Journal of Experimental Psychology:
Learning, Memory, and Cognition
2010, Vol. 36, No. 6, 1452–1465

0278-7393/10/$12.00 DOI: 10.1037/a0021040

Assessing the Concreteness of Relational Representation

Jonathan R. Rein and Arthur B. Markman
University of Texas at Austin

Research has shown that people's ability to transfer abstract relational knowledge across situations can be heavily influenced by the concrete objects that fill relational roles. This article provides evidence that the concreteness of the relations themselves also affects performance. In 3 experiments, participants viewed simple relational patterns of visual objects and then identified these same patterns under a variety of physical transformations. Results show that people have difficulty generalizing to novel concrete forms of abstract relations, even when objects are unchanged. This suggests that stimuli are initially represented as concrete relations by default. In the 2nd and 3rd experiments, the number of distinct concrete relations in the training set was increased to promote more abstract representation. Transfer improved for novel concrete relations but not for other transformations such as object substitution. Results indicate that instead of automatically learning abstract relations, people's relational representations preserve all properties that appear consistently in the learning environment, including concrete objects and concrete relations.

Keywords: knowledge representation, abstraction, similarity, analogy

Relational representation is a crucial part of the human ability to transfer information across situations and may be at the core of what separates human cognition from that of other species (Gentner, Bowdle, Wolff, & Boronat, 2001; Penn, Holyoak, & Povinelli, 2008). One reason why relational representations have been central to discussions of knowledge use is that they help to solve the problem of what is meant when a representation is called abstract. That is, it is typically assumed that to allow knowledge from one domain to be used to understand a second domain, there must be sufficient abstractness to allow that representation to apply to both domains. Research on analogical reasoning has demonstrated that relational representations are particularly well suited to serve this function (Gentner, 1983, 1989; Hummel & Holyoak, 1997).

In this article, we focus on a relatively neglected aspect of relations, which is the abstractness of the relational representations themselves. We start by discussing the distinction between object and relational similarity that is central to work on analogy. Then, we define ways that relations can differ in abstractness. Finally, we present three experiments that explored the degree to which relational representations may differ in abstractness and how that may affect people's ability to use those representations to process relational similarities.

Object and Relational Similarity

Models of analogy distinguish between the objects (or entities) in a domain and the relations that bind those objects together (Gentner, 1983; Markman, 1999). For example, the statement

The dog chases the yellow cat (1)

can be represented as

CHASES(dog, cat) & YELLOW(cat). (2)

In this case, CHASE(x, y) expresses a relation that can hold between pairs of objects (that fill the roles held by the variables x and y). Descriptive attributes of objects, like YELLOW(x) help to specify properties of the objects themselves.

Research on analogy has demonstrated that people are quite good at forming comparisons across domains that preserve relational similarities among objects that have different attributes (Gentner & Markman, 1997). The classic analogy between the atom and the solar system, for example, preserves the relational similarity that there is a central object with smaller objects that revolve around it, even though the objects in each domain (e.g., the sun vs. the nucleus of the atom) do not look at all alike.

It is often assumed that these relational similarities are what allows knowledge to be used abstractly (Gentner, 1983; Sloutsky, Kaminski, & Heckler, 2005). When a comparison preserves only the relations, knowledge from one domain can be used to understand a second domain that shares no object similarities with it. Indeed, studies of relational transfer across domains have often shown that object similarities actually get in the way of focusing on the relations. In some cases, this work has shown that significant information about objects in a base domain may make it hard for people to represent the relevant relational information (Sloutsky et al., 2005; Son & Goldstone, 2009). For example, Son and Goldstone (2009) found that when people get diagrams or simulations as examples of a complex process, they are better able to focus on the relevant relations when the diagrams are schematic (using circles and blobs to depict objects) than when they have more iconic pictures of the objects.

In other cases, this research has shown that misleading object similarities may cause people to put objects in correspondence in

This article was published Online First October 4, 2010.

Jonathan R. Rein and Arthur B. Markman, Department of Psychology, University of Texas at Austin.

Correspondence concerning this article should be addressed to Arthur B. Markman, Department of Psychology, University of Texas, Austin, 1 University Station A8000, Austin, TX 78712. E-mail: markman@psy.utexas.edu

ways that defy the relational similarity (Gentner & Toupin, 1986; Ross, 1989a). For example, Gentner and Toupin (1986) asked children to retell stories they had heard with a different set of characters. Thus, the children had to preserve the relationships among the character roles in the retelling, despite changes to the characters themselves. Gentner and Toupin found that children had particular difficulty with this task when the characters were cross-mapped from one story to the next. In a cross-mapping, objects that look similar play different roles in the relational structure in each situation. For example, in the original story, there might be a squirrel hero and an eagle villain. In a cross-mapping, the child would try to retell the story with a hawk playing the hero and a chipmunk playing the villain. Thus, children had to actively ignore the object similarities to succeed. Research suggests that cross-mappings are also quite difficult for adults (Markman & Gentner, 1993; Ross, 1989a).

On this view, then, the ability to focus on relations and to ignore object similarities is the crucial determinant of transfer from one domain to another. That is, this approach assumes that abstraction is basically determined by the ability to use relational similarities rather than object similarities in comparisons of pairs of items.

There is also some evidence, primarily from linguistics, that relations may also differ in their degree of abstractness. Verbs typically express relational meanings, and they can differ in the degree of constraint that they place on the action or relationship being described (Gentner, 1975; Miller & Fellbaum, 1991). For example, the verb *transfer* specifies only a change in possession of an object from one person to another. In contrast, the verb *trade* adds additional relations to describe the more concrete situation in which one item is transferred from one person to another contingent on a second item being transferred from the second to the first. A second way that verbs can differ in concreteness is when one includes selection restrictions on the arguments of the verb. So *trade* differs from *buy* because, in the buy relation, one of the objects being transferred must be money, while, in the trade relation, there is no such restriction. Just as the hero role can be filled with various concrete objects such as squirrels and hawks, the abstract relation transfer can be embodied in various concrete relations such as trade and buy.

In this article, we are particularly interested in the prospect that the abstractness of a relational system influences the ability to apply that system to a new situation. Our method is an analogue to studies of analogical reminding (e.g., Gentner, Rattermann, & Forbus, 1993; Markman, Taylor, & Gentner, 2007; Ross, 1987, 1989a; Wharton et al., 1994), employing visual stimuli instead of the traditional semantic stimuli. Studies of analogical reminding typically manipulate whether a retrieval cue shares object similarity and relational similarity with items in memory. These studies often demonstrate that object similarity has a greater influence on retrieval than does relational similarity. In our studies, we manipulated object and relational similarity to examine the effects of these forms of similarity on the access and use of prior relational information. In addition, we examined learning factors that may influence the abstractness of the relations.

Task Description and Predictions

A brief overview of the design of the forthcoming experiments is useful to ground our predictions. Full detail is presented in the Method section. Figure 1 illustrates the basic design. Participants had to learn to identify four different visual patterns each composed of two shapes presented in a grid. Each pattern had a label that described the pattern in abstract relational terms (e.g., "vertical line" and "horizontal line" for the patterns in Figure 1). In an initial identification training phase, each pattern was always composed of the same two shapes, and participants learned to associate the patterns with the labels.

