

The Effects of Complexity on Relational Recognition

Katherine A. Livins (klivins@ucmerced.edu)

University of California at Merced, Department of Cognitive Science,
5200 North Lake Road, Merced, CA, 95343

Leonidas A.A. Doumas (alex.doumas@ed.ac.uk)

University of Edinburgh, School of Philosophy, Psychology, and Language Sciences,
Psychology Building, 7 George Square, Edinburgh, EH8 9JZ

Abstract

Analogy is an important cognitive process that has been researched extensively. Functional accounts of it typically involve at least four stages of processing (access, mapping, transfer, and evaluation, e.g., see Kokinov & French, 2002), however, they take the way in which the base analog is understood, along with its relational structure, for granted. The goal of this paper is to open a discussion about how this process (which we will call “relational recognition”) may occur. To this end, this paper describes two experiments that vary the level of relational complexity across exemplars. It was found that relational recognition tasks benefit from increased complexity, while mapping tasks suffer from it.

Keywords: relational reasoning, analogy, recognition, relational complexity, mental representation

Imagine that you are a fighter pilot. As a pilot, you are highly trained in flying your jet, and you know, in detail, the mechanics of how your jet works. During a covert mission, you are stranded on foreign soil and need to find a way back home. You manage to locate a foreign plane, but its control panels look significantly different from your own, and none of the panels are labeled in a language that you understand. At first you panic, but then you remember that all planes follow the same mechanics of flight, and have some control for pitch, yaw, and roll. You reason that if you can determine which button or lever controls each one of those variables, that you will be able to fly the plane well enough to escape. In other words, even though the buttons or levers may not look like those in your own plane, they will do the same thing. After some trial and error you manage to figure out which levers and dials control those three variables, and manage to get the plane off the ground before flying to a safe location.

The human ability to make analogies is at the heart of this story, since recognizing that a lever in the first plane might be “like” a knob in the second requires one to focus on the roles that they are playing, rather than the features that each possesses. This role-sensitivity is the hallmark of analogical cognition (Holyoak, Gentner, & Kokinov, 2001; Hummel & Holyoak, 2003), and it not only allows for the identification of similarities, but also for powerful inferences to be drawn based on them. So, if you know that frenetic movements of a pitch controller may make a plane spiral out of control, then you can generalize this knowledge to any pitch controller (whether it be a knob, lever, or button).

Functional accounts of analogy-making specify the need for at least one *base* analog (the representation being mapped from) and at least one *target* analog (the

representation being mapped to). Once these are specified, the analogy-making process breaks down into four parts: retrieval, mapping, transfer, and evaluation (Kokinov & French, 2002). Thus, when one is faced with a base analog, one begins by *retrieving* potential analogical matches (targets) from long-term memory. Candidates compete, and the most likely target is mapped to the base under the constraints of the structural similarities shared between the analogs. Information is ultimately transferred from the base to the target, before the mapping is evaluated. If the analogy is judged to be appropriate, the process ends.

This general process is reasonably uncontroversial, however it does appear to take the process of understanding the base analog (and its relational structure) for granted. In other words, it assumes that one already knows what the base analog is and how to represent it when analogical processing begins.

Relational recognition (the term we will herein use for the process of recognizing the relational structure at play in a base analog) is not a trivial problem though. To the point, one must recognize a relation before one can map it to another relation, and so without recognition, the rest of the analogy-making process would not even get off the ground. However, as Gick and Holyoak (1980, 1983) have pointed out, people often fail to notice relations unless they are explicitly directed to do so, suggesting that recognition may not always take place, and that certain conditions can affect the course of learning (also see Doumas et al., 2008, for a discussion on the difficulty of learning previously unknown relations). Thus, we are interested in opening a discussion about how people may go about solving the problem of relational recognition. This paper will not be exhaustive, however, it will be a starting point with the goal of investigating how recognition is like or dislike other parts of the analogical reasoning process.

One factor that has been predominant throughout the analogy literature is relational complexity—to the point, experimental efforts have found that as relational complexity increases, analogical competency decreases (e.g., Halford et al., 1998). For example, Viskontas et al. (2004) demonstrated this trend exists across a longitudinal trajectory and Waltz et al. (1999) showed a similar trend across patients with various types of brain injuries. These studies employed relational tasks ranging from pictorial similarity mappings, to transitive inference problems, to Ravens Standard Progressive

Matrices; the trend was maintained across each paradigm.

