61]$, $z = 6.29$, $p < .001$; Passive-Adjacent: $OR = 2.59 [1.13, 5.96]$, $z = 3.10$, $p = 0.01$; Active: $OR = 3.60 [1.52, 8.52]$, $z = 4.04$, $p < .001$), indicating that the instructional manipulation had the intended effect.

Lastly, this analysis provided further evidence that the type of training affected whether Non-informed participants became aware of the underlying hierarchy during the task. There was no main effect of training condition on awareness ($\chi^2(2) = 2.27$, $p = 0.32$), but there was a significant interaction ($\chi^2(2) = 7.85$, $p = 0.02$). Within the Non-informed condition, post-task awareness was higher in the Passive-Adjacent condition than the Passive-Frequency condition ($OR = 2.63 [1.03, 6.73]$, $z = 2.79$, $p = 0.04$). No other pairwise comparisons were significant. In addition to enhanced inference and ranking performance, Passive-Adjacent training was associated with the highest post-task awareness among Non-informed participants, further suggesting that the chaining of premises facilitated the spontaneous discovery of the hierarchy.

Relationship to task performance.

The next analysis examined the relationship between post-task awareness and performance on the forced choice and ranking tests. Although explicit awareness among non-instructed participants has been linked to improved TI (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005), this relationship has not been observed in some studies (Frank et al., 2005; Greene et al., 2001). The regression models described in the previous sections were expanded to include awareness scores as a covariate, along with the full set of interactions between awareness, instructional condition, and training condition. Linear contrasts were then used to test whether there was a significant association between post-task awareness and performance in each condition. Because the results for endpoint and non-endpoint inference accuracy were comparable, they were combined to simplify the remaining analyses.

The estimated effects are shown in Table 1. Among Informed participants, higher awareness scores were consistently related to improved relational learning, including both inference accuracy and ranking accuracy, regardless of training condition. Awareness was also associated with higher ranking confidence in the Passive-Adjacent and Active conditions, fewer blocks to criterion in the Passive-Frequency condition, and higher recall accuracy in the Active condition. Even though all Informed participants were told about the hierarchical organization of the items, this result suggests that variation in attention or understanding of the task within that group were strongly related to task performance.

In contrast, among Non-informed participants, Passive-Adjacent training was the only condition with consistent associations between post-task awareness and relational learning, including higher inference accuracy, ranking accuracy, and ranking confidence (see Fig. S2 in the Supplementary Materials). In the Active condition, awareness was positively related to inference accuracy but no other dependent measures, while in the Passive-Frequency condition there were no significant associations between awareness and performance. In addition to discovery of the hierarchy being more likely in the Passive-Adjacent condition, these results show that individuals in that condition who exhibited explicit awareness were also better able to rank the items and make relational inferences at test.

Comparison of training sequences

The preceding results demonstrate that for Non-Informed participants, Passive-Adjacent training led to higher infer-

Table 1 Associations between post-task awareness and other dependent measures in each condition

Condition	Blocks to criterion	Recall accuracy	Inference accuracy	Ranking accuracy	Ranking confidence
Non-informed					
Passive-Frequency	−0.77 [−2.41, 0.87]	1.17 [0.60, 2.27]	0.83 [0.61, 1.13]	1.46 [0.74, 2.87]	0.26 [−0.50, 1.02]
Passive-Adjacent	−0.68 [−1.75, 0.38]	1.70 [0.99, 2.94]	1.31 [1.03, 1.68] *	1.68 [1.06, 2.64] *	0.63 [0.13, 1.12] **
Active	0.53 [−0.85, 1.91]	1.35 [0.72, 2.55]	1.33 [1.00, 1.76] *	1.55 [0.86, 2.78]	0.57 [−0.07, 1.21]
Informed					
Passive-Frequency	−1.06 [−2.02, −0.10] *	1.36 [0.90, 2.05]	1.42 [1.15, 1.74] ***	1.85 [1.22, 2.80] ***	0.41 [−0.03, 0.86]
Passive-Adjacent	−0.66 [−1.51, 0.19]	1.52 [0.93, 2.46]	2.74 [2.04, 3.68] ***	2.71 [1.73, 4.25] ***	0.70 [0.31, 1.09] ***
Active	−0.52 [−1.46, 0.42]	1.78 [1.10, 2.90] *	1.74 [1.40, 2.16] ***	3.49 [2.14, 5.69] ***	0.59 [0.16, 1.03] **