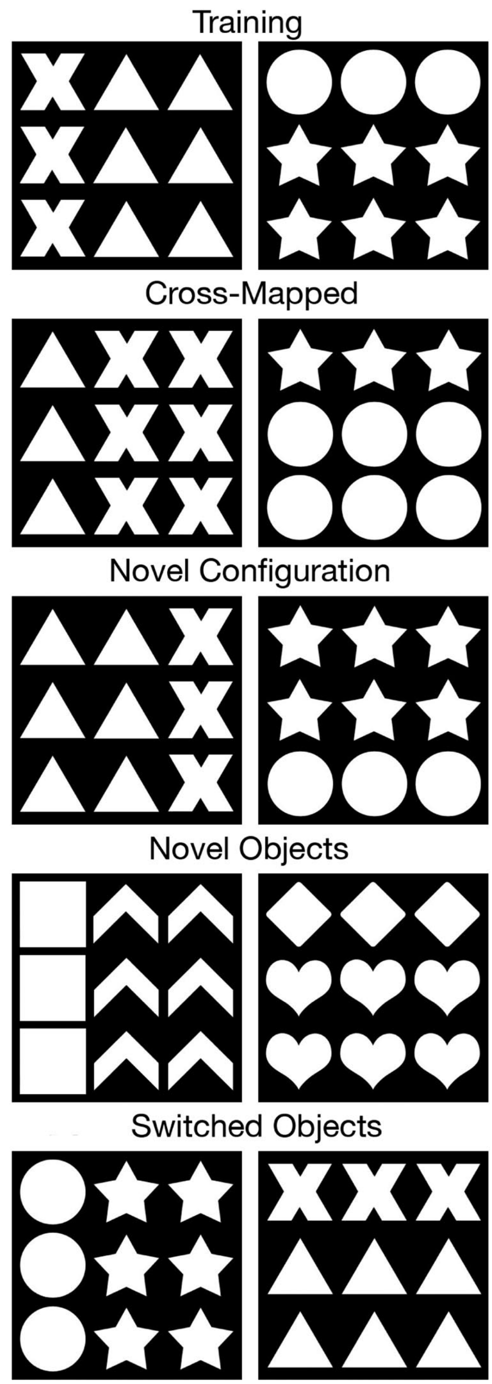

Figure 1. Examples of the five types of transfer stimuli. The bottom four are transformations of training stimuli.

A transfer phase followed in which participants made the same four identification responses to pattern stimuli that were transformations of the training stimuli. These transformations all retained abstract relational similarity to the original patterns but differed in their object similarity and the form of relational similarity to the training patterns.

There were five different transfer types, as depicted in Figure 1. The *training* stimuli were identical to those presented during the training phase. *Cross-mapped* stimuli contained the same object types for each pattern as in the training phase, but the roles of the two object types were reversed (Gentner & Toupin, 1986; Markman & Gentner, 1993; Ross, 1989b). In other words, the object that had previously been the figure became the ground, and vice versa. *Novel configuration* stimuli contained the same object types for each pattern as in the training phase, but the pattern was shifted or rotated. Training and novel configuration stimuli shared the same abstract relation and concrete objects but differed in their concrete relation.[1] In the example in Figure 1, the vertical line has shifted rightward from its original position in the training stimuli. It is still a vertical line in the most abstract sense, but the concrete form that the line takes has changed. Novel configuration stimuli were like cross-mapped stimuli in the sense that they maintained the same objects but disrupted the object–role bindings from the training stimuli. The crucial difference between these stimulus types is that cross-mapped stimuli contained the same concrete relation as training stimuli (e.g., left-justified vertical line), while novel configuration stimuli contained a different, novel concrete relation (e.g., right-justified vertical line). The difference between these two types of stimuli is critical. If participants have greater difficulty generalizing to novel configuration stimuli, this suggests that they encode the training stimuli as concrete relations, despite the instructions' emphasis on abstractness. We return to this point throughout the article. *Novel objects* stimuli contained object types that were not present during the pattern presentations of the training phase. *Switched objects* stimuli contained object types that had appeared in a different pattern during the training phase. In the example in Figure 1, the horizontal-line pattern is composed of the objects that had previously formed the vertical-line pattern.

Figure 2 presents schematic illustrations of predictions of several plausible mechanisms of relational memory, the majority of which do not contain concrete relational representations. This figure presents the predictions for the relative speed of correct identification of patterns for each type of transformation. We discuss these mechanisms in order of increasing complexity.

Figure 2A presents the response times that would be expected if people learn to identify the patterns based on pure relational abstractions. In this case, a stimulus would be encoded as an instance of the vertical-line relation. Transfer stimuli would be evaluated in an identical manner and compared to the training stimuli on this most abstract level. This form of clever indexing is embodied in some case-based reasoning models (e.g., Kolodner, 1984). Because all of the transfer patterns have one of the four abstract relations, they should all be rapidly assessed and responded to with equal ease.

Figure 2B presents the response times that would be observed if responses were based purely on abstract relations and object similarity. This is the simplest case of superficial features (objects) influencing transfer. Under this mechanism, all patterns that have the same objects as corresponding training stimuli with the same relation will be highly similar and processed most easily.

Figure 2C presents the response times that would be observed if pattern identification is governed by an overall sense of similarity, without regard to the distinction between object and relational similarity (Holyoak & Koh, 1987).[2] This subjective similarity would presumably be some weighted average of object similarity and relational similarity, possibly including information about the role that objects play within that relational structure.

Figure 2D presents the response times that would be observed if identification is facilitated by having similar objects playing similar roles in each situation, as in some mapping models (e.g., Falkenhainer, Forbus, & Gentner, 1989). On this view, a comparison is made between two relational systems, represented as full networks of objects and relations. This comparison is easiest for training and novel configuration stimuli because the objects play identical roles in identical relational systems to the original training stimuli. Novel objects and switched objects stimuli are less similar to the training stimuli because they have different objects, but they do not violate previous role bindings because those objects had not previously appeared in the relation. In contrast, cross-mapped stimuli are difficult to map because they contain familiar objects in the wrong relational roles. It is more difficult to break the original object–role bindings than to simply insert novel objects. This finding is akin to seeing a nature show for the first time and being much less confused seeing a lion chasing a gazelle (unfamiliar objects) than seeing a cat chasing a dog (familiar objects with reversed roles).

Figure 2E presents the response times that would be observed if similarity is determined jointly by object similarity and the roles those objects play in relations. In this case, difficulty will track the number of object correspondences. Objects with identical objects/features in identical roles within the relation, or *matches in place,* will contribute more to similarity than objects with identical features in different roles, *matches out of place* (Goldstone, 1994). For example, training stimuli have nine matches in place. For the vertical-line relation, novel configuration stimuli have three matches in place and six matches out of place, compared to the training stimulus. Cross-mapped stimuli have zero matches in place and nine matches out of place. Novel objects and switched objects do not have any matches, in or out of place, so they are substantially less similar.

Finally, Figure 2F presents the predictions of a related view in which similarity is most affected by how well objects and roles match a previously seen instantiation of the relation. In this case, the relative difficulty of transfer items is what was predicted in Figure 2E, except that the novel configuration stimuli are processed with greater difficulty because they do not match the

[1] We use the terms *concrete relation* and *configuration* interchangeably. Configuration is simply an intuitive way to think of relational concreteness for the stimuli used in this article.

[2] Subjective similarity ratings were obtained separately from 10 participants. These participants received a training phase identical to Experiment 1. They also received a single transfer block. On each transfer trial, participants pressed the keyboard button that corresponded to the pattern, then gave a numeric rating from 1 to 7 indicating how similar that pattern was to its instantiation during training.