Importantly, complexity has been defined in different ways: First one may consider the *arity* of the relations involved in the task (i.e., the number of slots the relations hold) (see Halford et al., 2005) such that processing higher arity relations are more complex than lower arity relations. For example, *bigger-than(John, Mary)* would be a lower arity, and therefore be less complex than *bigger-than(John, Mary, Sue)*. Secondly, relational complexity has been defined in terms of the number of relations that one must process simultaneously in order to deal with a given problem (Viskontas et al., 2004). For example, mapping *bigger-than(John, Mary)* to *bigger-than(Sue, Charlie)* would involve fewer relations, and therefore be less complex than mapping *bigger-than(John, Mary)* to *bigger-than(Sue, Charlie)* and also to *bigger-than(Los Angeles, Fresno)*. Interestingly though, both of these definitions boil down to the issue of how many individual elements must be bound to roles in order to process the given relation (i.e., the number of role bindings, see Doumas et al., 2008). Thus, in order to satisfy both existing definitions, this paper will define complexity in terms of the number role bindings in a given problem.

With this definition in mind, we can notice that mapping more complex relations requires one to not only keep more elements and their respective role-bindings in mind, but to make structural alignments based on those role-bindings. It has been argued that this process understandably taxes working memory (Viskontas et al., 2004; Cho, Holyoak & Cannon, 2007), so as complexity increases, performance decreases.

Not all stages of analogical processing are so structure-sensitive though, and so it may not be the case that all stages of relational processing interact with complexity in the same manner. For instance, models of retrieval suggest that retrieval (or the process of *access*) is more sensitive to object features than to relational structure. To the point, computational models such as MAC/FAC (Forbus, Gentner, and Law, 1994), and ARCS (Thagard et al., 1990) describe how access works by scanning long-term memory for objects that share features with the base analog, and many experiments (e.g., Gentner, Ratterman, and Forbus, 1993) have shown that people will remember (i.e., access) analogs that share surface similarities more than they will analogs that share structural ones. These results have lead to the widespread conclusion (e.g., Gentner, 1989, 2003) that the greater the featural (“surface”) match between a base and a target, the greater the likelihood of accessing that target. Admittedly, Gentner does not explicitly use the word “complexity” in her analysis, however, her claim does suggest that the more features in common between

analogs, the greater the ease of access. Thus, while less is more in the case of mapping, more information seems capable of boosting access.

While it may seem curious that more information can be useful in the case of access, but not in the case of mapping, remember that greater complexity is likely troublesome for mapping due to working memory: as complexity increases, the number of elements that have to be aligned and maintained in working memory increases, and so working memory is taxed and ultimately overloaded (e.g., Doumas et al., 2008; Halford et al., 1998, 2010, 2012). However, access is more focused on semantic similarities without the need to create explicit alignments, and so it is possible that a greater number of elements could carry a greater amount of semantic information, and thus promote access.

On the surface, relational recognition seems more like access than mapping. While the mechanisms of relational recognition are not yet specified, it presumably involves querying long-term memory for detected relational features, much like how access involves querying memory for objects. As a result, it seems reasonable to expect that recognition may be equally sensitive to those features, and so equally improved by greater amounts of information. The following studies aim to investigate this hypothesis.

Experiment 1

As discussed, the similarity between access and recognition suggests that increased complexity may boost relational recognition, despite the fact that it may hinder relational mapping. Essentially, the expectation is that if relational recognition is sensitive to features, then problems with greater complexity should simply provide a greater number of relational elements and so a higher concentration of relational features. A higher concentration of relational features should result in a greater probability of relational features being highlighted (e.g., Doumas et al., 2008). However, this reasoning also suggests that the specific presentation style of a relation (and not just the amount of complexity involved) should affect the way that recognition interacts with complexity. For instance, integrated relations have three or more relations engaged in the given relation. This structure means that one element is engaged in two instances of the same relation, and so fewer extraneous, object-specific features are present to distract from the relational one, thus creating a higher ratio of relation-specific features to element-specific features per relational exemplar.

