

Effects of “chained” study on spontaneous relational discovery

Douglas B. Markant (dmarkant@uncc.edu)

Department of Psychological Science
University of North Carolina at Charlotte
9201 University City Blvd., Charlotte, NC 28223 USA

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 associations. 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. 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: relational learning; relational discovery; transitive inference

Relational knowledge guides learning and generalization in novel circumstances, as when preexisting schemas allow a learner to rapidly encode and react to new events (Gilboa & Marlatte, 2017). A fundamental problem in cognitive science is to understand how such abstract, relational knowledge emerges from experience with fragments of a larger conceptual structure. This problem is especially crucial in light of evidence that, in the absence of explicit instruction or salient relational cues, people often fail to recognize the relational structure that underlies the events they observe (Goldwater & Gentner, 2015; Goldwater, Don, Krusche, & Livesey, 2018).

The present study examines how the dynamics of learning affect whether people discover common relational structure across a set of interrelated experiences. I focus on the problem of 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., $A < B < C < D$) and participants learn the relations between adjacent *premise pairs* (e.g., $A < B$, $B < C$). They are then tested on both their memory for directly experienced premises (e.g., $A ? B$) and their ability to make transitive inferences given novel pairings that were never seen during training (e.g., $A ? C$). Previous studies of TI have found that, in the absence of explicit instruction, relatively few people spontaneously discover the

hierarchical organization of the items during training, but this explicit awareness is associated with improved relational inference. At present, the factors which lead to explicit awareness in TI are unknown. In the following I examine whether two aspects of training facilitate such discovery: 1) experiencing “chained” sequences of overlapping premises, and 2) active control over the selection of premises for study.

Explicit vs. implicit transitive inference

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* learning strategies that entail awareness of the hierarchy and the reasoning process itself, which may involve logical deduction (Clark, 1969) or an integrated mental representation of the hierarchy (Hummel & Holyoak, 2001). The explicit nature of these strategies is supported by their reliance on working memory (Libben & Titone, 2008) and sensitivity to relational complexity (Clark, 1969).

More recent research has demonstrated that explicit awareness of the hierarchy is not necessary for transitive inference (Vasconcelos, 2008). These studies have relied on *discrimination-based* tasks in which items lack any overt cues to the underlying hierarchical organization (Dusek & Eichenbaum, 1997; Frank, Rudy, Levy, & O’Reilly, 2005; Greene, Spellman, Levy, Dusek, & Eichenbaum, 2001; Libben & Titone, 2008; Smith & Squire, 2005). Premises are instead learned through trial-and-error, where reinforcement feedback is determined by an item’s rank in the latent hierarchy ($A + B-$, $B + C-$, $C + D-$, etc., with + following the item that is reinforced in each premise). After learning to select the reinforced item in each premise, people exhibit a capacity for TI despite being unaware of the underlying hierarchy. These findings have lent support to implicit accounts of TI based on associative learning in the absence of explicit, logical reasoning.

Although explicit awareness is not necessary for TI, it is typically associated with both faster learning of premises and more accurate inference. Performance in discrimination-