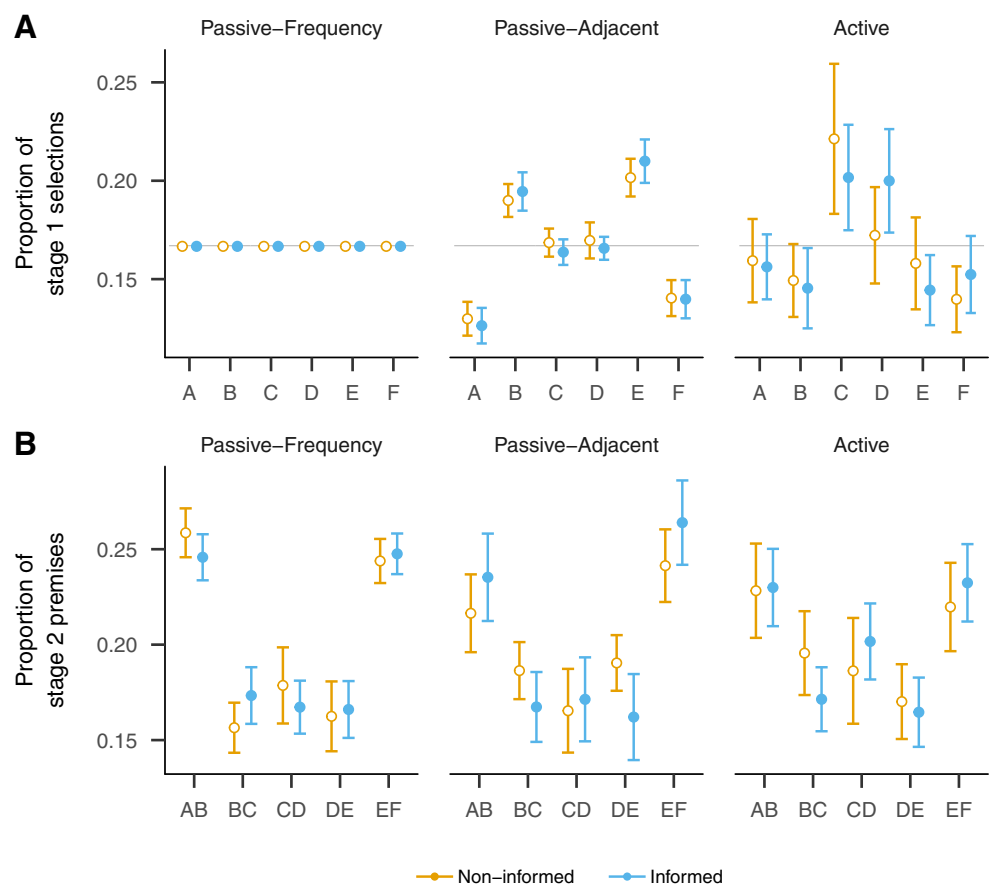
* $p < .05$; ** $p < .01$; *** $p < .001$, adjusted for multiple comparisons. For recall accuracy, inference accuracy, and ranking accuracy, estimates are relative odds ratios (OR). For blocks to criterion and ranking confidence, estimates are linear effects

ence accuracy, ranking ability, and post-task awareness compared to Passive-Frequency and Active training. This is consistent with the hypothesis that chained sequences of overlapping premises during study lead to both improved relational inference and explicit discovery of the hierarchy. However, in addition to differences in the amount of chaining, the training conditions may have also varied in the overall presentation frequency of individual premises. This is in contrast to typical studies of discrimination-based TI in

which premises occur with equal overall frequency across training. In this section I compare the makeup of training sequences across conditions and consider whether there are alternative explanations for the benefits of Passive-Adjacent training aside from chaining.

I first calculated the relative frequency of stage 1 selections by item rank during training (Fig. 3A). By design, Passive-Frequency training resulted in equal selection frequencies of each item. In the Passive-Adjacent condition,

Fig. 3 A: Proportion of stage 1 selections (mean \pm SE) during training by item rank. Horizontal lines indicate the expected frequency from random selection. **B:** Proportion of stage 2 presentations of each premise pair during training



items B and E were selected most often because they were next to the endpoints (e.g., whenever endpoint item A was selected, on the next trial item B would necessarily be selected), and the endpoint items (A and F) were selected least often. Interestingly, the pattern of stage 1 selections was markedly different among Active participants who were free to select any of the six items throughout training. Active participants preferred to select items in the middle of the hierarchy (C or D) in both the Informed and Non-informed conditions, potentially revealing a preference to learn about items with the most variable reinforcement history.