Thus, based on our predictions that (i) relational recognition is similar to access, and (ii) that relational recognition will interact with the structure of a given problem such that an integrated structure will provide a better relational-feature-to-element-feature ratio, this

experiment provided participants with a relational recognition task that varied exemplars based on complexity and integrated structure.

Specifically, participants were required to recognize relations in pictorial scenes and pick the relation out of a word list. Complexity and structure were varied across exemplars that depicted binary relations (two elements involved in a single relation), "integrated relations (three elements engaged in the same relation where one element was both an actor and a patient), and multi-relational exemplars (four elements in total, broken into two groups of two, where each pair is separately engaged in the same relation).

Ultimately, if it is the case that higher relational complexity is always associated with lower relational performance (as is the case with mapping tasks), then the binary relations should possess the fastest reaction times on the recognition task. However, if a greater number of elements speeds up relational recognition, then the multi-relational exemplars should show improved reaction times. Finally, if an integrated structure is itself helpful during recognition, then the integrated relational exemplars should show the fastest reaction times, with the multi-relational exemplars showing the second fastest results.

Participants: We recruited twenty-three undergraduate participants through the psychology department at the University of Hawaii at Manoa. The participants were between 18 and 30 years of age and all had normal to corrected-to-normal vision. They were compensated with course credit for their participation.

Stimuli: Stimuli consisted of pictorial scenes adapted from Richland, Morrison, & Holyoak (2006). Each stimulus contained six objects dispersed around a black and white, drawn image; all stimuli were 720 by 450 pixels in size and presented on a black background. They all included living and non-living objects.

Each stimulus depicted one of the relational structures of interest: (i) Binary relational images were created by depicting a single actor, and a single patient involved in some relationship with a collection of distractor items (e.g., Figure 1). ii) Integrated relations were created by depicting three items involved in a nested relationship, such that one item was the patient, one was the actor on that patient, and also a patient itself for yet another actor (e.g., Figure 2). And (iii) multi-relations were created by depicting two sets of two objects involved in the same binary relation, such that there were two independent actors and two independent patients (e.g., Figure 3). Twenty exemplars of each type were created, resulting in a total of sixty stimuli.

The relational items (those that were the actors and the patients) varied in all three conditions, and the order

in which each stimulus was presented was randomly generated for each participant. All participants saw all stimuli, thus making this experiment a repeated measures design.



Figure 1: Binary relation example; *chases*(boy, cat).

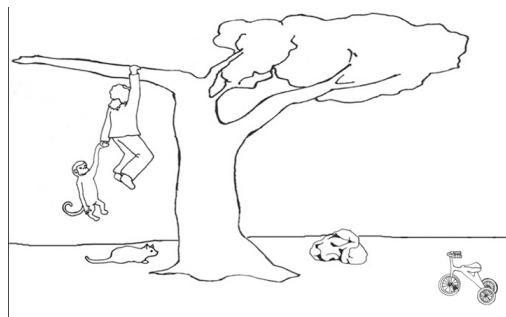


Figure 2: Integrated relation example; *hangs-from*(woman, tree)-and-*hangs-from*(monkey, woman).

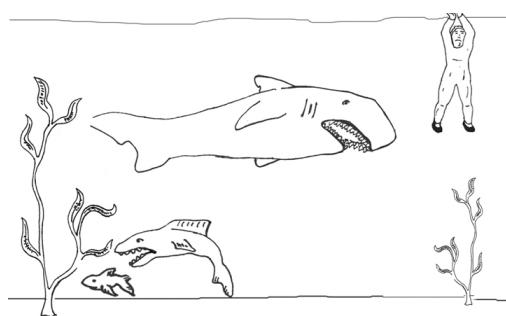


Figure 3: Multi-relation example; *hunts*(shark, human), *hunts*(fish1, fish2).

Stimuli were presented in the top center of a 1440 by 900-pixel screen, and depicted one of the following relations: *hunting*, *hanging*, *pulling*, *reaching*, *chasing*, *dropping*, *scolding*, *balancing*, *kissing*, and *talking*. The names of these possible relations were printed in text to the bottom right of each image in 22 pixel-high, Times font. The words were printed as a list, one per line, and each time a new stimulus appeared the words were randomly shuffled to new locations (in order to control for order effects). A fixation cross was placed on the

left side of the screen, across from the relational central word (see Figure 4). The cross was used as the starting point for the mouse for each trial (i.e., the mouse would automatically reposition to the cross at the point of presentation of each new stimulus).