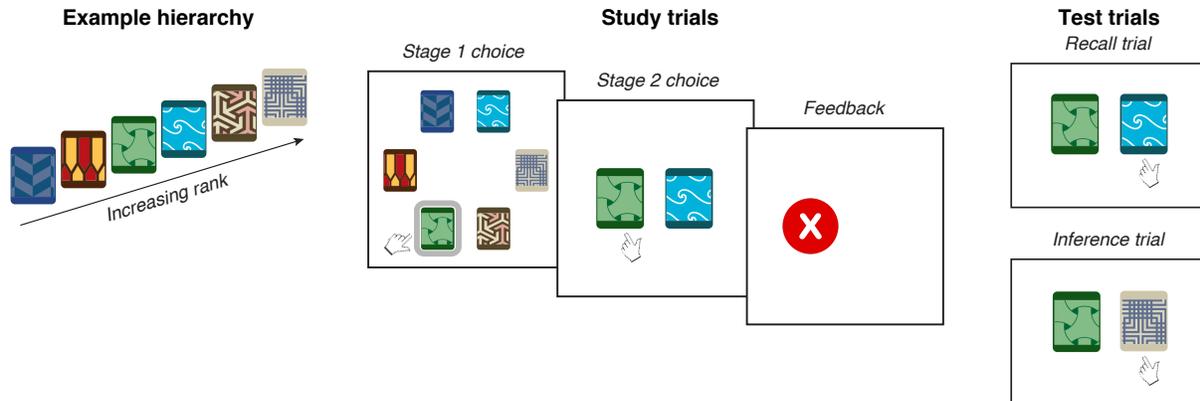


Figure 1: Depiction of transitive inference task. **Left:** Six cards were randomly arranged in a hierarchy that was unknown to participants. **Middle:** During training, participants learned about *premise pairs* composed of items that were adjacent in the hierarchy. In each study trial they first selected a card to learn about (*stage 1 choice*), which was randomly paired with an adjacent card. Participants then chose one card from the pair (*stage 2 choice*) to reveal whether a reward was hidden beneath it. **Right:** In each test trial, participants predicted 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.

based TI improves when participants are informed about the hierarchy prior to training (Greene et al., 2001; Libben & Titone, 2008) or when the framing of the task signals common relations across premises (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010). Without prior knowledge or cues to relational structure, some participants spontaneously *discover* the underlying hierarchy during discrimination-based TI, and this explicit awareness is similarly linked to faster learning and more accurate inference (Kumaran & Ludwig, 2013; Lazareva & Wasserman, 2010; Libben & Titone, 2008; Smith & Squire, 2005; although see Frank et al., 2005; Greene et al., 2001). However, serendipitous discovery is infrequent (typically occurring for a minority of participants) and little is known about when and why it occurs.

Chaining and spontaneous discovery

One factor which may contribute to relational discovery in TI is the order of premises during training. People may 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.). Chained sequences improve performance in the standard version of the task (Andrews, 2010; Waltz et al., 2004), but their effect on discrimination-based TI is less clear. Chained sequences may help the learner identify the common relational structure across multiple items (i.e., B is "better" than A, but "worse" than C; C is "better" than B, but "worse" than D). Indeed, this possibility has led researchers to minimize such overlap during training when they wish to focus on implicit forms of TI (Frank et al., 2005). However, no existing studies have directly tested whether chained sequences are more likely to result in explicit discovery of the hierarchy.

A second factor which may impact the discovery of relational structure is the opportunity to self-generate the train-

ing sequence. A recent study found that active control over the selection of premises during training led to improved performance in a standard TI task (Markant, 2019, 2020), largely due to participants' preference to chain premises during study. It is unknown, however, whether learners exhibit the same study behavior in discrimination-based TI or if active control facilitates relational discovery in the absence of prior knowledge of the hierarchy.

The present study examined the effects of chaining and learner control on the discovery of relational structure. A novel discrimination-based TI task (Figure 1) allowed participants to exert control over the selection of premises during training. Participants played a card game 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. Each training trial began with a *stage 1 choice* 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. In the Active training condition, participants were free to select any of the items in the hierarchy for study. There were two passive training conditions in which the training sequences were predetermined. In the Passive-Frequency condition, items were selected according to how often they had been experienced during training. In the Passive-Overlap condition, items were selected so as to create chained sequences of overlapping premises from trial to trial.

Participants were tested on their recall of studied premise pairs, their ability to make transitive inferences, and their explicit post-task awareness of the hierarchy. To better evaluate any changes in awareness (and associated differences in performance; see e.g., Greene et al., 2001), participants'

prior knowledge of the hierarchy 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 explicit awareness and a comparison of the effects of training condition when directly instructed about the hierarchy.

Experiment

Participants

$N = 252$ people were recruited from Amazon Mechanical Turk. Forty-seven people were excluded because they failed attention screening ($N = 24$) or instruction comprehension questions ($N = 23$), leaving $N = 205$ participants (age $M = 36.66$ years, $SD = 10.25$; 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 21.39 minutes on average ($SD = 8.02$).

Materials and Procedure

Participants learned about a 6-item hierarchy made up of cards with unique graphical patterns (Figure 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 consistent perceptual features which might serve as a cue to the cards’ 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-Overlap, or Active) as between-subjects factors.

Instructional manipulation. The between-subjects instructional manipulation occurred at the start of the task. Participants in the Non-informed condition saw the following:

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

In contrast, participants in the Informed condition were told there was an underlying hierarchy:

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 will depend on whether you learn the correct ranking of the six cards.

A quiz followed the instructions in which Informed participants had to verify they understood the hierarchical relationship. There were no further differences between the Informed and Non-informed conditions.