More central to understanding effects on learning is the consequences of these selections for how often each premise appeared in stage 2 choices, when participants chose one item from each pair and received feedback about whether it was hiding the reward. For instance, if non-endpoint items were simply experienced more often in the Passive-Adjacent condition, that might explain higher performance on non-endpoint inference trials, rather than the order in which those premises occurred during training. Figure 3B shows the proportion of stage 2 choices involving each premise. Despite some notable differences in stage 1 selections (Fig. 3A), the relative frequencies of premises in stage 2 was similar across conditions. Training trials were categorized according to whether they included an endpoint and the proportion of endpoint trials was modeled with a two-way ANOVA with training condition, instruction condition, and their interaction as between-subjects factors. There was a main effect of training condition ($F(2, 179) = 6.35$, $MSE = 0.00$, $p = .002$, $\hat{\eta}_G^2 = .066$), but no effect of instruction condition ($F(1, 179) = 2.78$, $MSE = 0.00$, $p = .097$, $\hat{\eta}_G^2 = .015$) and no interaction ($F(2, 179) = 2.30$, $MSE = 0.00$, $p = .103$, $\hat{\eta}_G^2 = .025$). The proportion of endpoint premises was higher in the Passive-Frequency condition than the Active condition ($\beta = 0.04$ [0.01, 0.07], $t = 3.61$, $p = 0.00$), consistent with the more frequent selection of middle items (C or D) in the Active group. However, there were no significant differences between the Passive-Adjacent group and the other conditions. A further analysis of items' reinforcement rates similarly showed that there were no systematic differences that could account for higher performance on inference and ranking tests in the Passive-Adjacent condition (see Supplementary Materials Section S1). These results instead suggest that it is the sequencing of premises—in particular, the chaining of overlapping premises in successive trials—that is responsible for the enhanced relational learning seen in the Passive-Adjacent condition.

Exploration of the hierarchy in the Active condition

The final set of analyses explored active learners' search behavior during training and how their selections related

to performance. Recent work indicates that active learners prefer to create chained sequences when learning how to rank items within a familiar hierarchy (Markant, 2020), but it is unknown whether people show a similar search preference in discrimination-based TI. Although Active groups had lower overall performance than the Passive-Adjacent condition in the present task, individual differences in search behavior (in particular, the tendency to chain premises) might be related to accuracy on tests of relational learning.

On every study trial Active participants were free to select any item from the hierarchy. Selections were scored by their absolute distance to the item selected on the previous trial (excluding the first trial of each block). A distance of 0 indicates that the same item was repeatedly selected, whereas a distance of 1 indicates that an item immediately adjacent in the hierarchy was selected. As in the Passive-Adjacent training, selecting items at a distance of 1 was most likely to produce chained premises from trial to trial.

Figure 4A shows the proportion of selections at each distance among Active participants, with horizontal lines marking the proportions expected from random search. Both Informed and Non-informed participants made repeated selections (distance = 0) more often than expected from random search, but repeated selections were less frequent among Informed participants than Non-informed participants ($OR = 0.52$ [0.45, 0.61], $z = -8.10$, $p < .001$). In contrast, Informed participants were more likely to select adjacent items (distance = 1; $OR = 1.34$ [1.13, 1.59], $z = 3.35$, $p < .001$) and items at a distance of two ($OR = 1.63$ [1.33, 1.99], $z = 4.72$, $p < .001$) and three ($OR = 1.40$ [1.09, 1.80], $z = 2.62$, $p = 0.009$) positions away. The proportions of more distant selections (4 or 5 positions away from the previous item) did not differ between Informed and Non-informed groups. Although both groups tended to repeatedly select the same item in successive trials, prior knowledge of the hierarchy led to a stronger preference to explore items that were nearby in the hierarchy to the item selected on the last trial, including adjacent (distance = 1) items that often resulted in chained premises.

Finally, I examined whether these search behaviors were related to learning performance. Regression models for each dependent variable were expanded to include terms for the proportion of selections at distances of 0, 1 and 2 in each instructional condition.² The estimated effects

²The proportion of selections at distances 3–5 were treated as a third category that was not included in the model. Because the proportions are constrained to sum to 1, an estimated parameter therefore indicates the effect of increasing the predictor while holding the other predictors in the model constant, and would necessarily correspond to a decrease in the proportion of more distant selections.