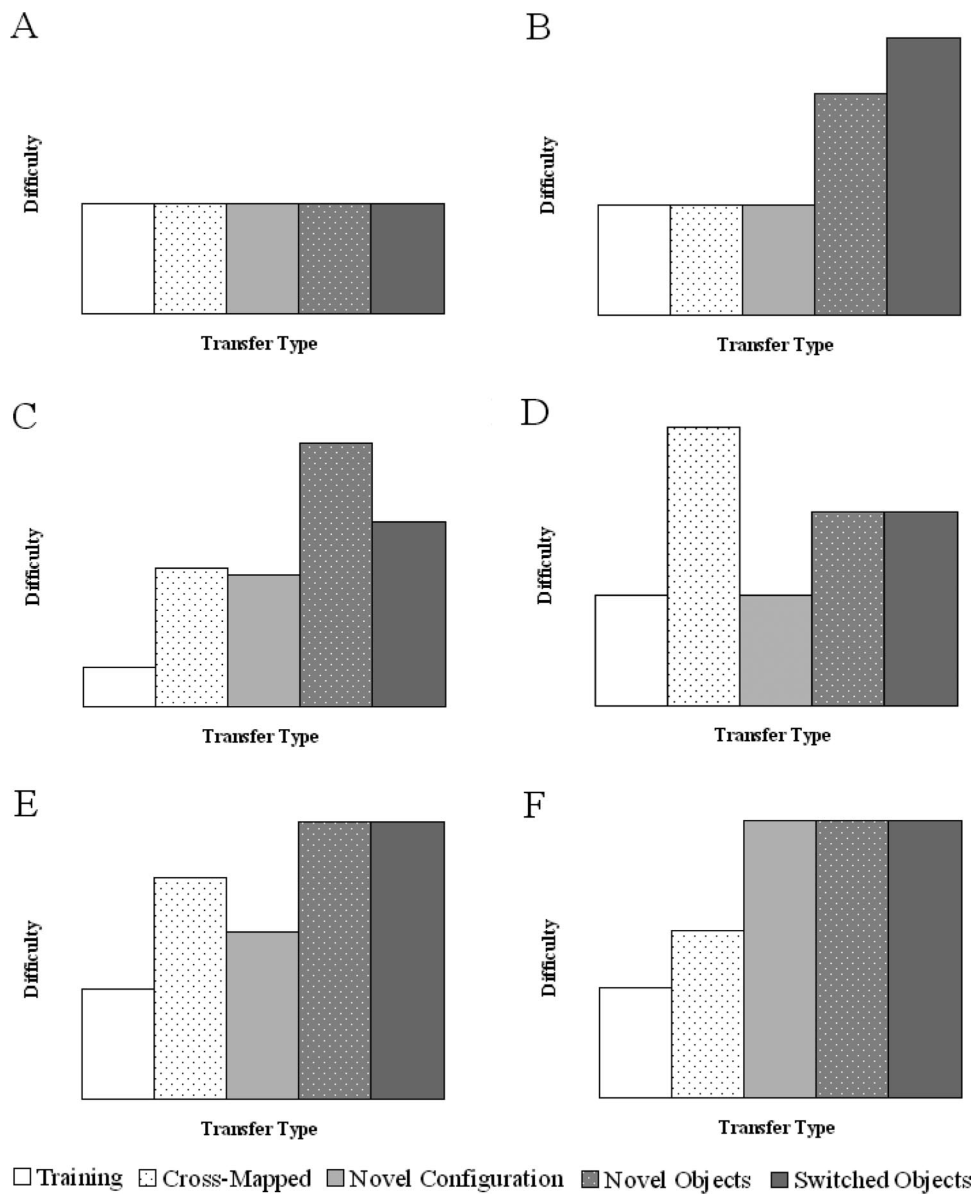

Figure 2. Schematic predictions of response time data during transfer, according to six plausible theories. A: Complete abstraction. B: Object similarity. C: Overall similarity. D: Object–role binding. E: Object correspondence. F: Object correspondence to specific configurations.

specific, concrete representation of the training patterns. Matches in place and matches out of place are similarly critical, but the definition of *match* is different. Specifically, comparison is made with reference to how the roles are situated in the concrete relation of the training stimuli. For instance, in Figure 1, the relation embodied in the training stimulus would not be represented abstractly as a vertical line but concretely as a left-justified vertical line. That is, the concrete relational structure places more constraint on what counts as an instance of this pattern. Consequently, an X that appears at the top of the line in the novel configuration stimulus is not considered a match in place because the entire line and all of its roles have been altered.

So, this study aimed to explore the contributions of object similarity, object–role binding, and relational concreteness to the ability to identify new pattern stimuli. The five types of transfer items systematically assessed the contribution of this information on performance. Most importantly, a decrement in performance for novel configuration stimuli would be evidence of concrete relational representation.

In addition to establishing the role of concrete relations in the representation of relational systems, we were also interested in how these relations might transition from concrete to abstract. One plausible determinant of the degree of abstraction in relations is the nature of the learning experience. This article investigates how different degrees of variability in the training input affect the relative abstractness of relational representation in memory. In Experiment 1, we presented each pattern with a unique set of objects in a single configuration throughout training. In Experiments 2 and 3, we presented the patterns in multiple configurations, assessing how variation during training affected the abstractness of learned relations.

Previous research has shown that multiple examples of a relation support more relational thinking that is less bound to superficial elements of individual examples (Catrambone & Holyoak, 1990; Gick & Holyoak, 1983; Kotovsky & Gentner, 1996; Loewenstein, Thompson, & Gentner, 1999). The studies in this article explored the limits of abstraction from multiple examples. Some computational models of relational learning predict that if the training examples have common surface similarities, they remain bound to the relational representation (Doumas, Hummel, & Sandhofer, 2008; Winston, 1982). If this principle holds, then presentation of a single configuration during training would result in configuration-specific representation of the stimuli, leading to a data pattern similar to Figure 2F. In contrast, presentation of multiple configurations during training would promote a more abstract mnemonic representation, making novel configuration transfer stimuli easier. However, varying training in this way would not affect the influence of object dissimilarity on performance. We further discuss the role of training variety in relational learning following Experiment 1.

Experiment 1

In this experiment, each participant went through a short training phase in which he or she learned the correct button response for each of four patterns. Throughout training, each presentation of a particular pattern contained the same two objects in the same configuration. After training, participants identified the novel pattern stimuli, transformed from the training stimuli as in Figure 1.

Method

Participants. There were 34 participants recruited from the University of Texas at Austin community. They participated for either credit in an introductory course or payment of $8. The participant pool was composed of people in their late teens to early 20s, was nearly balanced on gender, and was ethnically and racially diverse. All participants had normal or corrected-to-normal vision.

Materials. There were four total training stimuli, one for each pattern. Each of these four stimuli was composed of two unique objects (eight total). For example, a particular participant would always see the vertical-line pattern composed of Xs and triangles. Assignment of objects to patterns was randomized across participants. There were 20 total transfer stimuli, one for each pattern in each transfer type. Transfer stimuli were constructed as in Figure 1. The complete set of objects and patterns is shown in the Appendix.

All stimuli were presented on a 50.8-cm computer screen with a resolution of 1,024 $\times$ 768 pixels. Object patterns were constructed as in Figure 1. Each pattern was formed by the placement of two objects in a nine-object, 3 $\times$ 3 grid against a black background. There were four total patterns that, by taking one object as the figure and the other as the ground, could be verbally described as a horizontal line, a vertical line, a diagonal line, and a T shape. Each object was a simple, white, geometrical figure, subtending a maximum of 125 $\times$ 125 pixels (14.06 cm^2, 4.3° of visual angle). There were 60 pixels (1.8 cm) between each object; the entire 3 $\times$ 3 grid subtended 495 $\times$ 495 pixels (220.52 cm^2, 16.9° of visual angle).

The main advantage of these stimuli is their simplicity. The patterns embody intuitive abstract relations that are easily learned. They can also take multiple concrete relation forms and allow straightforward object substitution. Their simple visual form allowed us to attribute effects to the structural manipulations of interest, without the intrusion of more complicated factors such as semantic content and reading comprehension.

Procedure. Each participant was greeted by an experimenter and directed to sit at one of eight computer stations in an isolated room. Each computer station had barriers to reduce distractions. Participants sat approximately 50 cm from the monitor. All instructions were presented on the computer.

Training. Prior to the training phase, participants were told that they would be seeing visual patterns composed of two objects and would be identifying each pattern with a unique button response. They were given the verbal description for each pattern ("horizontal line", "vertical line", "diagonal line", and "T shape") and its corresponding keyboard button. Verbal descriptions were provided to ensure that participants represented the patterns with the intended abstract relation. This explicit instruction was crucial. If participants were free to represent the patterns with multiple unspecified relations, then it would be difficult to interpret their transfer behavior. The verbal descriptions were maximally abstract. The transfer phase was therefore a conservative test for the influence of concrete elements. Participants were also explicitly given the response key for each stimulus to reduce errors. They were further instructed to focus initially on being as accurate as possible. They were instructed that after they were able to identify the items reliably, they should focus on responding as quickly as possible while maintaining high accuracy.

Each trial began with a 200 $\times$ 200–pixel red fixation cross for 500 ms, followed immediately by the presentation of the pattern in the center of the screen. When the participant hit one of the four response buttons, he or she immediately received auditory and visual feedback. A buzzer or whoop sound played, "Incorrect" or "Correct" appeared above the pattern, and the correct response was given beneath the pattern in red or blue for incorrect and correct responses, respectively. Any response other than the response designated in the instructions for the pattern was considered incorrect. Feedback remained for 2 s, followed by a solid black break screen for a 750-ms intertrial interval.

The training phase was separated into blocks of 40 trials. Each pattern appeared 10 times within each block. The order of appearance of the patterns was determined randomly and uniquely for each participant. At the end of each block, participants were informed of their accuracy and mean response time for that block. There were two training blocks total.

