Active transitive inference: When learner control facilitates integrative encoding

Douglas B. Markant¹

¹ Department of Psychological Science, University of North Carolina at Charlotte

Author Note

Douglas B. Markant, Department of Psychological Science, University of North Carolina at Charlotte.

Word count: 13490

Declaration of competing interests: None

Correspondence concerning this article should be addressed to Douglas B. Markant, Department of Psychological Science, Colvard South Building, UNC Charlotte, 9201 University City Blvd., Charlotte, NC 28223. E-mail: dmarkant@uncc.edu

Abstract

A growing body of research indicates that active control of learning improves episodic memory for material experienced during study. It is less clear how active learning impacts the integration of those experiences into flexible, generalizable knowledge. This study presents a novel active transitive inference task to investigate how people learn a relational hierarchy through active selection of premise pairs. Active control improved memory for studied premises as well as transitive inferences involving items that were never experienced together during study. Active learners also exhibited a systematic search preference, generating sequences of overlapping premises that may facilitate relational integration. Critically, however, advantages from active control were not universal: Only participants with higher working memory capacity benefited from the opportunity to select premise pairs during learning. These findings suggest that active control enhances integrative encoding of studied material, but only among individuals with sufficient cognitive resources.

Keywords: active learning, transitive inference, relational learning, memory, information search

Active transitive inference: When learner control facilitates integrative encoding

A central function of cognition is the formation of generalizable knowledge from discrete experiences. Memory not only enables recollection of independent events from the past, but also lets us link those events together to identify common elements or structure. For example, if a child remembers that two animals, seen at different points in time, were given the same label (e.g., “look at the *DOG!*”), that common association provides the foundation for generalizable knowledge about a category of objects and its typical features.

Modern educational philosophies such as constructivism (Phillips, 1995) recognize the learner to be the locus of such knowledge construction: conceptual understanding is built from the ingredients of direct experience through cognitive processing. Constructivism was a response to traditional pedagogy which emphasized rote memorization of materials delivered by a teacher with little opportunity for learner-driven elaboration or conceptual integration (Dewey, 1938). Although constructivists may acknowledge that rote learning has a place in education (Von Glasersfeld, 2012), they argue that conceptual knowledge cannot be delivered whole via memorization. Instead, each learner must assemble a new concept by processing their experiences in relation to each other and existing prior knowledge.

Learner-centric approaches to instruction, including constructivism and its close relatives, share this assumption that knowledge formation occurs through internal cognitive processing. At the same time, advocates of these approaches have commonly argued that *external*, overt control over the learning experience—expressed through behaviors such as physical interaction, information search, and exploration—are integral to knowledge construction (Bruner, 1961; Dewey, 1929; Phillips, 1995). In contrast to the former, seemingly uncontroversial claim, this latter view has not been without criticism. Debate continues over whether overt learner control, as in instruction that prioritizes learner-driven discovery and exploration, is aligned with the goal of building conceptual knowledge (Kirschner, Sweller, & Clark, 2006; Mayer, 2004; Tobias & Duffy, 2009). A recent review of classroom interventions suggests that moderate amounts of control lead to better learning outcomes than either direct instruction or unassisted discovery (Alfieri, Brooks, Aldrich, & Tenenbaum, 2011). However, the diversity of instructional methods and outcomes in this literature

makes it difficult to draw clear conclusions about when and why learner control is advantageous. There is a need to better understand how overt control over study interacts with basic cognitive processes involved in memory and knowledge formation (Bernstein, 2018; T. M. Gureckis & Markant, 2012).

The present study addresses this need by examining the impact of active control in relation to the fundamental distinction between rote memory for experienced materials and the construction of integrated knowledge from that experience. I focus on a simple, well-known example of relational generalization from memory: transitive inference (TI). In TI people learn about an ordered hierarchy (e.g., $A < B < C$) by studying premises comprised of adjacent items (e.g., $A < B$, $B < C$). They are then tested on their memory for the ordering of studied pairs (*recall trials*; e.g., $A ? B$) and their ability to infer relationships between items that were never experienced together (*inference trials*; e.g., $A ? C$). TI has been the subject of a wealth of past research, but it has always been examined under passive conditions in which control is absent. This study introduces a novel, active TI task in which participants choose which premises to study during learning. In a passive control condition participants experience the same training procedure but make no study decisions. Active and passive study are compared in terms of both recall of directly experienced premises and the ability to draw inferences that require their relational integration. This comparison provides new insight into how overt control during study affects the construction of generalizable knowledge.

Effects of active control on memory and concept learning

A growing body of experimental research suggests that exerting control over study benefits learning and memory. Many of these studies have relied on *yoked* experiments in which the information observed during passive study is matched to the experience actively selected by other learners. In the domain of episodic memory, control over the selection of information leads to enhanced recognition memory compared to yoked observation of the same study sequences (Markant, DuBrow, Davachi, & Gureckis, 2014; Voss, Gonsalves, Federmeier, Tranel, & Cohen, 2011). Similar effects have been found across a range of materials and procedures, suggesting that control over the study experience tends to improve later memory of it. This enhancement might arise from a number of mechanisms depending on the kind of control afforded to the learner,

including improved coordination of attention with the flow of experience, additional metacognitive processing, or enriched encoding processes that result from volitional choice (for a review, see Markant, Ruggeri, Gureckis, and Xu (2016)).

A separate line of research has found that active control also leads to more efficient acquisition of abstract concepts, including causal relationships (Lagnado & Sloman, 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003) and categories (Markant & Gureckis, 2014). For instance, Markant and Gureckis (2014) found that control over the selection of training exemplars led to faster learning of unidimensional category boundaries compared to yoked observation of the same training sequences. Critically, however, these studies were not designed to distinguish between the acquisition of abstract, conceptual knowledge and memory for directly experienced training items, leaving open the question of what mechanism causes those advantages from active control. Do they reflect better memory for experienced information itself, which then supports generalization at the time of test? Or does active control change how information is processed during learning, enhancing the formation of relational knowledge that abstracts away from the study experience? Following Zeithamova, Schlichting, and Preston (2012), these alternatives can be conceptualized as two types of memory formation during study: *elemental encoding* of stimuli or associations that are directly experienced, and *integrative encoding* through which disparate study episodes are linked together into a unified representation. Whereas existing research has established that active control enhances elemental encoding in a variety of contexts, its impact on integrative encoding remains unclear.

Elemental vs. integrative encoding in transitive inference

Transitive inference is well-suited to examine whether active control has broader benefits for knowledge construction beyond enriched elemental encoding of studied items. Performing TI from memory involves comparing two items that have never been experienced together (e.g., $A ? C$) but are related by two or more overlapping premise pairs that were directly studied (e.g., $A < B$, $B < C$). It has long been recognized that multiple mechanisms can support TI, which for the present purposes can be distinguished by their reliance on either elemental or integrative encoding. Elemental encoding-based inference occurs by retrieving studied premises at the time of test and

reasoning across overlapping relations (Kumaran & McClelland, 2012). In this case, successful inference hinges on robust encoding of pairs during study to ensure later retrieval. This process implies that test performance will be best for direct recall of studied pairs (e.g., $A \geq B$) and will decline as the distance between test items increases (e.g., $A \geq C$, $A \geq D$, etc.). Larger distances between test items are predicted to lengthen response times since a greater number of intervening pairs must be traversed, and decrease accuracy since there are more opportunities for retrieval errors along the way.

In contrast, integrative encoding-based accounts of TI postulate the formation of a unified representation of the hierarchy during study (Dusek & Eichenbaum, 1997; Hummel & Holyoak, 2001; Shohamy & Wagner, 2008; Zeithamova & Preston, 2010). By combining information from premise pairs as they are experienced, people induce a spatial (De Soto, London, & Handel, 1965; Huttenlocher, 1968) or propositional (Hummel & Holyoak, 2001; Trabasso, Riley, & Wilson, 1975) representation by mapping items onto an integrated ordinal dimension (e.g., $A < B < C < D < E \dots$). Inference then simply entails comparing the positions of any two items along that dimension. Integrative encoding is associated with a different pattern of test performance from elemental encoding: Accuracy should *increase* (and response time decrease) with distance, as items that are further apart on that latent dimension are easier to distinguish. Such *symbolic distance effects* (SDEs) are a hallmark of integrative encoding (Acuna, Sanes, & Donoghue, 2002; Moyer & Landauer, 1967).