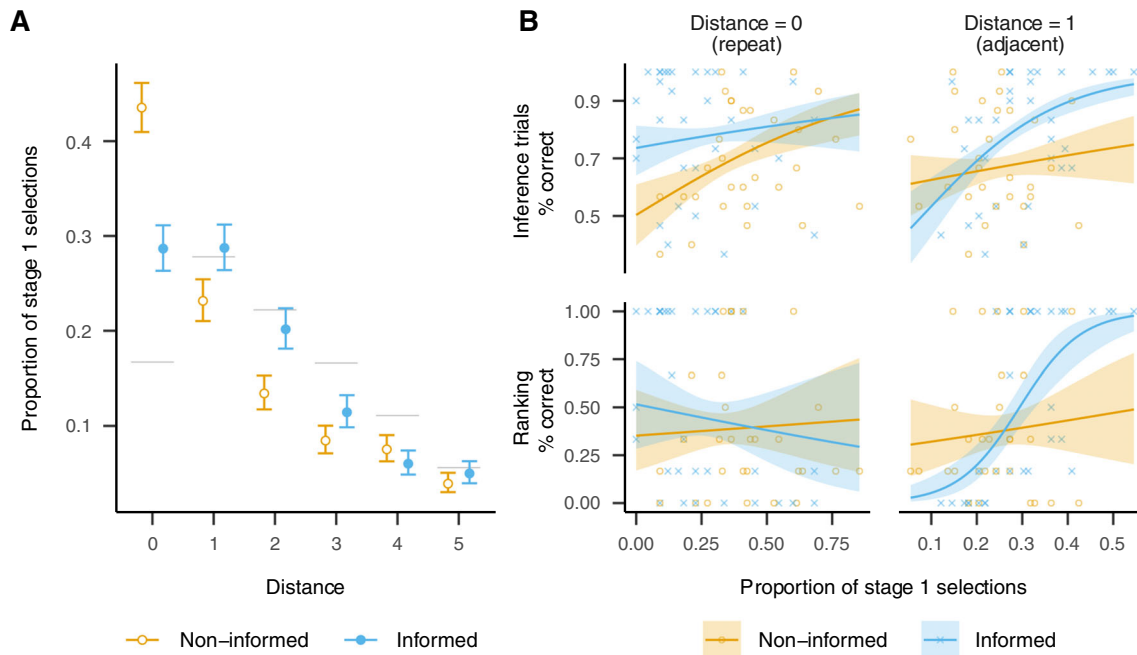


Fig. 4 **A**: Proportion of stage 1 selections in the Active condition at each absolute distance from item selected on previous study trial. Horizontal lines indicate the proportions at each distance expected from

are listed in Table 2. Among Non-informed participants, repeated selections appeared to aid relational learning, as the proportion of distance = 0 selections was positively related to inference accuracy (see Fig. 4B, left column). There were no other effects on performance among Non-informed participants at any distance. In contrast, a preference to select adjacent items was strongly predictive of relational learning among Informed participants. While selections at a distance of 0 and 2 had a negative impact on learning (specifically, greater number of blocks to criterion), the proportion of adjacent selections (distance = 1) was positively related to both inference accuracy and ranking accuracy (see Fig. 4B, right column). These results demonstrate that the impact of chained study was starkly different for active learners depending on their prior knowledge: Informed participants who frequently selected adjacent items also tended to exhibit highly accurate relational knowledge after training, whereas there was no evidence that chaining benefited Non-informed learners in the Active condition.

Discussion

It is well-established that organizing new experiences under a familiar relational schema supports rapid learning and generalization (Gilboa & Marlatte, 2017; Halford et al., 2010). Less clear is how people *discover* such abstract

relations in the first place, particularly in the absence of explicit instruction, hints, or salient relational features. This gap is exemplified by research on discrimination-based transitive inference in humans. People may learn the premises and make transitive inferences without an explicit understanding of the task structure, potentially by relying on implicit, associative mechanisms (Delius & Siemann, 1998; Dusek & Eichenbaum, 1997; Frank et al., 2005, 2006). Nevertheless, prior studies have consistently found that some learners become aware of the latent hierarchy over the course of training, and that this discovery may provide an immediate boost to relational learning (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005).

The present findings show that the makeup of the training experience—specifically, the order in which premises are encountered—is an important factor driving such relational discovery. Training in which learners frequently experienced overlapping premises in successive trials (the Passive-Adjacent condition) led to the highest test accuracy and post-task awareness among Non-informed participants. Passive-Adjacent training also produced the best overall performance in terms of ranking accuracy and confidence, providing further evidence that these participants became aware of the hierarchical organization of the items. Although similar training has been linked to improved inference in informed settings (Andrews, 2010; Halford, 1984; Markant, 2020; Waltz et al., 2004), this is the

random search. **B**: Effects of repeated selections (distance = 0, left column) and adjacent selections (distance = 1, right column) on inference accuracy (top row) ranking accuracy (bottom row) in the Active groups

Table 2 Associations between proportion of selections at each distance and other dependent measures in the Active training groups

Distance	Blocks to criterion	Recall accuracy	Inference accuracy	Ranking accuracy	Ranking confidence	Awareness
Non-informed						
0 (repeat)	-0.33 [-2.12, 1.45]	1.04 [0.51, 2.13]	1.59 [1.14, 2.21] **	1.09 [0.53, 2.26]	0.14 [-0.84, 1.12]	0.85 [0.29, 2.51]
1 (adjacent)	-0.61 [-2.10, 0.89]	1.09 [0.59, 2.01]	1.14 [0.86, 1.51]	1.18 [0.64, 2.15]	0.14 [-0.68, 0.96]	0.61 [0.23, 1.63]
2	-0.71 [-2.34, 0.93]	0.90 [0.46, 1.74]	1.30 [0.96, 1.77]	1.18 [0.61, 2.29]	0.25 [-0.65, 1.14]	0.64 [0.23, 1.82]
Informed						
0 (repeat)	2.23 [0.47, 3.99] **	1.56 [0.65, 3.73]	1.19 [0.81, 1.75]	0.79 [0.33, 1.93]	-0.15 [-1.11, 0.82]	1.12 [0.39, 3.21]
1 (adjacent)	0.47 [-0.80, 1.75]	1.18 [0.64, 2.17]	1.99 [1.44, 2.73] ***	4.56 [2.13, 9.74] ***	0.44 [-0.26, 1.14]	1.73 [0.79, 3.77]
2	1.45 [0.17, 2.73] *	0.92 [0.52, 1.63]	1.06 [0.79, 1.41]	1.06 [0.59, 1.93]	-0.01 [-0.72, 0.69]	1.15 [0.54, 2.43]