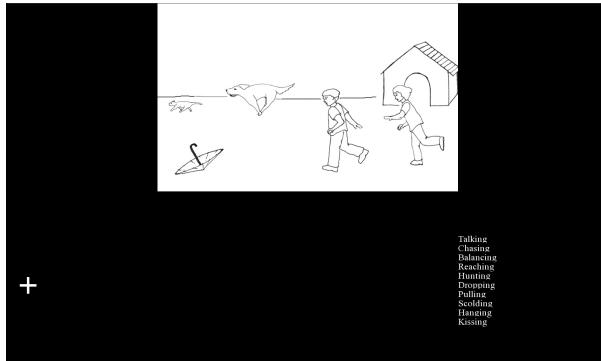


Figure 4: An example of a trial in Experiment 1.

Procedure: All participants were instructed to look at the images presented and determine which word in the word list best described the relationship depicted in it. They were also told that their chosen word should be the most central relationship to the image. Upon reading the instructions, participants were shown two training trials using exemplars and words that would not be part of the experiment itself, in order to get them accustomed to the mouse movements. Participants were then told to ready themselves for the actual experiment.

Once the experiment began, participants were self-paced, moving forward by clicking the word that they selected for each word. Clicking on a word would bring up the next stimulus and a new order of words.

Results: There was a ceiling effect across conditions on correctly identifying the relations ($M=19$, $SD=1.17$ for the integrated relations condition, $M=18.87$, $SD=1.22$ for the multiple relations condition, and $M=18.78$, $SD=1.20$ for the binary relation condition), however this result was expected given the simplicity of the task. However, a repeated measures ANOVA with a Greenhouse-Geisser correction revealed that reaction times across conditions differed significantly $F(1.33, 29.318) = 13.902$, $p < .01$. Post hoc testing with a Sidak correction showed that participants were significantly faster ($p < .01$) on the integrated exemplars (4.01 ± 0.64 sec) than they were on the multi-relational exemplars (4.26 ± 0.65 sec), and that they were also significantly faster ($p < .05$) on the multi-relational exemplars than they were on the binary exemplars (4.63 ± 1.07) (see Figure 5)¹.

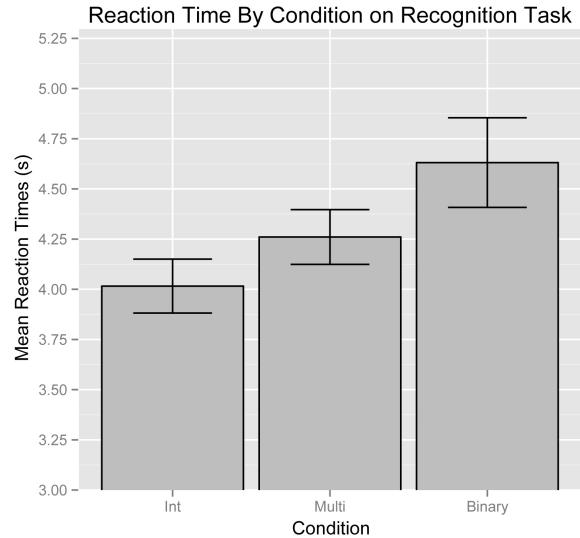


Figure 5: Results from Experiment 1, showing reaction times in seconds by condition. Error bars represent two standard errors.

Discussion: These results are consistent with our initial predictions: more complexity produced faster performance, however the integrated relational structure produced the fastest performance. This trend is what should be seen if recognition is sensitive to a greater amount of relational feature information present in a stimulus, the ratio of featural information to element information is also important.

That said it does seem necessary to compare these results to performance on a mapping task involving the same sort of stimuli. This comparison will be useful in ensuring that our data is not anomalous and that the previous trends in complexity (i.e., those reviewed in the opening section of this paper) replicate given the same type of stimuli. Thus, we should see a decrease in performance as complexity increases, regardless of integration.

Experiment 2

Participants: Participants in experiment two were analogous to those in experiment one and included twenty-four undergraduate participants, recruited through the psychology department at the University of Hawaii at Manoa. They ranged from 18 to 30 years of age, had normal to corrected-to-normal vision, and were compensated with course credit for their participation.