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 with the six items displayed in a circle (Figure 1, middle). Participants selected an item for study (stage 1 choice) according to their training condition (between-subjects):

- *Passive-Frequency condition.* A predetermined item was highlighted and participants were instructed to select it. The selected item was randomly sampled from the set of items that had been studied the least often up to that point. This condition produced sequences in which all items were selected with equal frequency by the end of training and selections of any given item tended to be spaced apart.
- *Passive-Overlap condition.* A predetermined item was highlighted and participants were instructed to select 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. As a result, training sequences in this condition featured a high proportion of chained premises in successive trials.
- *Active condition.* 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 an adjacent item to form a premise. All participants then chose one of the two items (stage 2 choice) and received feedback about whether it was hiding a reward (Figure 1, middle).

After 12 study trials, participants completed 10 recall trials (Figure 1, right), with two trials for each premise pair. On each recall trial a premise was presented in the center of the display and participants were instructed to pick 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 correct card for every premise pair twice.

Test phase. The test phase was comprised of 45 trials, with three repetitions of every pairing of items from the hierarchy. *Recall trials* involved premises that were experienced during the study phase, whereas *inference trials* involved novel pairings of non-adjacent items. In each test trial, two items were presented side-by-side in the center of the display. As with 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 during the test phase.

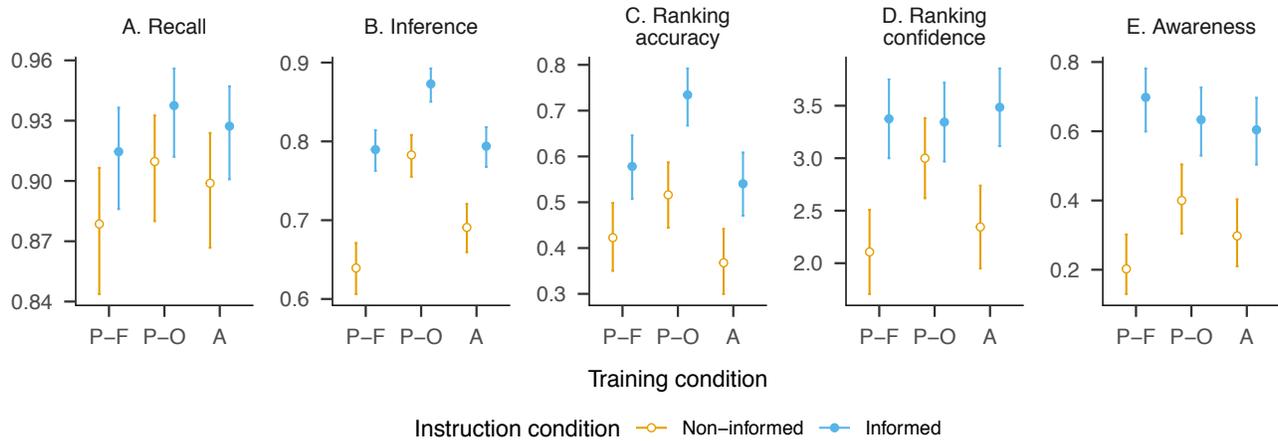


Figure 2: Dependent measures by condition. Error bars indicate 95% confidence intervals.

Post-task awareness questionnaire. Immediately following the test phase, participants were assessed for their explicit awareness of the hierarchy using questions adapted from Kumaran and Ludwig (2013). 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):

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 *awareness score* was the proportion of responses (out of three) which indicated awareness of the hierarchical organization of the items or the use of logical reasoning to draw inferences during the test.

Ranking elicitation. Finally, participants were instructed to rank the six items 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. Ranking accuracy was the proportion of items that

were in the correct position. Participants then rated their confidence that they ranked the items correctly on a scale from 1 (not at all confident) to 5 (completely confident).

Results

Test accuracy. Test responses were scored according to whether participants chose the higher-ranked item in each test pair (0 = incorrect, 1 = correct) and the proportion of correct responses was modeled using logistic regression. Recall accuracy (Figure 2A) 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 ($\chi^2(2) = 4.16, p = 0.12$) or interaction ($\chi^2(2) = 0.01, p = 0.99$). Pairwise comparisons were all non-significant.