Integrative encoding typically leads to better test performance than other strategies in TI, but at the cost of explicit, effortful reasoning during study (Acuna et al., 2002; Moses, Villate, Binns, Davidson, & Ryan, 2008). Integrative encoding is more likely when participants are aware there is an underlying hierarchy to be learned (Lazareva & Wasserman, 2010; Moses, Villate, & Ryan, 2006). Moreover, when people rely on integrative encoding their performance is correlated with working memory capacity (WMC) (Libben & Titone, 2008). Thus, although constructing an integrated representation during study leads to superior generalization, it also depends on explicit awareness and incurs larger cognitive costs, factors which may have important implications for the effects of active control on TI. If exerting control over the selection of information imposes

Figure 1. Transitive inference task. Participants learned about the "chain of command" at two 9-person companies, each made up of all men or all women (only one example shown here). In each learning trial two individuals were presented as options. One person was selected either through free choice (active condition) or predetermined choice (passive condition) to observe their direct supervisor. In each test trial participants decided which of two individuals was ranked higher in the company, with three types of trials that differed in the distance between items in the hierarchy.

additional processing costs (as has been suggested by critics of learner-driven instruction; see Kirschner et al. (2006)), individual differences in cognitive function (including WMC) may limit any potential advantage from the opportunity to control study.

Overview of the present study

Two experiments were conducted using a novel transitive inference task to investigate how active control impacts the construction of relational knowledge. The participants' goal in the task was to learn the "chain of command" at a set of 9-person companies, where each item in the hierarchy was an individual represented by a face image (see Figure 1). Each premise comprised an employee and their direct supervisor (the immediately superordinate item). In the active study condition, participants controlled the selection of premises by choosing a person to learn about their supervisor on each trial. In the passive condition all selections were predetermined. The test procedure involved judging for any possible pairing which person was higher in the corporate

hierarchy. Test performance is analyzed as a function of the distance between items in the hierarchy: adjacent items made up studied premise pairs that could be directly recalled whereas non-adjacent items required transitive inference. Because elemental and integrative encoding predict distinct relationships between distance and test performance, this design makes it possible to examine whether learner control affects how materials are processed. In particular, the presence of symbolic distance effects (SDEs), such that accuracy increases (or RT decreases) with greater distances between items, is taken as evidence for integrative encoding during study (see the General Discussion for discussion of alternative accounts for SDEs).

Active control was predicted to improve accuracy on inference trials, but this advantage might arise from enhancements to either elemental encoding or integrative encoding. Enhanced elemental encoding implies that active control improves memory for premise pairs themselves, in line with prior findings of better recognition of experienced materials following active study (Markant et al., 2016). A key reason to expect such an effect is that making selection decisions may prompt metacognitive processing about the premises that is absent during passive observation. For example, given a choice to learn about one of two items from the hierarchy, the learner may try to recall the superordinate item (i.e., the person’s direct supervisor) for each option. Enhanced memory for the premise pairs could then result either from the effects of retrieval practice (Roediger & Karpicke, 2006) or because learners allocate study effort to those pairs that are less well-learned (Metcalfe, 2009).

A related set of processes might underlie enhanced integrative encoding, with the key difference that active control promotes relational processing across premise pairs during study. For example, when choosing which option to study, a learner might attempt to relate each option to their existing knowledge of the hierarchy, including other premises that have already been experienced. Like retrieval practice, such elaborative processing involved in making a selection could on its own contribute to integrative encoding. Alternatively, it may allow the learner to generate study sequences that aid relational integration. For instance, passive training that involves successive presentation of overlapping premise pairs (e.g., $A < B$, $B < C$, $C < D$, ...) leads to faster learning than random sequences of pairs (Halford, 1984; Waltz et al., 2004). If studying

nearby premises helps learners to construct a representation of the hierarchy, active learners may benefit from choosing such options for study. The TI task was designed to explore whether people exhibit such a preference, with each selection involving a choice between a *near* and *far* option which differed in their distance from the pair studied on the previous trial (see Figure 4A).

To summarize, the task was designed to evaluate whether active control improves both recall and transitive inference compared to passive observation. The relationship between inferential distance and performance provides insight into whether any such advantage reflects enhanced elemental or integrative encoding. Experiment 1 examines the difference between active control and passive observation of randomly generated training sequences, both at an immediate test and after a week delay. In Experiment 2, active control is compared to yoked observation of training sequences generated by participants in Experiment 1. This provides a test of whether the effects of active control arise from the process of making study decisions (e.g., due to additional metacognitive or elaborative processing) or the consequences of those decisions in the form of more effective training sequences. Finally, to evaluate whether individual differences in cognitive resources affects the ability to construct relational knowledge through active selection, an assessment of WMC (operation span) was performed in each experiment and included as a covariate in the analyses.

Experiment 1

The methods described in this and the following experiment were approved by the Institutional Review Board at UNC Charlotte (IRB #17-0405).

Participants

$N = 100$ participants (60 women; age: $M = 21.94$ years, $SD = 5.60$) were recruited from the student population at UNC Charlotte. The sample size was chosen prior to the experiment based on a target of 25 participants for each of the four counterbalancing conditions (see below). Participants received either course credit or \$8 (\$4 per session) as compensation for participating in the study. All participants received an incentive payment based on their performance in the first test session ranging from \$0 ($< 50\%$ correct) to \$5 (90–100% correct). Payments were made in the form of Amazon gift cards. $N = 62$ participants (62%) returned for the second session.

Materials

Face stimuli were obtained from the 10k US Adult Faces Database (Bainbridge, Isola, & Oliva, 2013), a collection of pictures from Google Images designed to be a representative sample of the US adult population. The database includes subjective ratings of each face on a number of attributes, including judgments of perceived age, emotional affect, and memorability. For each sex, the stimulus set was filtered to include only faces that were non-famous and which had mean ratings within a 1-point interval centered on the midpoint of the rating scale for perceived age, emotional affect, and memorability. Thirty-six images (18 male, 18 female) were manually chosen from the filtered set to ensure high image quality and the absence of other distinctive features (e.g., jewelry, background objects). Two hierarchies were generated for each participant by randomly sampling 9 face images for each sex.

Procedure

There were two sessions. In the first session, participants completed the transitive inference task followed by the operation span task. The second session occurred 6-8 days after the first session and included only a second run of the test phases from the transitive inference task.

Transitive inference task. The transitive inference task (Figure 1) used a within-subjects design with two rounds. Participants learned about one 9-item hierarchy in the active condition and a second 9-item hierarchy in the passive condition. Each hierarchy was composed of either all female faces or all male faces in order to reduce interference between study conditions. The order of conditions and mapping of stimulus sex to condition were counterbalanced across participants. Each round was comprised of a learning phase (56 trials) followed by a test phase (72 trials).

Participants were instructed that the task involved learning about the “chain of command” at two different companies. During the instructions all 9 images from the to-be-learned hierarchy were presented in a horizontal array in random order. The instructions included an example of a 3-item hierarchy in which participants learned about two premise pairs (person A < person B, person B < person C) and were asked to infer the transitive relation (person A < person C). All participants were therefore informed of the hierarchical nature of the stimuli and were explicitly instructed to

learn to judge the relative rank of individuals in the company.

Learning phase. The learning phase involved a series of choices between two options corresponding to non-adjacent items in the present hierarchy (excluding the highest item in the hierarchy, which had no superordinate item and was never presented as a choice option). The options on the first learning trial were any two non-adjacent items sampled at random. On all subsequent trials, option sets were sampled in order to systematically vary their distance from the option selected on the previous trial: Each set included a *near* option that was 1–2 positions away from the option selected on the previous trial (either above or below), and a *far* option that was 3 or more positions away from the option selected on the previous trial. This manipulation of option distance was designed to test whether participants preferred to select items based on their distance in the active condition. In the passive condition selections were evenly divided between near and far options.

Learning trials: Active condition. Each learning trial began with the presentation of the two options in a vertical array in random order (Figure 1, middle). Participants were instructed to select an option using the mouse at their own pace in order to learn that person’s direct supervisor. Following their choice the unselected option disappeared and the premise pair (selected item and superordinate item) was displayed for 2 s. The options then disappeared and the experiment immediately proceeded to the next trial.

Learning trials: Passive condition. In the passive condition participants did not decide which option to select. As in the active condition, the trial began with the presentation of two options, one of which was already highlighted with a red border. Participants were instructed to select the highlighted option at their own pace, at which point the trial proceeded in the same manner as in the active condition.

Test phase: All conditions. In each test trial, two items were presented side-by-side and the participant was instructed to click on the item they judged to be ranked higher in the hierarchy. Test responses were self-paced and there was no time limit. The test phase was comprised of three types of trials: *recall* trials involving a choice between items from studied premise pairs (e.g., A ? B),

Table 1

Descriptive statistics from operation span task for Experiments 1 and 2.

	Math accuracy		Operation span		
	M	SD	M	SD	median
Experiment 1	0.92	0.06	14.86	11.29	12.5
Experiment 2	0.94	0.06	16.77	11.07	14

near inference trials involving items that were 2–3 positions apart (e.g., A ? C), and *far inference* trials involving items that were 4 or more positions apart (e.g., A ? E). There were 3 blocks of 24 trials, with each block made up of an equal number of recall, near inference, and far inference trials presented in random order. In the second session, participants completed a second run of the same test phase experienced during the first session, with test pairs presented in a new random order.