This robust training phase was crucial. Performance in the transfer phase required that these novel patterns were well learned and had strong mnemonic representations.

Transfer. Prior to the transfer phase, participants were told that they would be making the same responses to the same patterns but that the exact appearance of the patterns could be different from what they had experienced earlier. They were instructed to respond as fast as possible while maintaining a high level of accuracy. As in training, each trial began with a 500-ms fixation cross followed by the pattern. After the response was made, the 750-ms intertribal interval began. There was no trial-by-trial feedback during transfer.

The transfer phase was separated into blocks of 40 trials. Each pattern appeared 10 times within each block, twice in each of the five transfer types. Again, the order of stimuli was determined randomly. There were five total transfer blocks (200 trials total). As in training, participants were told their accuracy and mean response time following each block. This block feedback was given to encourage task vigilance and maintain the level of performance reached during training. Although it is possible that this feedback promoted some amount of learning during transfer, such learning was expected to be relatively small and uniform across transfer types (see below for a full discussion of learning effects during transfer).

Results

Accuracy scores lower than 50% on switched objects trials or lower than 50% across all transfer trials were considered evidence of misunderstanding instructions or poor learning. For instance, some participants consistently responded to switched objects stimuli with the response associated with the objects instead of the patterns. There were four such people, who were dropped from subsequent analyses.

For all analyses, we took the median response times of correct trials for each subject and then performed analyses on the means of these medians (Ratcliff, 1993). By the end of the second training block, the mean accuracy was 96.3%, and the mean median response time for all correct responses was 716 ms. The high level of accuracy and speed suggests that the patterns were learned quickly and were easily distinguishable. All of the remaining analyses in this section focus on the transfer phase.

Response times. Figure 3A depicts the mean of individuals' median response times for each transfer type. The effect of transfer type was qualitatively similar for all stimulus patterns throughout all five blocks, so we collapsed across pattern and block. As the graph indicates, there was an overall effect of transfer type on response times, $F(4, 116) = 22.82$, $p < .001$. Mean response times for training trials ($M = 746$ ms) were lower than response times for cross-mapped trials ($M = 805$ ms), $t(29) = 3.53$, $p = .001$, $d = .64$. Response times were lower for cross-mapped trials than novel configuration trials ($M = 866$ ms), $t(29) = 4.08$, $p < .001$, $d = .75$. Novel configuration trials did not reliably differ from novel objects ($M = 892$ ms), $t(29) = 1.14$, $p = .26$; nor did novel objects trials differ from switched objects ($M = 912$ ms), $t(29) = .86$, $p = .40$. Novel configuration response times were somewhat lower than switched objects, $t(29) = 2.09$, $p = .05$, $d = .38$, though this is not statistically reliable when corrected for multiple comparisons. The overall data pattern closely matches the predictions of the concrete relations hypothesis displayed in Figure 2F.

Error rates. Figure 3B depicts the error rate for each transfer type, collapsed across patterns and transfer blocks. In general, the error rate for each transfer type was relatively low. There was an overall effect of transfer type on error rate, $F(4, 116) = 4.59$, $p = .002$. A Scheffé test confirmed that there were fewer errors for training and cross-mapped responses than for novel configuration, novel objects, and switched objects responses, $F(4, 116) = 13.48$, $p = .01$. Novel objects error rates did not differ from novel configuration or switched objects, $F(4, 116) = 4.67$, $p = .33$.

Discussion

The difficulty of identifying the patterns clearly differed across transfer stimuli. These differences existed even though all of the pattern transformations preserved the abstract relations from training.

The relative difficulty of the transfer types sheds light on the memory processes and representations underlying performance. Object similarity improved performance, as displayed by the relatively slower responses to novel and switched objects stimuli compared to the other items. However, the poor correspondence between Figure 3A and Figure 2B illustrates that object similarity as a unitary construct does not capture the full data pattern. For example, the training and cross-mapped stimuli were composed of the same objects, but their respective response times differed significantly. This finding indicates that the patterns were encoded during training as objects bound to relational roles. Stimuli that violated the learned object–role bindings were identified more slowly than stimuli that preserved these bindings. While this violation of object–role binding had a significant effect on performance, it was not as drastic as predicted by some mapping models, as in Figure 2D. Previous work has shown that cross-mappings have a dramatically negative effect on analogical mapping (Gentner & Toupin, 1986; Markman & Gentner, 1993). In contrast, while the adverse impact of cross-mappings in the current study was relatively strong in terms of effect size, it was relatively weak in comparison to the other transformations, notably, changes in the concrete relation.

Object and object–role matches both influenced response times but not exactly in the manner predicted by the SIAM (similarity, interactive activation, and mapping) model (Goldstone, 1994; Goldstone & Medin, 1994), represented in Figure 2E. In particular, response times to novel configuration stimuli were markedly slower than to cross-mapped stimuli, despite having a greater number of object–role correspondences (matches in place) with the training stimuli. These data suggest that patterns were encoded during training in the specific configuration in which they were presented. Although cross-mapped stimuli had more violations of object–role bindings, they maintained the same configuration and were therefore identified more easily than novel configuration stimuli.

Most of the theories depicted in Figure 2 did not predict the difficulty of novel configuration items. Those theories embody the assumption, prevalent in research in similarity and relational memory, that relational systems are represented exclusively by concrete objects and abstract relations. In contrast, we argue that configurations involve both objects and concrete relations. The notion of concrete relations is not incompatible with these theories, but it is not an assumption that is typically incorporated into them. Given a single training stimulus for each pattern, participants' abstraction was limited. Their representations of the patterns retained specific object, object role, and concrete relation information. This limited abstraction led to relatively constrained generalization.

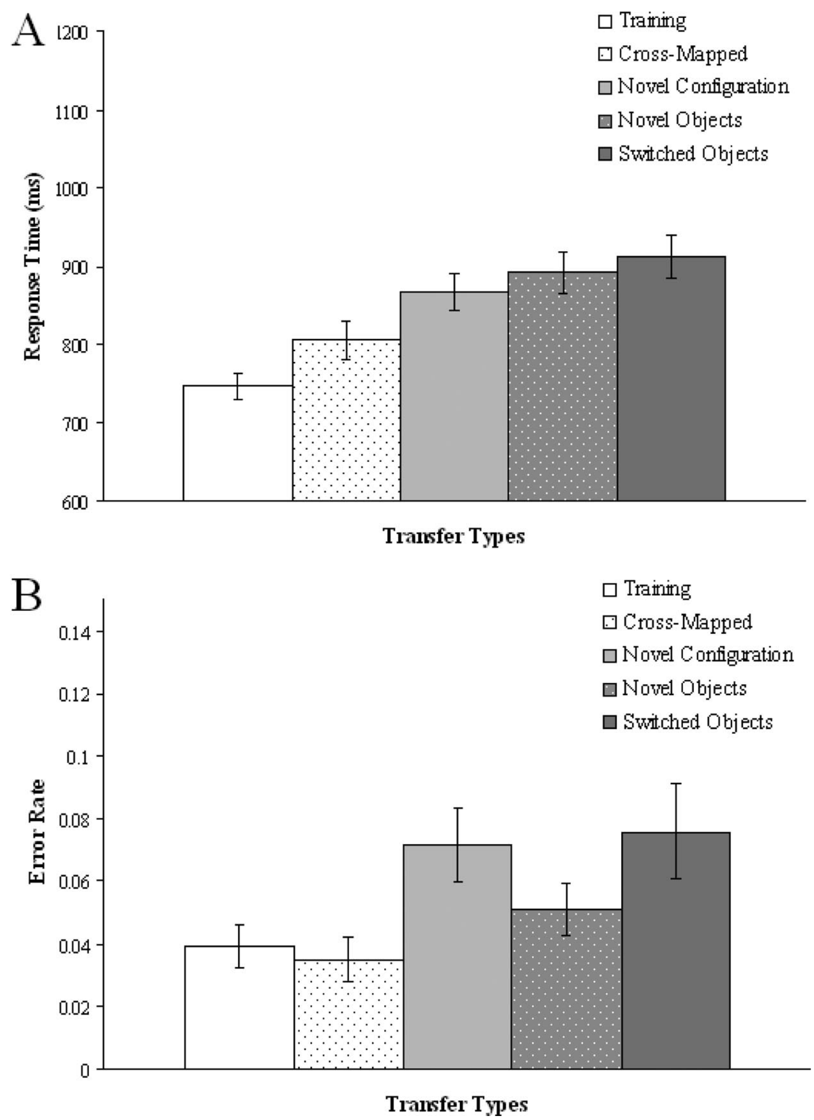

Figure 3. A: Mean response times for the five transfer types in Experiment 1. B: Mean error rates for the five transfer types in Experiment 1. Error bars indicate ±1 *SEM*.