There was a main effect of instructional condition on inference accuracy (Figure 2B) ($\chi^2(1) = 101.58, p < .001$), with accuracy higher among Informed participants than Non-informed in each training condition (all $p < .001$). In addition, there was an effect of training condition on inference accuracy ($\chi^2(2) = 74.47, p < .001$), but no interaction ($\chi^2(2) = 0.01, p = 0.99$). Within the Non-informed condition, inference accuracy was higher in the Passive-Overlap condition than both the Passive-Frequency ($OR = 2.03 [1.52, 2.72], z = 6.62, p < .001$) and Active conditions ($OR = 0.62 [0.46, 0.83], z = -4.42, p < .001$), whereas accuracy did not differ between the Active and Passive-Frequency conditions ($OR = 1.26 [0.95, 1.67], z = 2.26, p = 0.16$). Similarly, within the Informed condition, Passive-Overlap inference accuracy was higher than both the Passive-Frequency ($OR = 1.83 [1.30, 2.57], z = 4.83, p < .001$) and Active conditions ($OR = 0.56 [0.40, 0.79], z = -4.63, p < .001$) but there was no difference between Active and Passive-Frequency conditions ($OR = 1.03 [0.76, 1.39], z = 0.24, p = 1.00$).

Ranking accuracy and confidence. There were main effects of both instructional condition ($\chi^2(1) = 38.08, p < .001$) and training condition ($\chi^2(2) = 24.85, p < .001$) on ranking accuracy (Figure 2C), but no interaction ($\chi^2(2) = 1.23,$

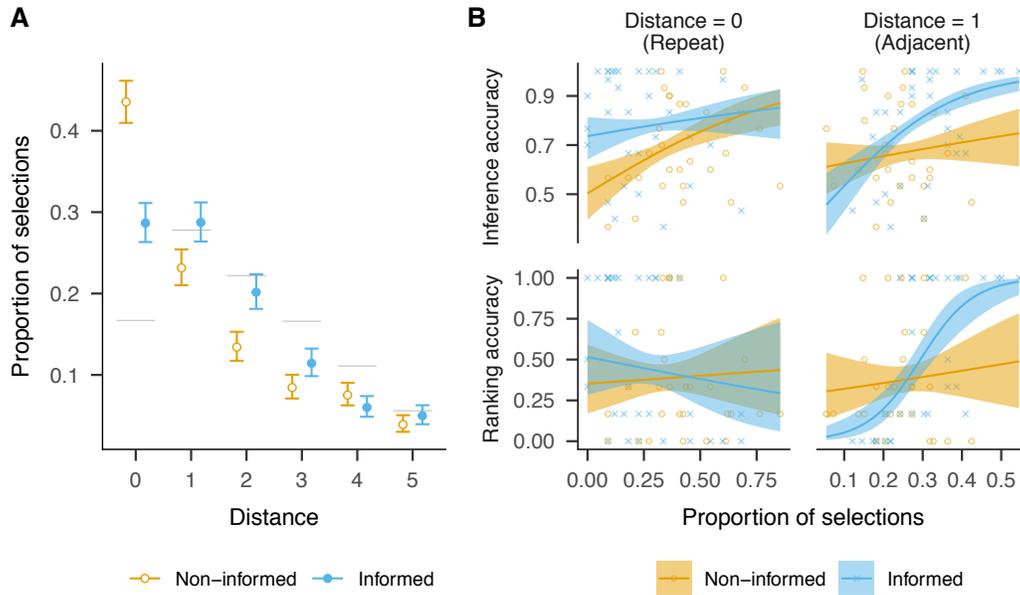


Figure 3: **A:** Proportion of selections (stage 1 choices) in the Active condition at each absolute distance from the item selected on previous trial. Horizontal lines indicate the proportions expected from random search. **B:** Effects of repeated selections (left column) and adjacent selections (right column) on inference (top row) ranking accuracy (bottom row) in the Active condition.

$p = 0.54$). Ranking accuracy was higher among Informed participants than Non-informed in each training condition (Passive-Frequency: $OR = 1.87$ [1.05, 3.35], $z = 2.93$, $p = 0.03$; Passive-Overlap: $OR = 2.59$ [1.43, 4.70], $z = 4.34$, $p < .001$; Active: $OR = 2.02$ [1.14, 3.60], $z = 3.31$, $p = 0.008$). Within the Informed condition, Passive-Overlap ranking accuracy was higher than both Passive-Frequency ($OR = 2.02$ [1.11, 3.66], $z = 3.20$, $p = 0.01$) and Active conditions ($OR = 0.43$ [0.24, 0.77], $z = -3.94$, $p < .001$), while there was no difference between Active and Passive-Frequency conditions ($OR = 0.86$ [0.49, 1.49], $z = -0.75$, $p = 0.96$). Within the Non-informed condition, ranking accuracy was lower following Active training than Passive-Overlap training ($OR = 0.55$ [0.30, 0.98], $z = -2.82$, $p = 0.04$) but there were no other pairwise differences.