Operation span. Operation span is a well-established measure of working memory capacity in which participants attempt to hold a sequence of items in memory while evaluating a set of interleaved math operations (Turner & Engle, 1989; Unsworth, Heitz, Schrock, & Engle, 2005). In the version used for this study (obtained from <http://www.cognitivetools.uk/cognition/tasks/Verbal-WM/operationSpan/>) participants attempted to remember sequences of numbers while judging the validity of arithmetic problems. The presentation of each number was followed by a math operation which participants judged to be either correct (e.g., $2 - 4 = -2$) or incorrect (e.g., $2 - 4 = 6$). At the end of a trial involving multiple such steps, they then attempted to recall the sequence of numbers in the order in which they had appeared. The set size (number of operations/numbers) ranged from 2–7, presented in increasing order. There were three trials of each set size for a total of 18 trials. Operation span was scored according to the summed number of digits recalled in the correct order for those trials in which no errors were made (see Table 1 for descriptive statistics from all three studies).

Results

The results described below are based on data from all participants, including those who did not return for the second session. Returning and non-returning participants did not differ in overall accuracy on the first test, $\Delta M = -0.04$, 95% CI $[-0.10, 0.02]$, $t(64.03) = -1.35$, $p = .182$, or operation span, $\Delta M = -0.54$, 95% CI $[-5.14, 4.06]$, $t(81.28) = -0.23$, $p = .816$. All error bars in figures represent within-subjects 95% confidence intervals calculated using the Cousineau-Morey method (Morey, 2008).

Test performance

Test performance was analyzed for trials that did not involve either endpoint of the hierarchy, as participants may rely on non-transitive strategies to respond in inference trials that include an endpoint (Dusek & Eichenbaum, 1997). For instance, the highest-ranked item only appeared as the superordinate outcome during study and never as a choice option, allowing a learner to infer that it is at the top of the hierarchy without relying on memory of its relationship to any other items. The following analyses therefore included only non-endpoint trials in order to focus on test trials that required transitive inference.

Accuracy. Test responses were scored according to whether participants correctly identified the superordinate item in each test pair (0 = incorrect, 1 = correct). Accuracy was modeled using mixed effects logistic regression (see Appendix A for the model selection procedure). Figure 2 displays test accuracy during the immediate test (top) and delayed retest (bottom), along with the prediction of the best-fitting model. Table 2 presents parameter estimates for fixed effects in the best-fitting model in terms of relative odds ratios (*OR*), which indicate the multiplicative change in the odds of responding correctly given a unit change in the predictor. An *OR* below 1 indicates a decrease in the probability of a correct response, whereas an *OR* above 1 indicates an increase in the probability of a correct response. Individual model contrasts, including those which involve combinations of factors, are presented in the text along with 95% confidence intervals and hypothesis tests.

In the initial test, passive study did not differ from active study for recall ($OR = 0.81$ [0.44,

Table 2

*Estimated fixed effects from models of test accuracy and RT for
Experiment 1.*

Predictor	Accuracy	RT
	OR	beta
(Intercept)	4.26***	3.2***
Session [retest]	0.75*	-0.71***
Condition [passive]	0.81	-0.11
Distance [near]	1	-0.17*
Distance [far]	1.89***	-0.41***
Operation span	1.85***	0.1
Session [retest] x Condition [passive]	0.76	
Session [retest] x Distance [near]	1.24*	
Session [retest] x Distance [far]	1.08	
Condition [passive] x Distance [near]	0.99	0.02
Condition [passive] x Distance [far]	0.87	0.29*
Condition [passive] x Operation span	0.57***	
Distance [near] x Operation span	1.09	
Distance [far] x Operation span	1.7***	
Condition [passive] x Distance [near] x Operation span	1.04	
Condition [passive] x Distance [far] x Operation span	0.7*	

Note. *: $p < .05$; **: $p < .01$; ***: $p < .001$

Figure 2. Test performance in Experiment 1. Accuracy by option distance in the test (top row) and delayed retest (bottom row).

1.50], $z = -1.05$, $p = 0.29$), near inference ($OR = 0.80$ [0.43, 1.49], $z = -1.11$, $p = 0.27$), or far inference ($OR = 0.71$ [0.36, 1.38], $z = -1.62$, $p = 0.11$). In the delayed retest, however, passive study led to lower accuracy for recall ($OR = 0.62$ [0.34, 1.12], $z = -2.52$, $p = 0.01$), near inference ($OR = 0.61$ [0.33, 1.12], $z = -2.55$, $p = 0.01$) and far inference ($OR = 0.53$ [0.28, 1.03], $z = -2.96$, $p = 0.003$). Overall, active study therefore produced a clear advantage over passive study at the delayed retest for both recall and inference.

There was strong evidence of SDEs in both study conditions. In the initial test, accuracy in the active condition was higher for far inference than both recall ($OR = 1.89$ [1.11, 3.23], $z = 3.72$,

$p < .001$) and near inference ($OR = 1.89 [1.20, 2.96]$, $z = 4.40$, $p < .001$). An even clearer effect was observed in the active condition in the delayed retest: far inference accuracy was higher than both recall ($OR = 2.04 [1.16, 3.60]$, $z = 3.93$, $p < .001$) and near inference ($OR = 1.64 [1.02, 2.65]$, $z = 3.22$, $p = 0.001$), while near inference accuracy was also higher than recall ($OR = 1.64 [1.02, 2.65]$, $z = 3.22$, $p = 0.001$). Weaker SDEs were also found in the passive condition. In the initial test, far inference accuracy was higher than both recall ($OR = 1.64 [0.98, 2.74]$, $z = 3.01$, $p = 0.003$) and near inference ($OR = 1.66 [1.09, 2.54]$, $z = 3.72$, $p < .001$). At the delayed retest, far inference accuracy was higher than both recall ($OR = 1.77 [1.03, 3.06]$, $z = 3.29$, $p = 0.001$) and near inference ($OR = 1.45 [0.92, 2.28]$, $z = 2.53$, $p = 0.01$), and near inference was higher than recall ($OR = 1.23 [0.89, 1.69]$, $z = 1.98$, $p = 0.05$). Consistent evidence of SDEs in both conditions suggests that integrative encoding drove performance and that the gap between active and passive performance was unlikely to have arisen from differences in encoding strategy.

Recall decreased from the test to retest in both the active ($OR = 0.75 [0.53, 1.06]$, $z = -2.62$, $p = 0.009$) and passive condition ($OR = 0.57 [0.40, 0.80]$, $z = -5.12$, $p < .001$). In the passive condition, accuracy also declined to a lesser extent for both near inference ($OR = 0.70 [0.49, 1.02]$, $z = -2.98$, $p = 0.003$) and far inference ($OR = 0.61 [0.39, 0.95]$, $z = -3.45$, $p < .001$). In contrast, performance on inference trials did not change from test to retest in the active condition (near inference: $OR = 0.93 [0.66, 1.31]$, $z = -0.68$, $p = 0.50$; far inference: $OR = 0.81 [0.52, 1.24]$, $z = -1.55$, $p = 0.12$). Thus, the emergence of larger SDEs in the retest reflected decreases in recall but relatively preserved inference accuracy after the week delay.

Effects of WMC. Working memory capacity had a dramatic influence on these results. The best-fitting model included significant predictors for operation span and its interaction with study condition and distance (Table 2). Operation span was positively related to performance in the active condition at any distance (recall: $OR = 1.85 [1.18, 2.91]$, $z = 4.23$, $p < .001$; near inference: $OR = 2.02 [1.21, 3.36]$, $z = 4.30$, $p < .001$; far inference: $OR = 3.14 [1.52, 6.49]$, $z = 4.91$, $p < .001$). This effect was greater for far inference trials than both recall ($OR = 1.70 [0.98, 2.94]$, $z = 2.99$, $p = 0.003$) and near inference ($OR = 1.55 [0.99, 2.44]$, $z = 3.04$, $p = 0.002$), whereas it did not differ between recall and near inference ($OR = 1.09 [0.80, 1.48]$, $z = 0.89$, $p = 0.38$). In contrast,

Figure 3. Accuracy by option distance in the test and delayed retest in Experiment 1, following a median split on operation span.

operation span was unrelated to performance in the passive condition at any distance (all $p > .21$). WMC was therefore selectively associated with performance in the active condition, particularly for more distant inferences.