* $p < .05$; ** $p < .01$; *** $p < .001$, all adjusted for multiple comparisons. For recall accuracy, inference accuracy, ranking accuracy, and awareness, estimates are relative odds ratios (OR). For ranking confidence, estimates are linear effects

first demonstration that it facilitates the discovery of the hierarchy in naive learners.

There was also strong evidence for a link between spontaneous discovery and performance. Among Non-informed participants, Passive-Adjacent training was the only condition in which post-task awareness was consistently related to tests of relational learning. Past studies with similar assessments of awareness have produced mixed results concerning this relationship in non-instructed settings, as post-task awareness has been associated with faster learning and more accurate inference in some cases (Lazareva & Wasserman, 2010; Libben & Titone, 2008; Moses, Ostreicher, & Ryan, 2010; Smith & Squire, 2005) but not others (Frank et al., 2005; Greene et al., 2001; Kumaran & Ludwig, 2013; Siemann & Delius, 1996). The present results suggest that Passive-Adjacent training not only facilitated the discovery of the hierarchy, but that this awareness had rapid effects on the ability to learn the rank ordering of items during training.

Additional analyses showed that the impact of Passive-Adjacent training could not be explained by differences in how often premises were studied or their reinforcement histories, suggesting that it was the order of premises which led to stronger performance in that condition. Although further work is necessary to clarify how chaining leads to spontaneous discovery, it is likely that chained study highlights common relational structure across premises. By experiencing overlapping premises in short succession, learners may realize that many of the items share an abstract, relational feature: They are reinforced in the context of one item but not another, a pattern that can be explained by a rank-ordered organization which determines which item in any given pair is reinforced. Although still possible, this realization may be unlikely when overlapping premises are spaced apart during training, as tended to be the case in the Passive-Frequency condition. A similar finding was recently reported by Don, Goldwater, Greenaway, Hutchings, and Livesey (2020) using a set of related discrimination tasks (patterning and biconditional discrimination). Participants in that study were more likely to learn an abstract relational rule when training sequences alternated between discrimination sets that belonged to the same relational category. For example, in negative patterning, two cues have a positive outcome when presented independently (A+, B+), but a negative outcome when presented together (AB-). Don et al. found that clustered presentations of related discrimination sets (e.g., A+, B+, AB-, C+, D+, CD-) led to greater rule learning compared to random or blocked sequences. Taken together, these findings suggest that comparison across related problems leads to the explicit discovery of relational structure (Doumas et al., 2008; Gentner, 2010; Goldwater & Gentner, 2015), even in domains where such awareness may not be necessary for learning the correct responses to individual problems.

The present results also reiterate the powerful influence of prior knowledge on relational learning. Being informed about the hierarchy had widespread benefits, including fewer blocks to criterion during training, greater recall of studied premises, and higher accuracy on tests of relational inference and ranking. These effects are in line with similar instructional manipulations in past work (Greene et al., 2001; Lazareva & Wasserman, 2010; Libben & Titone, 2008), as well as studies in which the framing of the task signals a hierarchical organization (Kumaran & Ludwig, 2013; Moses et al., 2010). It should be noted, however, that Passive-Adjacent training also improved performance among Informed participants, reinforcing previous findings that chaining facilitates the integration of relational knowledge given a known schema (Andrews, 2010; Halford, 1984; Markant, 2020; Waltz et al., 2004). Chained study may therefore be a broadly effective approach for sequencing study when training involves premises with overlapping elements, whether or not learners have prior knowledge of how those premises fit into a larger conceptual structure.

Active control and prior knowledge

In contrast to the consistent effects of chained study, self-directed control over the order of premises (Active condition) was in most cases no better than random presentation (Passive-Frequency condition). This is a surprising finding in light of growing evidence that active control is beneficial for many forms of concept learning (Gureckis & Markant, 2012; Markant et al., 2016), including in a standard TI task in which participants were instructed to learn to rank individuals within a social hierarchy (Markant, 2020). A closer analysis of search behavior suggests that prior knowledge of the hierarchy had a striking impact on how active learners explored the hierarchy during training. Although overall test performance was lower than the Passive-Adjacent condition, Informed participants who tended to select adjacent items achieved high levels of inference and ranking accuracy (Fig. 4B), echoing recent evidence of a similar search preference in a standard (non-discriminative) TI task (Markant, 2020).