A similar pattern was seen in confidence judgments of elicited rankings (Figure 2D). There was no effect of training condition ($F(2, 179) = 2.02$, $p = .135$), but there was an effect of instructional condition ($F(1, 179) = 32.49$, $p < .001$) and a significant interaction ($F(2, 179) = 3.31$, $p = .039$). Confidence was higher when Informed for both Passive-Frequency ($\beta = 1.27$ [0.50, 2.04], $z = 4.52$, $p < .001$) and Active training conditions ($\beta = 1.14$ [0.38, 1.90], $z = 4.13$, $p < .001$). For Passive-Overlap training only, there was no difference in confidence between Informed and Non-informed groups ($\beta = 0.34$ [-0.41, 1.09], $z = 1.26$, $p = 0.75$). For Non-informed participants, confidence was higher following Passive-Overlap training compared to Passive-Frequency training ($\beta = 0.89$ [0.12, 1.67], $z = 3.16$, $p = 0.01$).

Post-task awareness. Responses to the awareness questions were coded based on whether participants indicated explicit awareness of the hierarchy. The probability of making “aware” responses was modeled using logistic regression (0 = unaware, 1 = aware) based on three responses for each person. Awareness was higher among Informed participants in all three training conditions (all $p < .001$), indicating that the instructional manipulation had the intended effect of increasing explicit awareness of the hierarchy.

There was no effect of training condition ($\chi^2(2) = 2.27$, $p = 0.32$), but there was a significant interaction ($\chi^2(2) = 7.85$, $p = 0.02$). Within the Non-informed condition, awareness was higher following Passive-Overlap training than Passive-Frequency training ($OR = 2.63$ [1.03, 6.72], $z = 2.79$, $p = 0.04$). No other pairwise comparisons were significant. Thus, in addition to the effects on relational learning, Passive-Overlap training was associated with the highest post-task awareness among Non-informed participants.

Selections during training

The results indicate that Active participants performed worse in tests of relational learning compared to Passive-Overlap training. The final analyses examined the search behavior (stage 1 choices) of Active participants and how it related to performance. On every study trial Active participants were free to select any item from the hierarchy. Selections were classified by their absolute distance to the item selected on the previous trial. 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. Figure 3A shows the proportion of selections at each distance, with the proportions expected from random search marked by horizontal lines.

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 ($OR = 1.34 [1.13, 1.59]$, $z = 3.35$, $p < .001$) 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. Thus, prior knowledge of the hierarchy was associated with a preference to learn about items that were near to the item chosen in the last trial, including adjacent (distance = 1) selections that could produce overlapping premises in successive trials.

Finally, I examined how these search behaviors were related to learning and post-task awareness. Regression models for each dependent variable included terms for the proportion of selections at distances of 0, 1 and 2 in each condition. Among Non-informed participants, repeated selections (distance = 0) were associated with improved inference accuracy ($OR = 1.59 [1.14, 2.21]$, $z = 3.64$, $p = .001$) (Figure 3B, top left). The proportion of adjacent selections (distance = 1) was unrelated to relational learning in the Non-informed condition. In contrast, among Informed participants the proportion of adjacent selections (distance = 1) was strongly positively related to both inference accuracy ($OR = 1.99 [1.45, 2.73]$, $z = 5.57$, $p < .001$) and ranking accuracy ($OR = 1.56 [2.13, 9.74]$, $z = 5.17$, $p < .001$) (see Figure 3B, right column). Selections at any distance were not significantly related to post-task awareness (all $p > .29$).

Discussion

Previous studies of discrimination-based TI found that, even without the aid of explicit instruction or salient relational cues, some learners become aware of the hierarchical organization of the items. The present findings show that the sequence of events during training is an important factor which drives this relational discovery. The Passive-Overlap condition, in which training included a high proportion of chained premises in successive trials, led to the highest accuracy and post-task awareness among Non-informed participants. Although similar training has been linked to improved inference in informed settings (Andrews, 2010; Markant, 2020; Waltz et al., 2004), this is the first demonstration that it is more likely to cause explicit awareness of the hierarchy in naïve learners. Chained study likely highlights common structure across multiple items (i.e., the presence of opposing contingencies), a process that is consistent with the broader claim that comparison spurs relational discovery (Doumas, Hummel, & Sandhofer, 2008; Goldwater & Gentner, 2015).