In order to explore the effect of WMC further, a followup analysis was performed using a median split on operation span to divide participants into low WMC and high WMC groups (Figure 3). Among low WMC participants, there was no difference in overall accuracy (collapsing across distance levels) between study conditions at either the test ($OR = 1.42 [0.70, 2.87]$, $z = 1.33$, $p = 0.18$) or retest ($OR = 1.39 [0.74, 2.62]$, $z = 1.42$, $p = 0.15$), and accuracy declined from test to retest in both conditions (active: $OR = 0.74 [0.56, 1.00]$, $z = -2.73$, $p = 0.006$; passive: $OR = 0.73 [0.51, 1.04]$, $z = -2.37$, $p = 0.02$). Among high WMC participants, passive performance was lower than active in both the test ($OR = 0.46 [0.22, 0.95]$, $z = -2.87$, $p = 0.004$) and retest ($OR = 0.27 [0.14, 0.53]$, $z = -5.24$, $p < .001$). Accuracy declined from test to retest in the passive condition ($OR = 0.53 [0.37, 0.75]$, $z = -4.79$, $p < .001$), but did not change in the active condition ($OR = 0.89 [0.59, 1.33]$, $z = -0.79$, $p = 0.43$). In sum, the overall gap in accuracy between active and passive study was driven by high WMC participants, who showed an advantage from active study in both

test sessions and for whom performance in the active condition was sustained across a week delay.

Reaction time. Reaction time at test was analyzed using linear mixed effects regression. The parameters of the best-fitting model are presented in Table 2. RT declined from the test to the retest ($\beta = -0.71 [-0.85, -0.57]$, $z = -13.83$, $p < .001$). Responses were slower in the passive condition than the active condition for far inference ($\beta = 0.18 [-0.03, 0.40]$, $z = 2.32$, $p = 0.02$), but there were no other differences between conditions.

There was also evidence for SDEs based on RT. Within the active condition, responses were faster for far inference compared to both recall ($\beta = -0.41 [-0.63, -0.19]$, $z = -5.19$, $p < .001$) and near inference ($\beta = -0.24 [-0.46, -0.02]$, $z = -3.04$, $p = 0.002$). Responses were also faster for near inference compared to recall ($\beta = -0.17 [-0.39, 0.05]$, $z = -2.15$, $p = 0.03$). Similarly, within the passive condition, RTs were smaller for far inference than both recall ($\beta = 0.29 [-0.01, 0.60]$, $z = 2.63$, $p = 0.009$) and near inference ($\beta = 0.27 [-0.03, 0.58]$, $z = 2.45$, $p = 0.01$). In contrast to the accuracy results, operation span was not related to RT ($\beta = 0.10 [-0.10, 0.30]$, $z = 1.35$, $p = 0.18$).

Selections during learning

The next analysis examined participants' selections during learning and whether they could account for differences in test accuracy described above. Study condition was not related to premise frequency during study (multinomial logistic regression, likelihood ratio test: $\chi^2_{(1,7)} = 7.20$, $p = 0.41$), indicating that the distribution of experienced premise pairs was similar across active and passive study. The remaining analyses focused on the preference to select either near or far options during study.

Proportion of near option selections. Each learning trial involved a choice between a near option (1–2 positions away from the option selected on the previous trial) and a far option (3 or more positions away). By design, the proportion of near choices in the passive condition was approximately 50% ($M = 0.50$, $SD = 0.01$). In the active condition there was a small but robust preference to select the near option ($M = 0.56$, $SD = 0.07$; $M = 0.56$, 95% CI $[0.55, 0.58]$, $t(99) = 77.86$, $p < .001$). As a result, the average distance between successive selections was smaller in the active condition ($M = 2.63$, $SD = 0.13$) than the passive condition ($M = 2.83$, $SD = 0.13$; $M_d = -0.21$, 95% CI $[-0.26, -0.15]$, $t(99) = -7.53$, $p < .001$).

Figure 4. **A:** Depiction of procedure for sampling options based on distance to selection in previous learning trial. Each choice involved a near option (1-2 positions away from previous selection, see green region) and a far option (3+ positions away, see blue region). **B:** Selection results for Experiment 1 (left column) and Experiment 2 (right column) as a function of near option distance. The top row displays the proportion of trials in which the near option was chosen. The bottom two rows display the median selection RT on those trials in which the far option was chosen (middle row) and trials in which the near option was chosen (bottom row).

Near selections may be especially useful if they cause overlapping premise pairs to be experienced in successive trials. If representations of overlapping pairs are simultaneously active in memory, the learner may be better able to integrate them together into a unified hierarchy. I next examined whether the preference to select near items in the active condition depended on the distance between the near option and the item selected on the previous trial ($dist_{near} \in \{-2, -1, +1, +2\}$). When $dist_{near} = +1$, the near option was immediately superordinate to the previously selected item; that is, the near option belonged to an overlapping pair that appeared in the previous trial (see Figure 4A).

Figure 4B (top row, left column) shows the proportion of near selections as a function of near option distance. In the passive condition, the likelihood of selecting the near item was fixed at .5 regardless of its distance from the previous selection. In the active condition, the proportion of near selections did not differ from chance when $dist_{near} = -2$ ($OR = 1.10$ [0.95, 1.26], $z = 1.59$, $p = 0.11$). However, the proportion of near selections was greater than chance when $dist_{near} = -1$ ($OR = 1.22$ [1.07, 1.40], $z = 3.77$, $p < .001$), $dist_{near} = +1$ ($OR = 1.78$ [1.55, 2.04], $z = 10.47$, $p < .001$), or $dist_{near} = +2$ ($OR = 1.15$ [0.99, 1.32], $z = 2.33$, $p = 0.02$). Thus, in the active condition participants preferred to select the near option when it was close to the option selected on the previous trial, and this preference was strongest when the option had appeared as the superordinate outcome in that trial. Compared to the passive condition, participants generated study sequences during active study in which overlapping premise pairs were more likely to be experienced in successive trials.

Can this tendency to select overlapping pairs account for the higher accuracy seen in the active condition? A new model of test accuracy was fit which included predictors for the overall proportion of near selections and its interaction with condition, but this indicated that the proportion of near selections was not related to accuracy in the active condition ($OR = 1.14$ [0.88, 1.49], $z = 0.99$, $p = 0.32$). The proportion of near selections at any distance was also unrelated to operation span ($dist_{near} = -2$: $OR = 0.93$ [0.80, 1.07], $z = -1.31$, $p = 0.19$; $dist_{near} = -1$: $OR = 0.95$ [0.83, 1.09], $z = -0.91$, $p = 0.36$; $dist_{near} = +1$: $OR = 1.06$ [0.92, 1.21], $z = 1.00$, $p = 0.32$; $dist_{near} = +2$: $OR = 0.96$ [0.84, 1.11], $z = -0.62$, $p = 0.54$). Selection of overlapping pairs therefore appeared to be a general preference that did not vary with WMC, suggesting that this search behavior cannot account for the divergence between active and passive study among higher WMC participants.

Selection response time. The final analysis examined the amount of time taken to either choose an option for study (active condition) or to select a predetermined option (passive condition). Median response times during selection were higher in the active condition ($M = 3.15$ s, $SD = 1.54$) than the passive condition ($M = 2.49$ s, $SD = 1.26$; $\beta = 0.65$ [0.37, 0.93], $z = 5.52$, $p < .001$), suggesting that active selection decisions were associated with an additional processing load.

In addition, selection RT in the active condition was positively related to operation span ($\beta = 0.49$ [0.16, 0.81], $z = 3.55$, $p < .001$), whereas there was no such relationship in the passive condition ($\beta = 0.21$ [-0.12, 0.53], $z = 1.52$, $p = 0.13$). Higher WMC participants who tended to have high accuracy in the active condition therefore also took longer to make selection decisions in that condition.

Lastly, I examined how selection RT depended on the distance of the near option, which was shown in the previous section to be strongly preferred when it was immediately superordinate to the option selected in the previous trial ($dist_{near} = +1$). Figure 4B shows median RT on trials in which the far option was selected (middle row, left) and trials in which the near option was selected (bottom row, left). A three-way repeated measures ANOVA was performed with condition (active/passive), near option distance (-2/-1/+1/+2), and selection (near/far) as within-subjects factors. In addition to the main effect of condition noted above ($F(1, 1483) = 121.71$, $MSE = 1.65$, $p < .001$, $\hat{\eta}_G^2 = .076$), there was a main effect of selection ($F(1, 1483) = 4.46$, $MSE = 1.65$, $p = .035$, $\hat{\eta}_G^2 = .003$) with RT lower when the near option was selected. There was also a main effect of near option distance ($F(3, 1483) = 6.22$, $MSE = 1.65$, $p < .001$, $\hat{\eta}_G^2 = .012$) and an interaction between selection and near distance ($F(3, 1483) = 7.79$, $MSE = 1.65$, $p < .001$, $\hat{\eta}_G^2 = .016$). Post-hoc comparisons indicated that when the far was option was selected (Figure 4B, middle left), there were no differences in RT as a function of distance in either condition (Tukey HSD, all $p > .4$). When the near option was selected in the active condition (Figure 4, bottom left), $dist_{near} = +1$ options were selected faster than all other types (all $p < .001$) but there were no other differences between option distances. In the passive condition, $dist_{near} = +1$ options were also selected faster than $dist_{near} = -2$ options ($p < .001$) but there were no other differences between near option distances. In sum, the preference to select options that had appeared as the superordinate item on the previous trial was also evident in faster selection decisions. The primacy of these $dist_{near} = +1$ options was even apparent in the passive condition in which participants merely had to select predetermined options.