Interestingly, there was no corresponding relationship between chained study and relational knowledge among Non-informed participants. These participants were also less likely to select adjacent items, instead favoring the repeated selection of the same item in successive trials. Given that they were unaware of the latent hierarchy, Non-informed participants may have sought to mass study in order to master individual premises, a search behavior which was in fact associated with improved inference accuracy in this group. This finding highlights a potential

risk of active control during learning: It may lead to study strategies that are well-suited to near-term learning goals (e.g., memorization of the premises) while being less effective for the discovery of abstract, relational concepts. This lends some support to the argument that “pure discovery learning”—i.e., the freedom to explore without prior knowledge or familiarity with a domain—may be less effective than more direct forms of instruction (Mayer, 2004).

The overall disadvantage from Active training is also notable in light of work showing that active control improves memory for studied materials (Markant et al., 2016; Murty et al., 2015; Ruggeri et al., 2019; Voss et al., 2011). These choice-related enhancements may result from multiple mechanisms, including enriched encoding due to a sense of agency (Murty et al., 2015) or greater attention to the outcomes of self-generated actions (Markant, DuBrow, Davachi, & Gureckis, 2014a). The two-stage structure of the present task likely precluded many of these effects: Differences between active and passive training occurred solely at stage 1 when items were selected for study, while all participants made volitional stage 2 choices to generate the feedback that was crucial for encoding the premises. A more likely candidate for a choice-related enhancement would be any additional metacognitive processing required to make stage 1 choices. For example, if active participants test their memory for premises when deciding which item to select, this might speed learning of the premises or allow individuals to tailor their study to focus on the most uncertain items. There were some signs of strategic search in the Active condition, including a tendency to select items in the middle of the hierarchy with more variable reinforcement histories (Fig. 3A). However, it is unclear to what extent metacognitive monitoring guided those decisions or had broader effects on learning in the Active condition. A promising direction for future work is to consider how relational discovery is caused by this type of interplay between uncertainty monitoring and exploration of interrelated materials.

Limitations and future directions

An important limitation of the current study is that the timecourse of relational discovery is uncertain. For those Non-informed participants who discovered the hierarchy, it is unknown at what point during the task they became aware that the items could be organized into a linear ranking. For instance, it is possible that participants remained unaware of the hierarchy until they were presented with novel non-adjacent pairs during the test phase, which could have prompted further reasoning about the relations between premises. Reconstruction of the hierarchy at test is a plausible strategy for performing TI (Kumaran & McClelland, 2012), but it is unlikely to

account for the differences in performance between training conditions. Given that all Non-informed groups had high performance on recall trials during the test, it is unclear why Passive-Adjacent learners would be more likely at that point to discover the hierarchy and show immediate benefits in terms of inference accuracy. This points to an emergence of relational knowledge at some point during training.

A related factor to consider is how the potential for discovery depends on the amount of training. The current study provides one snapshot of performance based on a somewhat liberal learning criterion: Participants had to reach 100% accuracy in a single training block (with two repetitions of each premise) in order to end the training phase, but this criterion might be met without perfect knowledge of the premises (as shown by the errors made on some premises during the subsequent test phase). It is an open question how the results would change with a longer training phase or overlearning of the premises. Although overlearning might provide more opportunities to discover the hierarchy during training, it could also *inhibit* such discovery if people achieve high recall performance by simply relying on associative memory for the premises. If relational discovery is a byproduct of explicit reasoning, it may be most likely early in training before the individual premises are overlearned. Accordingly, chained study may be most likely to catalyze the discovery of relational structure at those points in training when learners are attempting to make sense of confusing or conflicting experiences.

Appendix

Post-task awareness questionnaire

Questions were adapted from a questionnaire used in past studies of awareness in transitive inference (Kumaran & Ludwig, 2013; Moses, Villate, Binns, Davidson, & Ryan, 2008). Responses consistent with awareness of the hierarchy or a logical ordering of the items were coded as aware (1); all other responses were coded as unaware (0). In the questions, “Phase 1” refers to the initial training phase whereas “Phase 2” refers to the final test phase which included both recall trials (premises experienced during study) and inference trials (novel pairings of non-adjacent items).

Q1: “Do you think there was a correct answer for all of the pairs that you experienced during Phase 2?”

- No (0)
- Not sure (0)
- Yes (1)

Q2: “In Phase 2 when you were presented with different pairs of cards, what reason did you have for choosing one as opposed to the other?”

- There is a logically correct choice. (1)
- One just seemed right but I can’t explain why. (0)
- I guessed. There may be a correct answer but I don’t know what it is. (0)
- I made a random choice because there is no correct choice. (0)

Q3: “What strategy (if any) did you use in Phase 1 to learn which card was correct in each pairing?”

- I tried to figure out the correct ordering of all the cards. (1)
- I memorized the right choice for each pair. (0)
- I just chose randomly and eventually got it. (0)
- No strategy. (0)
- Other. (0)

The following question also appeared on the awareness questionnaire but was not included in the analysis because it did not assess understanding of the relationship between items:

Q4: “In Phase 2 you were asked to make choices between pairs of cards. Were all of the pairs in Phase 2 the same as the pairs you had already experienced during Phase 1?”