One might argue that participants did indeed represent the training patterns abstractly, just not exclusively with the relations we intended. For example, one could represent the vertical-line pattern as a vertical line composed of any given object left from a rectangle of any given different object. If this were the case, novel configuration stimuli would be interpreted as completely novel relations or would provide additional information to disambiguate which relation was correct. If so, this would undermine the claim that the novel configuration effect is evidence of concrete relational representation. However, this alternative explanation is unlikely. We explicitly defined the relations in the instructions in the most general abstract terms. Considering the extremely low error rates, participants appear to have correctly interpreted these instructions and successfully identified the relations in all of their transfer forms, albeit at different speeds. Furthermore, the low error rates and response time pattern depicted in Figure 3 persisted throughout transfer. If participants had realized the intended relation only at transfer, there would have been sharp changes in performance across transfer blocks. This persistent high performance throughout transfer also indicates that the feedback provided after each transfer block did not lead to qualitatively different new learning in any substantial way.

Experiment 2

One interpretation of the influence of concrete relations is that abstraction is simply difficult and limited. A second possibility is that learning was efficiently optimized to the training set; participants represented all of the information that reliably co-occurred with a pattern, including concrete objects and relations. If this were the case, a training regimen with more diverse elements could promote greater abstraction. This would result in improved generalization during transfer, as compared to Experiment 1.

Other work has shown that the relative flexibility exhibited in transfer performance is influenced by the variability of the training environment. Kotovsky and Gentner (1996) showed that children's ability to make abstract relational matches (e.g., of symmetrical patterns like XYX) were better when they had been trained with multiple dimensions (e.g., size and color) than with a single dimension. Gerken (2006) showed that infants who are exposed to artificial grammars with unique words in each stimulus will generalize purely relationally (e.g., Marcus, Vijayan, Bandi Rao, & Vishton, 1999). However, if a single word reliably occurs in the same position within the grammar, they will not generalize to stimuli that have the grammatical relation but not the diagnostic word. Similar effects occur in nonrelational settings. For example, during classification learning, people may learn separate features or combine them into one holistic feature, depending on how those features co-occur (Schyns, Goldstone, & Thibaut, 1998; Schyns & Rodet, 1997). Similarly, the number of dimensions of variability determines whether people learn the regularities of whole objects or only their values on a single dimension during statistical learning (Turk-Browne, Isola, Scholl, & Treat, 2008).

In this study, we explored the effects of stimulus variability during training on generalization during transfer. To do so, we present two configurations for each pattern during training (e.g., vertical line left-justified and right-justified). As in Experiment 1, each pattern was composed of two objects in both configurations. The only additional source of variability was the number of configurations.

Figure 4 depicts three possible data patterns. One possibility is that an additional training configuration will simply be another concrete memory representation that does not promote any additional abstract processing (see Figure 4A). This would lead to a pattern similar to that in Figure 2F and obtained in Experiment1. Another possibility is that training on multiple configurations will create a more abstract relational representation, one that retains the specific objects encountered during training (see Figure 4B). This is in line with data from children (Cohen & Oakes, 1993; Sheya & Smith, 2006) and computational models that learn abstract relations from examples by eliminating nonoverlapping surface elements while retaining common elements, including surface similarities like objects (Doumas et al., 2008; Winston, 1982). A third possibility is that multiple training examples will create a relatively broad, abstract schema (Schank, 1982) that covers all instances of similar relations (see Figure 4C).

Method

Participants. There were 34 participants recruited from the University of Texas at Austin community. They participated for either credit in an introductory course or payment of $8.

Materials. Stimuli were identical to those used in Experiment 1, with a few noteworthy exceptions. The diagonal-line pattern was replaced with a square pattern. This allowed all four patterns to have at least three configurations. There were eight total training stimuli, two configurations per pattern. Again, each pattern was composed of two object types in both configurations. Transfer stimuli were constructed in an identical fashion to Experiment 1. Only one configuration was used for each pattern's training stimuli during transfer. Novel configuration stimuli were created with one of the configurations not used during training.

Procedure. The procedure was identical to Experiment 1, except that the training phase included eight unique stimuli—four patterns in two configurations. All eight were presented in random order.

Results

Again, participants whose accuracy was lower than 50% on switched objects trials or lower than 50% across all transfer trials were considered to be using some nonrelational strategy. There were four such people, who were omitted from further analysis.

By the end of the second training block, the mean accuracy was 94.8%, and the mean of individual subjects' median response times was 805 ms. Accuracy and speed remained high, though lower than performance in the training phase of Experiment 1. All remaining analyses relate to the transfer phase.

Response times. Figure 5A depicts the mean of individuals' median response times for each transfer type, collapsed across patterns and transfer blocks. As the graph indicates, there was an overall effect of transfer type on response times, $F(4, 116) = 40.16$, $p < .001$. Mean response times for training trials ($M = 775$ ms) were lower than response times for cross-mapped trials ($M = 844$ ms), $t(29) = 4.89$, $p < .001$, $d = .89$. Training response times were also lower than novel configuration trials ($M = 828$ ms), $t(29) = 3.83$, $p < .001$, $d = .70$. Both cross-mapped and novel configuration had faster response times than novel objects ($M = 886$ ms), $t(29) = 2.50$, $p = .02$, $d = .46$, and $t(29) = 3.86$, $p < .001$, $d = .70$. Novel objects response times were lower than those for switched objects ($M = 969$ ms), $t(29) = 4.70$, $p < .001$, $d = .86$. The overall data pattern closely matches the predictions displayed in Figure 4B.

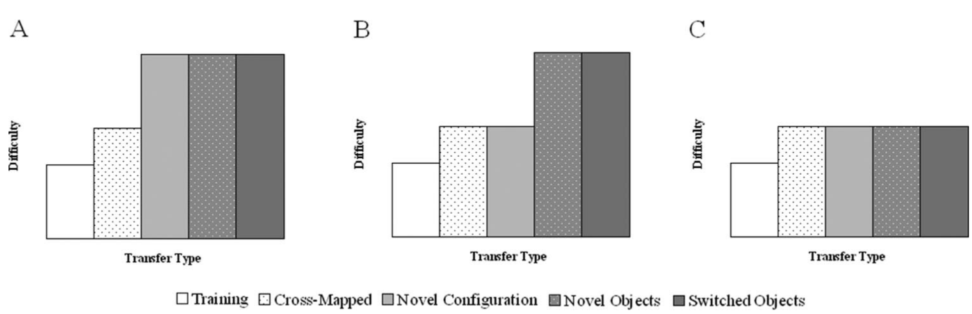

Figure 4. Schematic predictions of response time data during transfer, according to three plausible theories.