Discussion

Active control over the selection of premise pairs improved TI performance relative to passive observation in both an immediate test and a retest one week later. Symbolic distance effects, whereby accuracy increased (and RT decreased) with greater distances between test items, were found in both the active condition and, to a lesser extent, in the passive condition. This pattern suggests that participants engaged in integrative encoding of a unified representation of the hierarchy during study (Acuna et al., 2002; Zeithamova et al., 2012) and that this process was enhanced by control over the training sequence.

This study provides the first evidence of systematic search in TI: There was a strong preference to select options that appeared as the superordinate item on the previous trial ($dist_{near} = +1$). Participants thereby naturally generated “chained” sequences of overlapping premise pairs. This preference was widespread: 82% of participants chose the $dist_{near} = +1$ option in more than half of trials in which one appeared. Participants were also fastest to select $dist_{near} = +1$ options in both study conditions. Previous research on TI suggests this preference is adaptive, as chained sequences produce more efficient learning under passive conditions compared to random sequences (Andrews, 2010; Halford, 1984; Waltz et al., 2004). Studying overlapping premises in successive trials may facilitate integrative encoding for learners with limited capacity to remember experienced premises. If a learner sees the sequence $A < B$, $B < C$, $C < D$ and all three premises are simultaneously held in short-term memory, they can combine the premises to construct a single ordering of four items ($A < B < C < D$). In contrast, consider the training sequence $A < B$, $C < D$, $E < F$. Although the premises themselves might be encoded successfully, they cannot be used to order items from different pairs with respect to one another. By the time an overlapping premise is experienced (e.g., $B < C$) the earlier pairs may no longer be maintained in short-term memory, reducing the opportunity for integrative encoding.

Although selection of overlapping pairs would seem to promote integrative encoding, only those participants with higher WMC benefited from this search behavior. Among higher WMC participants, active control produced a ~10% initial advantage over passive study (increasing past 20% in the delayed retest). This is a striking demonstration of how the opportunity to control

study may not benefit individuals without the cognitive resources to make use of the information that is generated. Lower WMC participants selected overlapping premises in the active condition with the same frequency, but did not derive any advantage over passive observation of randomly generated sequences.

In contrast to the active condition, WMC was unrelated to accuracy in the passive condition. This finding conflicts with prior reports that WMC is correlated with TI accuracy in passive settings (e.g., Libben & Titone, 2008), but this discrepancy is likely due to the relative difficulty of passive study in the present experiment. Previous studies showing effects of WMC on TI have typically involved smaller hierarchies (ranging from 3–6 items) and scaffolded training sequences. For instance, Libben and Titone (2008) used clustered training sequences in which participants were likely to experience overlapping pairs in successive trials. Other studies linking TI to working memory (Fales et al., 2003; Waltz et al., 2004) have focused on relational integration while all stimuli were simultaneously displayed, removing any demands on maintenance of premise pairs in memory. With larger hierarchies, sequential presentation, and training sequences with comparatively large distances between successive premises, the passive condition used here may have been difficult even for participants with higher WMC. In order to better match training experiences across conditions, Experiment 2 adopts a yoked design in which the passive condition involves observation of sequences generated by participants in the active condition of Experiment 1.

The inclusion of a delayed retest provides additional insight into the longer term benefits of active control for constructing relational knowledge. Among higher WMC participants there was no significant change in accuracy for the active condition, whereas performance declined in all other conditions across a week delay. In general, drops in performance were greater for direct recall than for inference, thereby producing stronger SDEs in the retest. This pattern is consistent with recent studies reporting the emergence of relational knowledge after a delay due to sleep-dependent consolidation (Ellenbogen, Hu, Payne, Titone, & Walker, 2007; Stickgold & Walker, 2013; Werchan & Gómez, 2013). In a TI task involving training by reinforcement, Ellenbogen et al. (2007) found that performance on more distant inferences was at chance immediately following training but emerged after 12- and 24-hour delays that included sleep. Werchan and Gómez (2013) also found

that inference performance increased after a delay that included sleep, but only when training involved reinforcement. In conditions involving observational training (most similar to the passive condition here) inference accuracy did not exceed chance after the delay. It is not possible to directly compare these results to the present study given numerous differences in design, including their focus on inference in the absence of explicit awareness of the hierarchy and use of training to criterion to ensure equivalent recall of premise pairs by the end the study phase. Further work is necessary to understand whether the longlasting benefit from active study among higher WMC participants simply reflects more successful encoding during study, or if active control causes changes in the encoding process that promote the consolidation of relational knowledge.

Experiment 2

The results of the first experiment revealed a striking advantage from active study among higher WMC participants. As noted above, the relatively poor accuracy in the passive condition may be due to the pseudo-random training sequences in that condition. That is, the benefit from controlling study may be that it allows learners to generate more useful training sequences in which overlapping premises are more likely to occur in successive trials. To evaluate this possibility, Experiment 2 was designed to more closely match the training sequences across active and passive study using a yoked design. During passive study, participants were yoked to sequences generated by active participants in Experiment 1. The active condition was the same as in Experiment 1. Because an advantage from active control was evident among higher WMC participants in the immediate test in Experiment 1, this experiment did not include a delayed retest.

Participants

$N = 100$ participants (56 women; age: $M = 20.60$ years, $SD = 4.07$) were recruited from the student population at UNC Charlotte. Participants received either course credit or \$4 as compensation for participating in the study. All participants received an incentive payment based on their performance in the test ranging from \$0 ($< 50\%$ correct) to \$5 (90–100% correct).

Table 3

Estimated fixed effects from models of test accuracy and RT for Experiment 2.

Predictor	Accuracy	RT
	OR	beta
(Intercept)	4.72***	3.11***
Condition [yoked]	0.91	-0.02
Distance [near]	0.82*	-0.11
Distance [far]	1.07	-0.26***
Operation span	1.54*	-0.07
Condition [yoked] x Operation span	0.97	0.19*
Distance [near] x Operation span	1.08	
Distance [far] x Operation span	1.25	
Condition [yoked] x Distance [near]	1.04	
Condition [yoked] x Distance [near] x Operation span	0.75*	
Condition [yoked] x Distance [far]	1.12	
Condition [yoked] x Distance [far] x Operation span	0.63***	

Note. *: $p < .05$; **: $p < .01$; ***: $p < .001$

Materials and Procedure

All materials were identical to those of the first session in Experiment 1. Participants completed a single session comprised of the TI task followed by the operation span assessment. There was no second session in this experiment. When a participant began the TI task they were assigned a partner from Experiment 1. The study sequence generated by the partner during the active condition in Experiment 1 was reproduced exactly for the new participant's passive study round, including the set of faces, their order in the hierarchy, the choice options in each selection trial, and the selected premise pairs. As in Experiment 1, the order of study conditions was counterbalanced across participants.

Figure 5. Accuracy by condition and option distance in the test in Experiment 2.

Results

Test performance

Accuracy. The parameters of the fitted mixed effects logistic regression model are shown in Table 3. There were no differences in accuracy (Figure 5) between active and yoked conditions for recall ($OR = 1.10 [0.61, 1.99]$, $z = 0.46$, $p = 0.64$), near inference ($OR = 1.05 [0.58, 1.91]$, $z = 0.26$, $p = 0.80$), or far inference ($OR = 0.98 [0.51, 1.87]$, $z = -0.09$, $p = 0.93$). Within the active condition, near inference accuracy was lower than both recall ($OR = 0.82 [0.63, 1.06]$, $z = -2.31$, $p = 0.02$) and far inference ($OR = 1.31 [0.96, 1.79]$, $z = 2.56$, $p = 0.01$), but there was no difference between recall and far inference ($OR = 1.07 [0.79, 1.46]$, $z = 0.68$, $p = 0.50$). Within the yoked condition, there were no pairwise differences between recall, near inference, and far inference. Thus, there were only partial evidence of a SDE for accuracy in the active condition.