- No.
- Yes.
- Not sure.

Screening questionnaire

A questionnaire was used to identify participants who were not engaged in the task, who were not fluent in English, and other invalid respondents such as bots which may affect data quality on Amazon Mechanical Turk (Chmielewski & Kucker, 2019). Following the training phase, participants were instructed to respond to a set of naturalistic decision making prompts (see examples below). For each situation, participants generated a possible course of action. Two research assistants coded responses based on whether they were meaningfully related to the prompt. Participants were excluded when both raters agreed that the generated course of action was not responsive to the prompt (proportion of agreement was 99%). Participants were not excluded based on the validity/feasibility of responses or writing quality (e.g., spelling or grammar errors).

Prompt text: Imagine you find yourself in the following situation. What is a potential course of action that you could take in this situation?

- On the morning of an important job interview, you wake up with fever and chills.
- You have procrastinated on starting a term paper and the due date is tomorrow.
- You stop by a store to grab an item you need. You are in a hurry but the cashier is nowhere to be seen.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.3758/s13421-021-01201-1>.

Acknowledgements The author thanks Shaina Glass and Mitra Mostafavi for their assistance with response coding.

Open Practices Statement The data for this study are available at the OSF repository: <https://osf.io/hnp2j/>.

References

- Andrews, G. (2010). Belief-based and analytic processing in transitive inference depends on premise integration difficulty. *Memory & Cognition*, *38*(7), 928–940.
- Birnbaum, M. S., Kornell, N., Bjork, E. L., & Bjork, R. A. (2013). Why interleaving enhances inductive learning: The roles of discrimination and retrieval. *Memory & Cognition*, *41*(3), 392–402.
- Bruner, J. (1961). The act of discovery. *Harvard Educational Review*, *31*(1), 21–32.
- Carvalho, P. F., & Goldstone, R. L. (2015). The benefits of interleaved and blocked study: Different tasks benefit from different schedules of study. *Psychonomic Bulletin & Review*, *22*(1), 281–288.
- Chmielewski, M., & Kucker, S. C. (2019). An MTurk crisis? Shifts in data quality and the impact on study results. *Social Psychological and Personality Science*, *1948550619875149*.
- Clark, H. H. (1969). Influence of language on solving three-term series problems. *Journal of Experimental Psychology*, *82*(2), 205.
- De Soto, C. B., London, M., & Handel, S. (1965). Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology*, *2*(4), 513.
- Delius, J. D., & Siemann, M. (1998). Transitive responding in animals and humans: Exaptation rather than adaptation? *Behavioural Processes*, *42*(2–3), 107–137.
- Don, H. J., Goldwater, M. B., Greenaway, J. K., Hutchings, R., & Livesey, E. J. (2020). Relational rule discovery in complex discrimination learning. *Journal of Experimental Psychology, Learning, Memory, and Cognition*.
- Doumas, L., Hummel, J., & Sandhofer, C. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115*(1), 1–43.
- 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.
- Elio, R., & Anderson, J. R. (1984). The effects of information order and learning mode on schema abstraction. *Memory & Cognition*, *12*(1), 20–30.
- Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723–7728.
- Frank, M. J., Rudy, J. W., Levy, W. B., & O'Reilly, R. C. (2005). When logic fails: Implicit transitive inference in humans. *Memory & Cognition*, *33*(4), 742–750.
- Frank, M. J., O'Reilly, R. C., & Curran, T. (2006). When memory fails, intuition reigns: Midazolam enhances implicit inference in humans. *Psychological Science*, *17*(8), 700–707.
- Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, *34*(5), 752–775.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*(1), 1–38.
- Gilboa, A., & Marlatte, H. (2017). Neurobiology of schemas and schema-mediated memory. *Trends in Cognitive Sciences*, *21*(8), 618–631.
- Goldwater, M. B., & Gentner, D. (2015). On the acquisition of abstract knowledge: Structural alignment and explication in learning causal system categories. *Cognition*, *137*, 137–153.
- Goldwater, M. B., Don, H. J., Krusche, M. J., & Livesey, E. J. (2018). Relational discovery in category learning. *Journal of Experimental Psychology: General*, *147*(1), 1.
- Greene, A. J., Spellman, B. A., Levy, W. B., Dusek, J. A., & Eichenbaum, H. B. (2001). Relational learning with and without awareness: Transitive inference using nonverbal stimuli in humans. *Memory & Cognition*, *29*(6), 893–902.
- Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, *7*(5), 464–481.
- Halford, G. S. (1984). Can young children integrate premises in transitivity and serial order tasks? *Cognitive Psychology*, *16*(1), 65–93.
- Halford, G. S., Wilson, W. H., & Phillips, S. (2010). Relational knowledge: The foundation of higher cognition. *Trends in Cognitive Sciences*, *14*(11), 497–505.
- Henriksson, M. P., & Enkvist, T. (2016). Learning from observation, feedback, and intervention in linear and nonlinear task environments. *The Quarterly Journal of Experimental Psychology*, *1–57*.
- Hothorn, T., Bretz, F., & Westfall, P. (2008). Simultaneous inference in general parametric models. *Biometrical Journal*, *50*(3), 346–363.
- Hummel, J. E., & Holyoak, K. J. (2001). A process model of human transitive inference. In *Spatial schemas in abstract thought* (pp. 279–305).
- 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.
- Kumaran, D., & Ludwig, H. (2013). Transitivity performance, relational hierarchy knowledge and awareness: Results of an instructional framing manipulation. *Hippocampus*, *23*(12), 1259–1268.
- Lazareva, O. F., & Wasserman, E. A. (2010). Nonverbal transitive inference: Effects of task and awareness on human performance. *Behavioural Processes*, *83*(1), 99–112.
- Lazareva, O. F. (2012). Transitive inference in nonhuman animals. In Wasserman, E. A., & Zentall, T. R. (Eds.) *The Oxford handbook of comparative cognition*: Oxford University Press.
- Lazareva, O. F., Gazes, R. P., Elkins, Z., & Hampton, R. (2020). Associative models fail to characterize transitive inference performance in rhesus monkeys (*Macaca mulatta*). *Learning & Behavior*.
- Leo, P. D., & Greene, A. J. (2008). Is awareness necessary for true inference? *Memory & Cognition*, *36*(6), 1079–1086.
- Libben, M., & Titone, D. (2008). The role of awareness and working memory in human transitive inference. *Behavioural Processes*, *77*(1), 43–54.
- Markant, D., DuBrow, S., Davachi, L., & Gureckis, T. M. (2014a). Deconstructing the effect of self-directed study on episodic memory. *Memory & Cognition*, *42*(8), 1211–1224.
- Markant, D., & Gureckis, T. M. (2014b). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, *143*(1), 94–122.
- Markant, D., Ruggeri, A., Gureckis, T. M., & Xu, F. (2016). Enhanced memory as a common effect of active learning. *Mind, Brain, and Education*, *10*(3), 142–152.