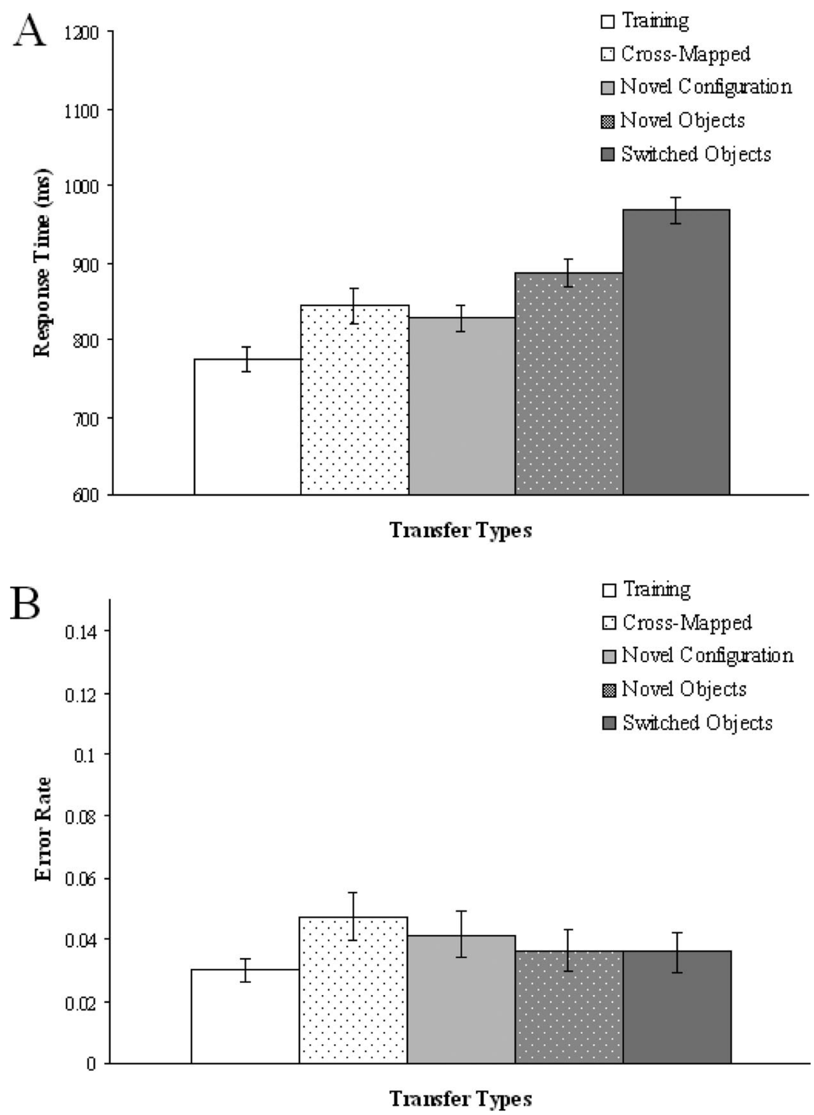

Figure 5. A: Mean response times for the five transfer types in Experiment 2. B: Mean error rates for the five transfer types in Experiment 2. Error bars indicate ±1 *SEM*.

Error rates. Figure 5B depicts the error rate for each transfer type, collapsed across patterns and transfer blocks. There was no overall effect of transfer type on error rate, $F(4, 116) = 1.59$, $p > .05$. The extremely low error rates indicate that participants learned the abstract relations and were responding accordingly.

Discussion

As in Experiment 1, the differences in response times across transfer stimuli reflect the influence of the different types of similarity. Again, participants responded more quickly to stimuli that shared objects with the training stimuli than to those that did not. This benefit of object similarity was again modulated by relational factors. The cross-mapped and novel configuration stimuli both had significantly longer response times than training stimuli, despite having identical objects.

Following the presentation of multiple configurations during training, response times were drastically reduced for novel configuration stimuli compared to novel objects and switched objects stimuli. Nevertheless, they remained higher than response times for training stimuli. That is, there was a processing advantage for the specific configurations experienced during training, but this advantage was relatively small following the variable training. The improved generalization indicates that the relational representation of patterns formed during training became more abstract. Nevertheless, the difference between training and novel configuration stimuli suggests that this representation retained concrete properties from training.

This interpretation is reinforced by the novel and switched objects stimuli, which appear unaffected by the training manipulation. The restriction of the effect to novel configurations suggests that variable training improves generalization only for the dimension that is varied. Multiple training configurations presumably created more abstract pattern representations. These abstract representations were highly similar to the novel configuration stimuli, resulting in more fluent processing. Although these representations were more abstract in the relational information, they nonetheless retained specific objects and therefore remained relatively low in similarity to novel and switched objects stimuli.

Experiment 3

Comparing the findings from Experiments 1 and 2, there appears to be a marked and isolated effect of multiple training configurations on novel configuration transfer stimuli. While the pattern is striking, we should be cautious about interpreting results across experiments. Accordingly, this experiment provided individual participants with both single- and multiple-configuration training sessions.

Method

Participants. There were 31 participants recruited from the University of Texas at Austin community. They participated for either credit in an introductory course or payment of $8.

Materials. Stimuli were identical to those used in Experiments 1 and 2, with several noteworthy exceptions. There were eight total 4×4 object patterns to allow for four distinct patterns in each of two phases. The patterns were horizontal bars, vertical bars, T shape, cross, arrow, dumbbell, ring, and alternating lines. Due to screen-size constraints, the display size of the pattern was reduced: Each object subtended a maximum of 83×83 pixels ($6.25\ \text{cm}^2$, 2.9° of visual angle). There were 40 pixels (1.2 cm) between each object; the entire 4×4 grid subtended 452×452 pixels ($183.87\ \text{cm}^2$, 15.4° of visual angle).

Procedure. There were two parts to this study. The procedure for Part 1 was identical to Experiment 1. Part 2 immediately followed Part 1 and was identical to Experiment 2. The object types used in Part 2's training stimuli were those unused during Part 1's training phase (i.e., Part 1's novel objects).

Results

Again, participants whose accuracy was less than 50% on switched objects trials or less than 50% across all transfer trials in either part were considered to be using some nonrelational strategy. There were six such people, who were omitted from further analysis.

By the end of the second training block in Part 1, the mean accuracy was 94.3%, and the mean of individual subjects' median response times was 826 ms. By the end of the second training block in Part 2, the mean accuracy was 96.0%, and the mean of individual subjects' median response times was 839 ms. These values are similar to Experiments 1 and 2. Neither accuracy nor response time differed reliably across Parts 1 and 2. All remaining analyses relate to the transfer phase.

Response times. Figure 6A depicts the mean of individuals' median response times for each transfer type in each study part, collapsed across patterns and transfer blocks. In keeping with the results from the previous experiments, there was a significant main effect of transfer type on response times, $F(4, 96) = 17.25$, $p < .001$. As with the training phases, there was no main effect of training variability. As predicted from the findings from Experiments 1 and 2, there was a significant Transfer Type ×Training Variability interaction, $F(4, 96) = 5.19$, $p < .001$. This interaction was driven by the large decrease in response times for novel configuration stimuli in Part 2.

Response times for novel configuration trials in Part 2 ($M = 899$ ms) were reliably faster than the corresponding trials in Part 1 ($M = 1{,}028$ ms), $t(24) = 2.71$, $p = .01$, $d = .54$. There was no statistically significant difference across parts for training trials (Part 1 $M = 859$ ms, Part 2 $M = 815$ ms), $t(24) = 1.57$, $p = .13$; cross-mapped trials (Part 1 $M = 875$ ms, Part 2 $M = 899$ ms), $t(24) = 0.51$, $p = .62$; novel objects trials (Part 1 $M = 1{,}014$ ms, Part 2 $M = 997$ ms), $t(24) = 0.38$, $p = .71$; or switched objects trials (Part 1 $M = 1{,}003$ ms, Part 2 $M = 981$ ms), $t(24) = 0.50$, $p = .62$. In Part 1, novel configuration trials did not differ from novel objects trials ($M = 1{,}014$ ms), $t(24) = 0.35$, $p = .73$, but were significantly slower than cross-mapped trials ($M = 875$ ms), $t(24) = 4.22$, $p < .001$, $d = .84$. In Part 2, novel configuration trials did not differ from cross-mapped trials ($M = 899$ ms), $t(24) = 0.01$, $p = .99$, but were significantly faster than novel objects trials ($M = 997$ ms), $t(24) = 2.47$, $p = .02$, $d = .49$. Overall, the data pattern in Part 1 corresponds to the predicted pattern in Figure 2F, while the pattern in Part 2 corresponds to the predictions in Figure 4B.

Error rates. Figure 6B depicts the mean error rate for each transfer type in each study part, collapsed across patterns and transfer blocks. As with the response times, there was a significant main effect of transfer type, $F(4, 96) = 7.55$, $p < .001$. There was no main effect of training variability, nor was there a significant interaction between transfer type and training variability. As in Experiments 1 and 2, error rates were low across transfer.

Discussion

As in Experiments 1 and 2, concrete transformations affected response times for transfer stimuli. Object similarity increased processing fluency, though this was again modulated by relational structure. Cross-mapped and novel configuration stimuli, which did not preserve the object–role bindings of the training stimuli, were responded to more slowly than training stimuli.