There were significant interactions between operation span, distance, and condition (Table 3). In the active condition, operation span was positively related to performance for recall ($OR = 1.54 [0.95, 2.49]$, $z = 2.67$, $p = 0.008$), near inference ($OR = 1.67 [1.03, 2.71]$, $z = 3.15$, $p = 0.002$), and far inference ($OR = 1.93 [1.13, 3.29]$, $z = 3.67$, $p < .001$). In the yoked condition, operation span was positively related to performance on recall ($OR = 1.49 [0.98, 2.28]$, $z = 2.82$, $p = 0.005$), but

not near inference ($OR = 1.21 [0.79, 1.85]$, $z = 1.35$, $p = 0.18$) or far inference ($OR = 1.18 [0.75, 1.88]$, $z = 1.08$, $p = 0.28$). For far inference only, the effect of operation span on accuracy was significantly larger in the active condition than the yoked condition ($OR = 1.63 [0.84, 3.17]$, $z = 2.18$, $p = 0.03$). Thus, although there were no overall differences in performance between active and yoked study, WMC was more consistently related to performance in the active condition, particularly for more distant inferences. This can be seen in the prediction from the regression model in Figure 5, where accuracy in the passive condition is relatively constant but accuracy in the active condition increases with operation span.

Reaction time. There were no differences in median RT (Table 3) between the active and yoked conditions ($\beta = 0.02 [-0.17, 0.20]$, $z = 0.21$, $p = 0.83$). Differences in RT with distance were consistent with a limited SDE: RTs were significantly smaller on far inference than recall trials ($\beta = -0.26 [-0.49, -0.03]$, $z = -2.93$, $p = 0.003$) but not significantly different from near inference trials ($\beta = -0.15 [-0.38, 0.08]$, $z = -1.69$, $p = 0.09$). There was no difference in RT between recall and near inference trials ($\beta = -0.11 [-0.34, 0.12]$, $z = -1.24$, $p = 0.21$). Finally, operation span was not related to RT in either the active ($\beta = -0.07 [-0.31, 0.17]$, $z = -0.72$, $p = 0.47$) or yoked condition ($\beta = 0.13 [-0.11, 0.37]$, $z = 1.36$, $p = 0.17$).

Selections during learning

Study condition was not related to premise frequency (multinomial logistic regression, likelihood ratio test: $\chi^2_{(1,7)} = 3.68$, $p = 0.82$), indicating that the aggregate distribution of experienced premise pairs was similar across active and yoked study.

Proportion of near option selections. The overall tendency to select the near option in the active condition ($M = 0.55$, $SD = 0.07$) was not different from that of their partners from Experiment 1 ($OR = 0.94 [0.87, 1.02]$, $z = -1.54$, $p = 0.12$). In the active condition, the proportion of near selections did not differ from chance when $dist_{near} = -2$ ($OR = 1.07 [0.92, 1.25]$, $z = 1.19$, $p = 0.23$), but there was a significant preference for the near option when $dist_{near} = -1$ ($OR = 1.14 [0.99, 1.31]$, $z = 2.47$, $p = 0.01$), $dist_{near} = +1$ ($OR = 1.55 [1.35, 1.79]$, $z = 8.40$, $p < .001$), or $dist_{near} = +2$ ($OR = 1.14 [0.98, 1.33]$, $z = 2.37$, $p = 0.02$). The proportion of near selections (Figure 4B, top center) was higher for $dist_{near} = +1$ than $dist_{near} = -1$ options ($OR = 1.37 [1.12,$

1.67], $z = 4.21$, $p < .001$), again demonstrating a strong preference to select options that had appeared as the superordinate item on the previous trial. The proportion of near selections was not related to operation span at any distance (all $p > .09$).

A new model of test accuracy was fit which included predictors for the overall proportion of near selections and its interaction with condition. This revealed that accuracy in the active condition was higher for those participants who more often selected the near option ($OR = 1.46$ [1.03, 2.07], $z = 2.45$, $p = 0.01$). In contrast, the proportion of near selections in the yoked condition (i.e., based on a previous participant's choices) did not relate to performance ($OR = 0.96$ [0.72, 1.27], $z = -0.36$, $p = 0.72$). This further suggests that nearby options were more likely to benefit performance when they were actively chosen than when observed in a sequence generate by someone else.

Selection response time. Median selection RT was higher in the active condition ($M = 2.68$ s, $SD = 1.26$) than the yoked condition ($M = 2.35$ s, $SD = 1.34$; $\beta = 0.33$ [0.13, 0.54], $z = 3.78$, $p < .001$). Selection RT in the active condition was positively related to operation span ($\beta = 0.26$ [-0.04, 0.56], $z = 2.03$, $p = 0.04$), whereas there was no relationship between operation span and selection RT in the yoked condition ($\beta = 0.23$ [-0.08, 0.53], $z = 1.76$, $p = 0.08$).

Figure 4 shows median RT on trials in which the far option was selected (middle row, center) and trials in which the near option was selected (bottom row, center). A three-way repeated measures ANOVA was performed with condition (active/yoked), near option distance (-2/-1/+1/+2), and selection (near/far) as within-subjects factors. In addition to the main effect of condition noted above ($F(1, 1483) = 34.90$, $MSE = 1.27$, $p < .001$, $\eta_G^2 = .023$), there was a main effect of selection ($F(1, 1483) = 6.09$, $MSE = 1.27$, $p = .014$, $\eta_G^2 = .004$) with RT lower when the near option was selected. There was also a main effect of near option distance ($F(3, 1483) = 7.11$, $MSE = 1.27$, $p < .001$, $\eta_G^2 = .014$) and an interaction between selection and near distance ($F(3, 1483) = 2.70$, $MSE = 1.27$, $p = .044$, $\eta_G^2 = .005$). Post-hoc comparisons indicated that when the far was option was selected, there were no differences in RT as a function of near option distance in either condition (Tukey HSD, all $p > .4$). When the near option was selected in the active condition, $dist_{near} = +1$ options were selected faster than all other types (all $p < .02$) but

there were no other differences between option distances. In the yoked condition, $dist_{\text{near}} = +1$ options were selected faster than $dist_{\text{near}} = +2$ options ($p < .01$) but there were no other differences between near option distances. Thus, responses in both conditions were again fastest when selecting an option that overlapped with the premise pair experienced in the previous trial.

Discussion

Overall, test accuracy did not differ between active and yoked conditions, suggesting that the generation of a more useful training sequence may largely account for the benefit of active study seen among higher WMC participants in Experiment 1. As in that experiment, WMC had a strong influence on performance in the active condition, as WMC was positively related to accuracy for both recall and inference. However, the relationship between WMC and accuracy in the yoked condition was notably weaker, as higher WMC was only associated with improved accuracy on recall trials. In addition, while there was some evidence of an SDE for accuracy in the active condition, there was no corresponding effect in the yoked condition. Thus, although yoking appeared to abolish any deficit in accuracy compared to active study, there was nevertheless evidence that WMC was positively related to the success of integrative encoding in the active condition (Figure 5).

Why might WMC predict inference performance in the active condition but not the yoked condition? In the aggregate, there were no apparent differences in training experiences that could account for this effect. Premise frequency and the proportion of near selections did not differ between conditions, and the proportion of near selections was unrelated to WMC. If the only role of WMC is to support integrative encoding of a training sequence (however it is generated), then it should be equally predictive of performance in both the active and yoked conditions. This suggests a further role for WMC in active selection. One possibility is that higher WMC participants engage in additional metacognitive processing when deciding which option to learn about (e.g., relating each option to their existing knowledge of the hierarchy), and this processing either contributes to integrative encoding or allows them to optimize their training sequence beyond what is possible from observing a sequence generated by someone else. This interpretation is indirectly supported by the finding that WMC was positively related to selection RT in the active, but not yoked, condition

(a relationship also found in Experiment 1).

Finally, study behavior in the active condition was highly consistent with that of Experiment 1. Learners preferred to select the near option, particularly when it had appeared as the superordinate item on the previous trial, and this search preference was unrelated to WMC. Participants were also faster when choosing $dist_{\text{near}} = +1$ options, even in the yoked condition where the selection was predetermined. In one notable departure from Experiment 1, in this experiment there was also evidence that the proportion of near selections was positively related to accuracy in the active, but not yoked, condition. This provides further indication of a benefit from actively choosing to study nearby options, even when the resulting study sequences are highly similar to those experienced during passive observation.

General Discussion

How does the construction of knowledge depend on the freedom to control one’s own learning experience? Proponents of self-directed, “active” instruction argue that learner control is critical to knowledge construction (Bruner, 1961; Dewey, 1929; Kolb, 1984). By interacting with the environment, testing hypotheses through experimentation, or otherwise directing the flow of information during study, active instruction lets learners shape their experience according to their own goals and existing knowledge. Despite the intuitive appeal of such personalized learning, however, evidence for the efficacy of learner-directed instruction is mixed (Alfieri et al., 2011; Bernstein, 2018). Educational psychologists generally endorse the constructivist view that conceptual knowledge is assembled by the learner, but increasingly question whether the learner always benefits from collecting the raw materials themselves (R. E. Clark, 2009).