- Markant, D. (2020). Active transitive inference: When learner control facilitates integrative encoding. *Cognition*, 200.
- Mayer, R. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist*, 59(1), 14–19.
- McGonigle, B. O., & Chalmers, M. (1977). Are monkeys logical?. *Nature*, 267(5613), 694.
- Moses, S. N., Villate, C., Binns, M. A., Davidson, P. S., & Ryan, J. D. (2008). Cognitive integrity predicts transitive inference performance bias and success. *Neuropsychologia*, 46(5), 1314–1325.
- Moses, S. N., Ostreicher, M. L., & Ryan, J. D. (2010). Relational framework improves transitive inference across age groups. *Psychological Research PRPF*, 74(2), 207–218.
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *The Journal of Neuroscience*, 35(16), 6255–6264.
- Phillips, D. C. (1995). The good, the bad, and the ugly: The many faces of constructivism. *Educational Researcher*, 24(7), 5–12.
- R Core Team. (2018). *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing.
- Ruggeri, A., Markant, D., Gureckis, T., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, 186, 82–94.
- Sana, F., Yan, V. X., & Kim, J. A. (2017). Study sequence matters for the inductive learning of cognitive concepts. *Journal of Educational Psychology*, 109(1), 84–98.
- Siemann, M., & Delius, J. D. (1996). Influences of task concreteness upon transitive responding in humans. *Psychological Research Psychologische Forschung*, 59(2), 81–93.
- Smith, C., & Squire, L. R. (2005). Declarative memory, awareness, and transitive inference. *Journal of Neuroscience*, 25(44), 10138–10146.
- Sobel, D., & Kushnir, T. (2006). The importance of decision making in causal learning from interventions. *Memory & Cognition*, 34(2), 411.
- Vandierendonck, A., & De vooght, G. (1997). Working memory constraints on linear reasoning with spatial and temporal contents. *The Quarterly Journal of Experimental Psychology Section A*, 50(4), 803–820.
- Vasconcelos, M. (2008). Transitive inference in non-human animals: An empirical and theoretical analysis. *Behavioural Processes*, 78(3), 313–334.
- Voss, J., Gonsalves, B., Federmeier, K., Tranel, D., & Cohen, N. (2011). Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nature Neuroscience*, 14(1), 115–120.
- Waltz, J. A., Knowlton, B. J., Holyoak, K. J., Boone, K. B., Back-Madruga, C., McPherson, S., & Miller, B. L. (2004). Relational integration and executive function in Alzheimer's disease. *Neuropsychology*, 18(2), 296.
- Wynne, C. D. L. (1995). Reinforcement accounts for transitive inference performance. *Animal Learning & Behavior*, 23(2), 207–217.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.