The results also illustrate the powerful influence of prior knowledge on relational learning, as being informed of the

hierarchy led to higher accuracy on tests of recall, transitive inference, and ranking. These effects are consistent with studies that have used similar instructional manipulations (Greene et al., 2001; Lazareva & Wasserman, 2010; Libben & Titone, 2008) or materials which naturally evoke a hierarchical organization (Kumaran & Ludwig, 2013). More broadly, the results provide further evidence of the facilitative effects of schematic knowledge on learning (Gilboa & Marlatte, 2017). It should be noted, however, that Passive-Overlap training also improved performance among Informed participants, indicating that chaining facilitates the integration of relational knowledge given a preexisting schema (Andrews, 2010; Markant, 2019, 2020; Waltz et al., 2004).

The importance of chaining was further demonstrated by Informed participants who actively selected premises during study. Although overall performance was lower than the Passive-Overlap condition, Informed participants who tended to select adjacent items achieved high levels of inference and ranking accuracy (Figure 3B), echoing recent evidence of a similar search preference in a standard (non-discriminative) TI task (Markant, 2020). In contrast, there was no such relationship between the proportion of adjacent selections and performance among Non-informed participants. This group was also less likely to chain premises, 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. This search behavior was in fact associated with improved inference accuracy but had no effects on post-task awareness or the ability to rank items, reinforcing the idea that this group relied on an implicit form of TI (Frank et al., 2005; Greene et al., 2001; Smith & Squire, 2005). For naïve learners, active control may lead to study strategies which help them reach immediate learning goals (e.g., memorizing the premises) and support associative inference, while being less effective for the explicit discovery of abstract, relational concepts.

The relative disadvantage from Active training is also notable in light of work showing that active control improves memory through a number of mechanisms (Markant, Ruggeri, Gureckis, & Xu, 2016; Ruggeri, Markant, Gureckis, & Xu, 2019; Murty, DuBrow, & Davachi, 2015; Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). For instance, a sense of agency alone may augment memory formation, as perceived control over study enhances recognition of items delivered in a predetermined sequence (Murty et al., 2015). The two-stage structure of the present task likely obscured any such enhancement: 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. Alternatively, active choice during stage 1 might enhance memory if it involves metacognitive processing (e.g., judgments of learning) that is absent during passive training. Although there were signs of strategic search

based on items' relationship to previously studied premises (Figure 3A), it is unclear to what extent metacognitive monitoring guided those decisions or had broader effects on memory formation. In general, these findings suggest that control over the order in which interrelated materials are studied may produce inconsistent outcomes depending on a learner's prior knowledge and search strategy. In this context, chained study (whether self-generated or directed by an external source) may be most effective for supporting transitive inference and catalyzing relational discovery.

References

- Andrews, G. (2010). Belief-based and analytic processing in transitive inference depends on premise integration difficulty. *Memory & Cognition*, *38*(7), 928–940.
- Clark, H. H. (1969). Influence of language on solving three-term series problems. *Journal of Experimental Psychology*, *82*(2), 205.
- 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.
- 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.
- Gilboa, A., & Marlatte, H. (2017). Neurobiology of schemas and schema-mediated memory. *Trends in Cognitive Sciences*, *21*(8), 618–631.
- 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.
- 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.
- 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.
- 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., & 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.
- Libben, M., & Titone, D. (2008). The role of awareness and working memory in human transitive inference. *Behavioural Processes*, *77*(1), 43–54.
- Markant, D. (2019). Navigating the chain of command: Enhanced integrative encoding through active control of study. In V. Goel, C. Seifert, & C. Freska (Eds.), *Proceedings of the 41st Annual Meeting of the Cognitive Science Society* (pp. 770–775). Austin, TX: Cognitive Science Society.
- Markant, D. (2020). Active transitive inference: When learner control facilitates integrative encoding. *Cognition*, *200*.
- 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.
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *The Journal of Neuroscience*, *35*(16), 6255–6264.
- Ruggeri, A., Markant, D., Gureckis, T., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, *186*, 82–94.
- Smith, C., & Squire, L. R. (2005). Declarative memory, awareness, and transitive inference. *Journal of Neuroscience*, *25*(44), 10138–10146.
- 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.