The current study aims to advance this debate by investigating the cognitive basis of self-directed learning (T. M. Gureckis & Markant, 2012; Markant et al., 2016). Whereas recent studies have demonstrated that learner control enhances elemental encoding of information directly experienced during study (Markant et al., 2014; Murty, DuBrow, & Davachi, 2015; A. Ruggeri, Markant, Gureckis, & Xu, 2019; Voss et al., 2011), less is known about how active control impacts integrative encoding, through which experienced information is combined into generalizable

knowledge during study. To address this gap, the present study used a transitive inference task to compare active control with passive observation when learning a relational hierarchy (Figure 1). The results indicate that active control enhances integrative encoding, but that this improvement is far from universal. Active control facilitated the formation of relational knowledge for learners with higher WMC, but among lower WMC participants did not differ from passive observation either for pseudo-random training sequences (Experiment 1) or for sequences generated by other participants (Experiment 2). WMC predicted performance on both direct recall of studied premises and transitive inference, but in both experiments was most strongly related to the ability to infer the relation between more distant items in the hierarchy. These findings provide further evidence of the importance of working memory in transitive inference (Libben & Titone, 2008; Waltz et al., 2004; Wendelken & Bunge, 2010) and, more generally, in the construction of mental representations from disparate elements (Hummel & Holyoak, 2003; Oberauer, Weidenfeld, & Hörnig, 2006; Vandierendonck & De Vooght, 1997).

There was consistent evidence of symbolic distance effects (SDEs) such that accuracy increased and RT decreased with greater distances between test items. SDEs are a hallmark of integrative encoding, reflecting the principle that it is easier to judge the relative positions of items that are more distant on an integrated latent dimension (Acuna et al., 2002; Moyer & Landauer, 1967; Trabasso, 1977). In contrast, inference based solely on elemental encoding—i.e., where premises are retrieved independently at test and used to reason across overlapping relations—should show the opposite pattern, with lower accuracy and longer RT for more distant items (Kumaran & McClelland, 2012). SDEs in the present task are in line with prior findings that learners rely on integrative encoding when they are aware of the hierarchical nature of the materials (Lazareva & Wasserman, 2010; Moses et al., 2006) and that performance in those conditions are related to WMC (Libben & Titone, 2008).

It is important to note that alternative learning mechanisms based on associative or reinforcement learning can also produce SDEs (Frank, Rudy, & O'Reilly, 2003; Frank, Rudy, Levy, & O'Reilly, 2005; Leo & Greene, 2008; Moses et al., 2006; Von Fersen, Wynne, Delius, & Staddon, 1991). Value-transfer theory (Von Fersen et al., 1991) provides one such account of TI that does not

require explicit reasoning about the hierarchy or coordination of multiple premises. Value-transfer theory proposes that the learner assigns a value to each item based on the rate at which it is reinforced during training, and that these values generalize between items that appear together as premise pairs. Transitive inference entails comparing the learned values of any two items and choosing the item with a higher value. Like integrative encoding, this process produces an SDE since distant items (which tend to be closer to endpoints with the highest or lowest values) exhibit larger disparities in value than adjacent items. Although this account cannot be ruled out entirely, there are several reasons to conclude that participants relied on integrative encoding in the present study. Prior work shows that explicit reasoning is more likely when learners are aware of the underlying hierarchy (Frank, O'Reilly, & Curran, 2006; C. Smith & Squire, 2005) and that awareness is associated with errors that are inconsistent with value transfer theory (Lazareva & Wasserman, 2010; Libben & Titone, 2008; Moses et al., 2006). Moreover, the influence of WMC on performance strongly suggests a reliance on explicit reasoning rather than a simpler associative learning process. The findings of Libben and Titone (2008) illustrate this distinction most clearly in their comparison of participants who differed in their awareness of the underlying hierarchy. Among unaware participants, WMC was unrelated to inference accuracy and the pattern of errors was consistent with value-transfer theory. In contrast, among aware participants, WMC was positively related to accuracy and performance was most consistent with a reliance on integrative encoding. Thus, it appears most likely that current task encouraged the use of integrative encoding and that this process was enhanced by the opportunity to control study.

A preference for “chained” study sequences

Taken together, the results suggest that a principal benefit of active control in TI is that it allows learners to generate study sequences that are more amenable to integrative encoding. Participants strongly preferred to learn about options that were nearby to items that had just been studied, and this preference was strongest when the near option was part of an overlapping premise pair. As a result, participants generated sequences of chained premises, a type of training that typically leads to faster learning than random sequences under passive conditions (Andrews, 2010). The strong preference for overlapping options even appeared to speed responses during passive

observation where selections were predetermined.

If chained sequences of overlapping pairs aid integrative encoding, why did only higher WMC participants benefit? It may be that only higher WMC individuals capitalize on chained sequences because they are more likely to maintain representations of premises from trial to trial and integrate information from multiple premises. In other words, those individuals with higher WMC simply have the resources necessary for integrative encoding, regardless of whether the study sequence is self-generated or provided by an external source. This interpretation is supported by the finding that higher WMC participants performed equally well in the active and yoked conditions in Experiment 2.

As noted above, however, active and yoked study differed in an important respect in that experiment: WMC was positively related to accuracy on inference trials in the active, but not yoked, condition. This suggests that the effect of active control cannot be completely accounted for by the makeup of the study sequence. One intriguing possibility to be examined in future work is that higher WMC individuals draw on their knowledge of the hierarchy when deciding which option to study. This may explain why selection RT was positively related to WMC in both experiments, reflecting additional time prior to choosing an option and before the superordinate outcome was revealed. If controlling the selection of premises causes learners to reactivate their existing knowledge, this decision process may on its own pave the way for integrative encoding. This would lend support to the popular notion that knowledge construction is more effective when the learner decides how to study materials for themselves.

Implications for learner-directed instruction

The present findings offer a note of caution regarding the effects of learner control in the formation of relational knowledge. Lower WMC participants showed no advantage from active control compared to passive observation, despite being aware of the hierarchy and having the opportunity to select premises for themselves. This finding may be unsurprising to researchers who have argued that having to search for information or independently discover concepts imposes an extraneous cognitive load that disadvantages lower WMC learners (Kirschner et al., 2006). It seems

unlikely, however, that lower WMC participants here were limited by the demands of making selection decisions, as WMC was unrelated to the frequency with which participants selected nearby options. Moreover, performance in the yoked condition of Experiment 2 indicates that release from the need to control study did not have any apparent benefit. It appears instead that lower WMC participants were less able to engage in integrative encoding and, as a consequence, could not take advantage of more useful study sequences.

This finding highlights the risk of conflating overt control over the study experience—a ubiquitous feature of “active” instruction—with the internal construction of knowledge. To physically interact with or to generate useful evidence from the environment does not guarantee that the learner extracts generalizable knowledge from that experience. For lower WMC learners in particular, active control may be most effective in environments that support integration across study episodes, including prompts to reactivate related materials (Kesteren, Krabbendam, & Meeter, 2018) or simultaneous presentation of materials (Eglington & Kang, 2018; Son, Smith, & Goldstone, 2011). The construction of relational knowledge is a fundamental learning goal that spans educational contexts (Dumas, Alexander, & Grossnickle, 2013). In order to meet the needs of diverse populations of learners, including those with impairments in cognitive functioning, there is a need for instructional approaches that guide conceptual integration during the course of self-directed study.

References

- Acuna, B. D., Sanes, J. N., & Donoghue, J. P. (2002). Cognitive mechanisms of transitive inference. *Experimental Brain Research*, 146(1), 1–10.
- Alfieri, L., Brooks, P., Aldrich, N., & Tenenbaum, H. (2011). Does discovery-based instruction enhance learning? *Journal of Educational Psychology*, 103(1), 1.
- Andrews, G. (2010). Belief-based and analytic processing in transitive inference depends on premise integration difficulty. *Memory & Cognition*, 38(7), 928–940.
- Bainbridge, W. A., Isola, P., & Oliva, A. (2013). The intrinsic memorability of face photographs. *Journal of Experimental Psychology: General*, 142(4), 1323.
- Bernstein, D. A. (2018). Does active learning work? A good question, but not the right one. *Scholarship of Teaching and Learning in Psychology*, 4(4), 290.
- Bruner, J. (1961). The act of discovery. *Harvard Educational Review*, 31(1), 21–32.
- Clark, R. E. (2009). How much and what type of guidance is optimal for learning from instruction? In S. Tobias & T. M. Duffy (Eds.), *Constructivist instruction: Success or failure* (pp. 158–183). New York: Routledge, Taylor; Francis.
- De Soto, C. B., London, M., & Handel, S. (1965). Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology*, 2(4), 513.
- Dewey, J. (1929). The quest for certainty.
- Dewey, J. (1938). *Experience and education*. Kappa Delta Pi.
- Dumas, D., Alexander, P. A., & Grossnickle, E. M. (2013). Relational reasoning and its manifestations in the educational context: A systematic review of the literature. *Educational Psychology Review*, 25(3), 391–427.
- 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.
- Eglington, L. G., & Kang, S. H. (2018). Retrieval practice benefits deductive inference. *Educational Psychology Review*, 30(1), 215–228.
- 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.
- Fales, C. L., Knowlton, B. J., Holyoak, K. J., Geschwind, D. H., Swerdloff, R. S., & Gonzalo, I. G. (2003). Working memory and relational reasoning in Klinefelter syndrome. *Journal of the International Neuropsychological Society*, 9(6), 839–846.
- 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.
- Frank, M. J., Rudy, J. W., & O'Reilly, R. C. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus. II. A computational analysis. *Hippocampus*, 13(3), 341–354.
- 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.
- Gureckis, T. M., & Markant, D. (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.
- Hummel, J. E., & Holyoak, K. J. (2001). A process model of human transitive inference. In *Spatial schemas in abstract thought* (pp. 279–305).
- Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference

- and generalization. *Psychological Review*, 110(2), 220.
- Huttenlocher, J. (1968). Constructing spatial images: A strategy in reasoning. *Psychological Review*, 75(6), 550.
- Kesteren, M. T. R. van, Krabbendam, L., & Meeter, M. (2018). Integrating educational knowledge: Reactivation of prior knowledge during educational learning enhances memory integration. *Npj Science of Learning*, 3(1), 11.
- Kirschner, P., Sweller, J., & Clark, R. (2006). Why minimal guidance during instruction does not work: An analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching. *Educational Psychologist*, 41(2), 75–86.
- Kolb, D. (1984). *Experiential learning: Experience as the source of learning and development*. Englewood Cliffs, NJ: Prentice-Hall.
- Kumaran, D., & McClelland, J. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, 119(3), 573.
- Lagnado, D. A., & Sloman, S. A. (2004). The advantage of timely intervention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(4), 856–876.
- Lazareva, O. F., & Wasserman, E. A. (2010). Nonverbal transitive inference: Effects of task and awareness on human performance. *Behavioural Processes*, 83(1), 99–112.
- 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., & Gureckis, T. M. (2014). 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., DuBrow, S., Davachi, L., & Gureckis, T. M. (2014). Deconstructing the effect of self-directed study on episodic memory. *Memory & Cognition*, 42(8), 1211–1224.
- 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.
- Mayer, R. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist*, 59(1), 14–19.
- Metcalfe, J. (2009). Metacognitive judgments and control of study. *Current Directions in Psychological Science*, 18(3), 159–163.
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4(2), 61–64.
- Moses, S. N., Villate, C., & Ryan, J. D. (2006). An investigation of learning strategy supporting transitive inference performance in humans compared to other species. *Neuropsychologia*, 44(8), 1370–1387.
- 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.
- Moyer, R. S., & Landauer, T. K. (1967). Time required for judgements of numerical inequality. *Nature*, 215, 1519–1520.
- Murty, V. P., DuBrow, S., & Davachi, L. (2015). The simple act of choosing influences declarative memory. *The Journal of Neuroscience*, 35(16), 6255–6264.
- Oberauer, K., Weidenfeld, A., & Hörnig, R. (2006). Working memory capacity and the construction of spatial mental models in comprehension and deductive reasoning. *The Quarterly Journal*

- of Experimental Psychology*, 59(2), 426–447.
- Phillips, D. C. (1995). The good, the bad, and the ugly: The many faces of constructivism. *Educational Researcher*, 24(7), 5–12.
- Roediger, H., & Karpicke, J. (2006). The power of testing memory: Basic research and implications for educational practice. *Perspectives on Psychological Science*, 1(3), 181–210.
- Ruggeri, A., Markant, D. B., Gureckis, T., & Xu, F. (2019). Memory enhancements from active control of learning emerge across development. *Cognition*, 186, 82–94.
- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: Hippocampal-midbrain encoding of overlapping events. *Neuron*, 60(2), 378–389.
- Smith, C., & Squire, L. R. (2005). Declarative memory, awareness, and transitive inference. *Journal of Neuroscience*, 25(44), 10138–10146.
- Son, J. Y., Smith, L. B., & Goldstone, R. L. (2011). Connecting instances to promote children’s relational reasoning. *Journal of Experimental Child Psychology*, 108(2), 260–277.
- Steyvers, M., Tenenbaum, J., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3), 453–489.
- Stickgold, R., & Walker, M. P. (2013). Sleep-dependent memory triage: Evolving generalization through selective processing. *Nature Neuroscience*, 16(2), 139.
- Tobias, S., & Duffy, T. M. (2009). *Constructivist instruction: Success or failure?* Routledge.
- Trabasso, T. (1977). The role of memory as a system in making transitive inferences. *Perspectives on the Development of Memory and Cognition*, 333–366.
- Trabasso, T., Riley, C. A., & Wilson, E. (1975). The representation of linear order and spatial strategies in reasoning: A developmental study. *Reasoning: Representation and Process in*

Children and Adults, 201–229.

Turner, M. L., & Engle, R. W. (1989). Is working memory capacity task dependent? *Journal of Memory and Language*, 28(2), 127–154.

Unsworth, N., Heitz, R. P., Schrock, J. C., & Engle, R. W. (2005). An automated version of the operation span task. *Behavior Research Methods*, 37(3), 498–505.

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.

Von Fersen, L., Wynne, C., Delius, J. D., & Staddon, J. (1991). Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(3), 334.

Von Glasersfeld, E. (2012). A constructivist approach to teaching. In *Constructivism in education* (pp. 21–34). Routledge.

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.

Wendelken, C., & Bunge, S. A. (2010). Transitive inference: Distinct contributions of rostrolateral prefrontal cortex and the hippocampus. *Journal of Cognitive Neuroscience*, 22(5), 837–847.

Werchan, D. M., & Gómez, R. L. (2013). Generalizing memories over time: Sleep and reinforcement facilitate transitive inference. *Neurobiology of Learning and Memory*, 100, 70–76.

Zeithamova, D., & Preston, A. R. (2010). Flexible memories: Differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. *Journal of Neuroscience*, 30(44),

14676–14684.

Zeithamova, D., Schlichting, M. L., & Preston, A. R. (2012). The hippocampus and inferential reasoning: Building memories to navigate future decisions. *Frontiers in Human Neuroscience*, 6.

Appendix A: Model selection

Mixed effects regression was used to model test accuracy and RT. The formulae for the best-fitting models are listed in Table 4. The procedure for selecting the best-fitting model was as follows. First, a baseline model was defined in each study which included relevant predictors as independent fixed effects (Experiment 1: Session, Condition, Distance, and Operation span; Experiment 2: Condition, Distance, and Operation Span). I then generated all possible combinations of interactions between these factors, leading to 113 models in Experiment 1 and 9 models in Experiment 2. Best-fitting models were selected based on the minimum AIC.

For RT models there was only one observation (median RT) for each combination of factors. As a result, the random effects structure only included random intercepts to capture between-subject variability. For the accuracy models, a random effects structure was chosen to include as many predictors (including interactions) as possible without producing convergence failures (see Table 4). In Experiment 2, including Distance as a random effect led to convergence failures, so this predictor was excluded from the fitted models for test accuracy.

	Outcome	Model formula in <i>R</i> syntax
Experiment 1	Accuracy	$\begin{aligned} \text{Correct} \sim & \text{Session} + \text{Condition} + \text{Distance} + \text{OperationSpan} \\ & + \text{Session:Condition} + \text{Session:Distance} + \text{Condition:Distance} \\ & + \text{Condition:OperationSpan} + \text{Distance:OperationSpan} + \text{Condi-} \\ & \text{tion:Distance:OperationSpan} + (1 + \text{Session} + \text{Condition} + \text{Distance} \\ & + \text{Session:Condition Subject}) \end{aligned}$
	RT	$\begin{aligned} \text{MedianRT} \sim & \text{Session} + \text{Condition} + \text{Distance} + \text{OperationSpan} + \\ & \text{Condition:Distance} + (1 \text{Subject}) \end{aligned}$
Experiment 2	Accuracy	$\begin{aligned} \text{Correct} \sim & \text{Condition} + \text{Distance} + \text{OperationSpan} + \text{Condi-} \\ & \text{tion:Distance} + \text{Condition:OperationSpan} + \text{Distance:OperationSpan} \\ & + \text{Condition:Distance:OperationSpan} + (1 + \text{Condition Subject}) \end{aligned}$
	RT	$\begin{aligned} \text{MedianRT} \sim & \text{Condition} + \text{Distance} + \text{OperationSpan} + \text{Condi-} \\ & \text{tion:Distance} + \text{Condition:OperationSpan} + \text{Distance:OperationSpan} \\ & + \text{Condition:Distance:OperationSpan} + (1 \text{Subject}) \end{aligned}$

Table 4

Formulae in R syntax for best-fitting mixed effects models in each experiment.