The results also confirmed the qualitative comparison of Experiments 1 and 2. Multiple configurations during training led to improved performance on novel configuration stimuli. Furthermore, this improvement in relational generalization did not apply to the other types of transfer stimuli.

General Discussion

In this article, we have explored the determinants of abstraction in relational processing. The relational stimuli used in these experiments allowed us to assess the contributions of object similarity and relational concreteness. Previous research demonstrated that an important way that relational structures are made concrete is by grounding them in specific situations by binding objects to the relational roles. The present results are certainly consistent with this mode of specifying a relational structure.

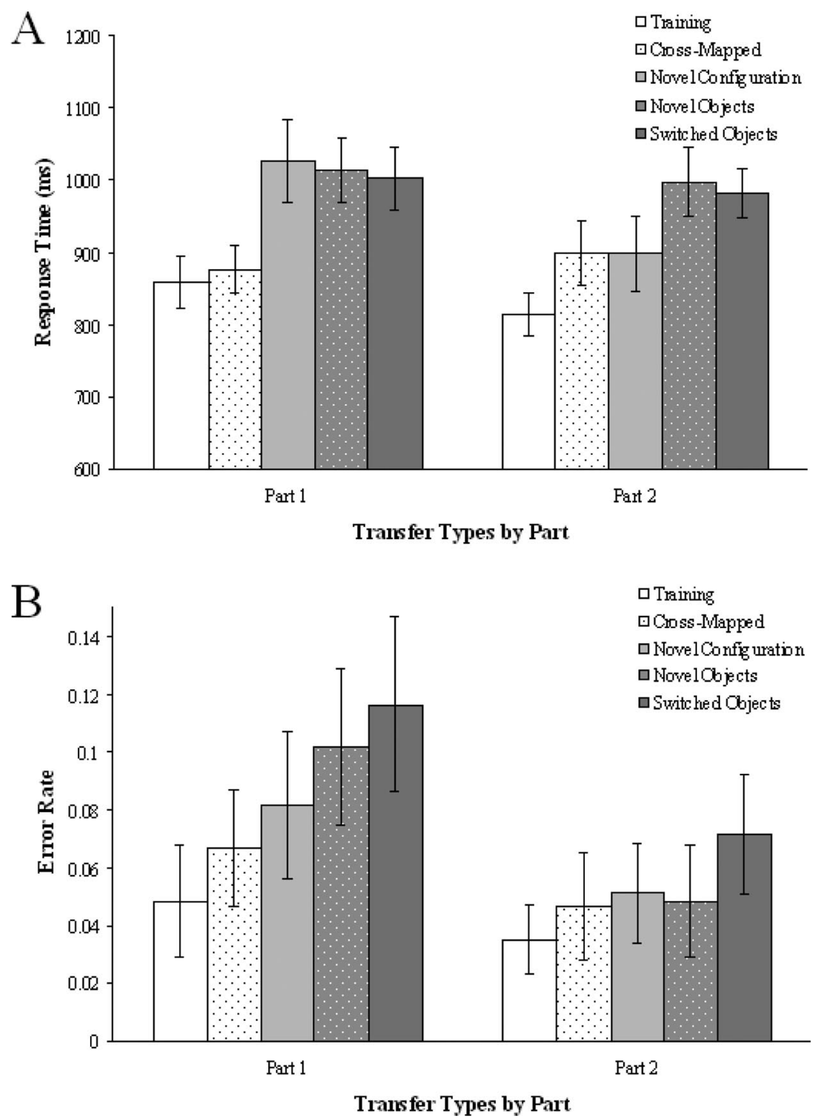

Figure 6. A: Mean response times for the five transfer types in both parts of Experiment 3. B: Mean error rates for the five transfer types in both parts of Experiment 3. Error bars indicate ±1 *SEM.*

Across all three experiments, there was a powerful effect of object similarity on response times to transfer stimuli. Transfer to items was best when the objects seen during training were present in the transfer patterns. Instantiating a relation with new objects was difficult. A similar effect of object similarity has been found in studies of problem solving (e.g., Ross, 1989a), analogical reminding (e.g., Gentner et al., 1993), and explicit comparison (e.g., Kroger, Holyoak, & Hummel, 2004). There was also an effect of object–role binding. While Ross (1989a) found no effect of object–role binding on memory for previous examples, we did find such a difference between our cross-mapped and training stimuli, possibly because the continuous nature of response time makes it a more sensitive measure than recall of probability rules in a problem-solving task. However, this effect was more modest than that found in explicit comparison or mapping tasks, in which people have more difficulty with cross-mappings (e.g., Gentner & Toupin, 1986; Goldstone, 1994, Experiment 1; Markman & Gentner, 1993).

Our findings go beyond previous research by demonstrating that the concreteness of relations themselves contributes to the concreteness of relational structures. As indicated in Figure 2, most existing theoretical frameworks would not predict the effect of configuration in Experiments 1 and 3. In general, formal models of similarity or analogy determine the degree of correspondence between two things according to shared concrete objects/features (e.g., Tversky, 1977), abstract relations (e.g., Kolodner, 1984), or both (e.g., Falkenhainer et al., 1989; Forbus, Gentner, & Law, 1995; Hummel & Holyoak, 1997, 2003), but not concrete relations.

When each relation was presented in only one form during training, concrete information about this relation was retained and slowed processing of novel forms of the relation. In other words, participants were fairly conservative when learning new relational structures. Their relational representations retained all of the constraints on relational structure that appeared in the input. This pattern of data is consistent with a broader strain of research on the influence of superficial concrete properties on the application of abstract principles, as in rule-based classification (Allen & Brooks, 1991; Erickson & Kruschke, 1998) and a variety of other reasoning tasks (Sloman, 1996).

In Experiments 2 and 3, we fostered a more abstract representation of the training patterns by presenting them in multiple configurations. We found that participants did respond to novel configuration transfer stimuli more quickly after viewing two training configurations. This finding is consistent with other studies showing that multiple instantiations of a relation or relational role increase perceived similarity (Jones & Love, 2007) and relational transfer (Catrambone & Holyoak, 1990; Gick & Holyoak, 1983; Kotovsky & Gentner, 1996; Loewenstein et al., 1999).

Although more abstract generalization occurred, it was limited. Improvement did not extend beyond novel configurations to any other transfer stimuli. This result is consistent with models that learn relations from experience by removing surface dissimilarities but retaining all commonalities, including superficial similarities (e.g., Doumas et al., 2008; Winston, 1982).

A similar pattern of limited abstraction has been observed in the cognitive development literature. Young children often name or make inferences about objects based on irrelevant surface properties instead of defining features or relations (Keil, 1989; Landau, Smith, & Jones, 1988; Sloutsky, Kloos, & Fisher, 2007). A similar bias in identifying relations based on surface properties has also been found in pigeons (Gibson & Wasserman, 2003). Although adult humans are capable of abstract thought, our data indicate that adults are still influenced by the objects bound to relations as well as the constraining relations.

Finally, this work has implications for pushing forward research on analogical reasoning. Models of analogy have made a clear distinction between objects and relations and have focused on two key representational questions. First, they have examined the role of object similarity in relational mapping and retrieval. Second, they have explored the influence of cross-mappings on the ability to form analogies. Because people have a relatively poor relational vocabulary, though, it has been difficult to study similarities among relations (e.g., Sieck, Quinn, & Schooler, 1999). This work suggests that examining the ease of comparisons following relational learning may allow researchers to better understand the formation of relational representations and the role of constraining relational structure in relational learning.

References

Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General, 120,* 3–19. doi:10.1037/0096-3445.120.1.3

Catrambone, R., & Holyoak, K. J. (1990). Learning subgoals and methods for solving probability problems. *Memory & Cognition, 18,* 593–603.

Cohen, L. B., & Oakes, L. M. (1993). How infants perceive a simple causal event. *Developmental Psychology, 29,* 421–433. doi:10.1037/0012-1649.29.3.421

Doumas, L. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review, 115,* 1–43. doi:10.1037/0033-295X.115.1.1

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140. doi:10.1037/0096-3445.127.2.107

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41,* 1–63. doi:10.1016/0004-3702(89)90077-5

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19,* 141–205.

Gentner, D. (1975). Evidence for the psychological reality of semantic components: The verbs of possession. In D. A. Norman, D. E. Rumelhart, & the LNR Research Group (Eds.), *Explorations in cognition* (pp. 211–246). San Francisco, CA: Freeman.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7,* 155–170. doi:10.1207/s15516709cog0702_3

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199–241). New York, NY: Cambridge University Press. doi:10.1017/CBO9780511529863.011

Gentner, D., Bowdle, B., Wolff, P., & Boronat, C. (2001). Metaphor is like analogy. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 199–253). Cambridge, MA: MIT Press.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52,* 45–56. doi:10.1037/0003-066X.52.1.45

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology, 25,* 524–575. doi:10.1006/cogp.1993.1013

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10,* 277–300. doi:10.1207/s15516709cog1003_2

Gerken, L. A. (2006). Decisions, decisions: Infant language learning when multiple generalizations are possible. *Cognition, 98,* B67–B74. doi:10.1016/j.cognition.2005.03.003

Gibson, B. M., & Wasserman, E. A. (2003). Pigeons learn stimulus identity and stimulus relations when both serve as redundant, relevant cues during same–different discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes, 29,* 84–91. doi:10.1037/0097-7403.29.1.84

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15,* 1–38. doi:10.1016/0010-0285(83)90002-6

Goldstone, R. L. (1994). Similarity, interactive activation, and mapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 3–28. doi:10.1037/0278-7393.20.1.3

Goldstone, R. L., & Medin, D. L. (1994). Similarity, interactive activation, and mapping: An overview. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections* (pp. 321–362). Westport, CT: Ablex Publishing.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15,* 332–340.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427–466. doi:10.1037/0033-295X.104.3.427

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110,* 220–264.

Jones, M., & Love, B. C. (2007). Beyond common features: The role of roles in determining similarity. *Cognitive Psychology, 55,* 193–231. doi:10.1016/j.cogpsych.2006.09.004

Keil, F. C. (1989). *Concepts, kinds and cognitive development.* Cambridge, MA: MIT Press.

Kolodner, J. L. (1984). *Retrieval and organizational structures in conceptual memory: A computer model.* Hillsdale, NJ: Erlbaum.

Kotovsky, L., & Gentner, D. (1996). Comparison and categorization in the development of relational similarity. *Child Development, 67,* 2797–2822. doi:10.2307/1131753

Kroger, J. K., Holyoak, K. J., & Hummel, J. E. (2004). Varieties of sameness: The impact of relational complexity on perceptual comparisons. *Cognitive Science, 28,* 335–358.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3,* 299–321. doi:10.1016/0885-2014(88)90014-7

Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin and Review, 6,* 586–597.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999, January 1). Rule learning by seven-month-old infants. *Science, 283,* 77–80. doi:10.1126/science.283.5398.77

Markman, A. B. (1999). *Knowledge representation.* Mahwah, NJ: Erlbaum.

Markman, A. B., & Gentner, D. (1993). Structural alignment during similarity comparisons. *Cognitive Psychology, 25,* 431–467. doi:10.1006/cogp.1993.1011

Markman, A. B., Taylor, E., & Gentner, D. (2007). Auditory presentation leads to better analogical retrieval than written presentation. *Psychonomic Bulletin and Review, 14,* 1101–1106.

Miller, G. A., & Fellbaum, C. (1991). Semantic networks of English. *Cognition, 41,* 197–229. doi:10.1016/0010-0277(91)90036-4

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman animals. *Behavioral and Brain Sciences, 31,* 109–130. doi:10.1017/S0140525X08003543

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin, 114,* 510–532. doi:10.1037/0033-2909.114.3.510

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 629–639. doi:10.1037/0278-7393.13.4.629

Ross, B. H. (1989a). Distinguishing types of superficial similarities: Different effects on the access and use of earlier examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 456–468. doi:10.1037/0278-7393.15.3.456

Ross, B. H. (1989b). Remindings in learning and instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 438–469). New York, NY: Cambridge University Press. doi:10.1017/CBO9780511529863.021

Schank, R. C. (1982). *Dynamic memory.* New York, NY: Cambridge University Press.

Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences, 21,* 1–17. doi:10.1017/S0140525X98000107

Schyns, P. G., & Rodet, L. (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 681–696. doi:10.1037/0278-7393.23.3.681

Sheya, A., & Smith, L. B. (2006). Perceptual features and the development of conceptual knowledge. *Journal of Cognition and Development, 7,* 455–476. doi:10.1207/s15327647jcd0704_2

Sieck, W. R., Quinn, C. N., & Schooler, J. W. (1999). Justification effects on the judgment of analogy. *Memory & Cognition, 27,* 844–855.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3–22. doi:10.1037/0033-2909.119.1.3

Sloutsky, V. M., Kaminski, J. A., & Heckler, A. F. (2005). The advantage of simple symbols for learning and transfer. *Psychonomic Bulletin and Review, 12,* 508–513.

Sloutsky, V. M., Kloos, H., & Fisher, A. V. (2007). When looks are everything: Appearance similarity versus kind information in early induction. *Psychological Science, 18,* 179–185. doi:10.1111/j.1467-9280.2007.01869.x

Son, J. Y., & Goldstone, R. L. (2009). Fostering general transfer with specific simulations. *Pragmatics and Cognition, 17,* 1–42.

Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 399–407. doi:10.1037/0278-7393.34.2.399

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327–352. doi:10.1037/0033-295X.84.4.327

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., Wickens, T. D., & Melz, E. R. (1994). Below the surface: Analogical similarity and retrieval competition in reminding. *Cognitive Psychology, 26,* 64–101. doi:10.1006/cogp.1994.1003

Winston, P. H. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence, 19,* 321–350. doi:10.1016/0004-3702(82)90004-2

Appendix

Experiment Stimuli

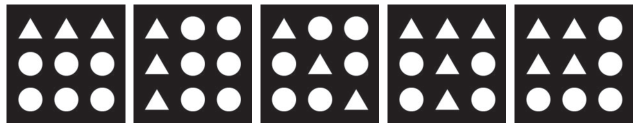

Figure A1. 3 × 3 patterns.

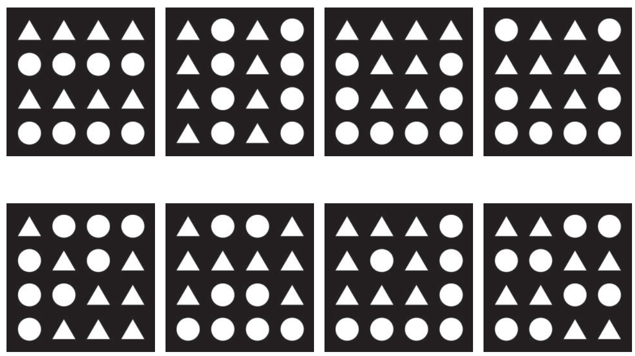

Figure A2. 4 × 4 patterns.

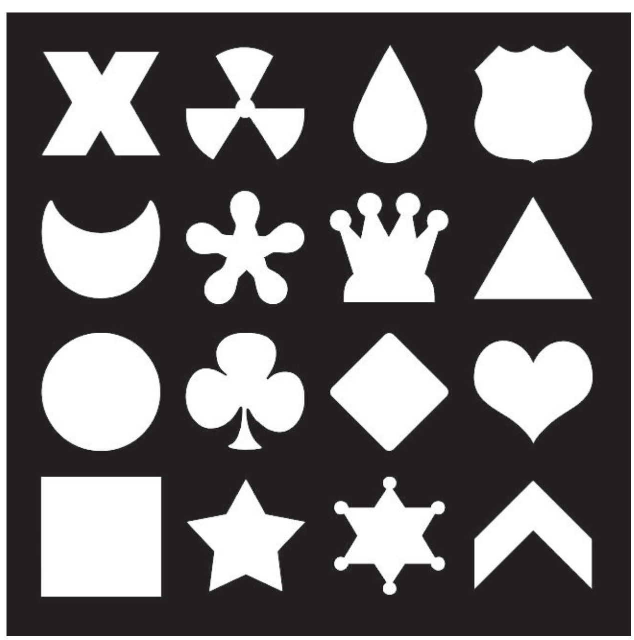

Figure A3. Object library.

Received July 22, 2008
Revision received June 16, 2010
Accepted June 21, 2010 ■