Design: Like experiment one, experiment two used the pictorial images adapted from Richland et al. (2006).

¹ Note that reaction times greater than 3 standard deviations from the mean were discarded for the purposes of this

calculation and that a Greenhouse-Geisser test was used because sphericity was violated

Thus the images were black and white drawings that were 720 by 450 pixels in size. Each image possessed six elements spread over the image space and were presented on a black background. Once again, stimuli depicted relational situations involving the following relations: *kissing*, *hunting*, *hanging-from*, *towing*, *reaching*, *pulling*, *chasing*, *dropping*, *scolding*, and *balancing*. Each relation was represented in each condition; in other words, it was represented as a binary relation involving two elements, an integrated ternary relation, and a multi-relational exemplar where the given relation was depicted twice in the same image. Stimuli consisted of two images of the same relation in the same condition, which were paired in order to make a base analog (the image to be mapped from) and a target analog (the image to be mapped to). There were ten pairs in each condition, creating thirty pairs overall.

Each trial presented the base analog in the top half of the screen, while the target was presented directly underneath it. The base analog image had one item circled in red, while the target analog image had four objects with red numbers printed beside them (see Figure 6).

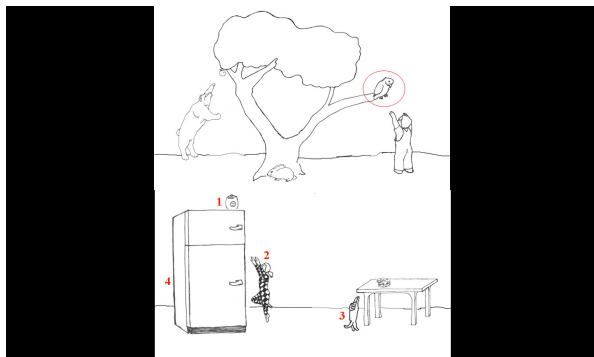


Figure 6: An example of a trial in Experiment 2.

Each condition was controlled for problem difficulty, and involved 9 cross-mapping problems. These problems required participants to reason over features more explicitly by including an item in both the base image and the target image.

Importantly, every question only had one possible answer and a number of distractor items. This fact was true even for the multi-relational trials, where only one item that was playing the correct role in the target image would be offered as a possible answer option, along with distractors. For example, imagine that “chases” was the relation in a given multi relational trial. The base image might depict *chases(boy,girl)* and *chases(dog,cat)*, with the boy circled, indicating that it was the object to be mapped. The target image might then depict *chases(bear, man)* and *chases(bird, worm)*, however, only the bear, the man, and the worm would be offered as answers, along with a distractor item such as an on-looking person.

Procedure: Participants were told that they were going to see two images at the same time, and that they were to match the circled item in the top image to one of the numbered items in the bottom image. Specifically, they were told to pick the item that they thought was “doing the same thing” as the circled item. Thus, participants needed to select the relational match between the base and target images.

Participants began by completing a single training example involving images that were not present in the rest of the experiment. Once the experiment began, participants were self-paced, moving forward by selecting one of the numbered items by pressing the keyboard key that matched the numbered item.

Results: Given the simplicity of the task involved, participants that completed less than an average of 15 out of 20 (75%) correctly across conditions were eliminated; five participants fell below this criterion and were eliminated. Unsurprisingly, as a result of this criterion, there was no significant difference between conditions for the number of correct responses ($F(2, 54) = 2.425, p = .118$), with performance on the integrated condition being almost equal ($M = 17.03, SD = 2.01$) to the multi-relational condition ($M = 17.21, SD = 1.90$), and only slightly higher in the binary relational condition ($M = 18, SD = 1.53$)

However, experiment two replicated the previous work on relational complexity with regard to reaction times. A repeated-measures Greenhouse-Geisser ANOVA was used due to a violation of sphericity, and it revealed a significant difference between conditions ($F(1.319, 23.742) = 22.970, p < .01$). Post hoc testing with a Sidak correction showed that participants were significantly faster ($p < .01$) on the binary relations (5.79 ± 1.92 sec) than they were on the multi-relational exemplars (8.47 ± 3.80 sec), however, the binary exemplars did not evoke significantly faster reaction times ($p = 0.48$) than the integrated exemplars (6.10 ± 2.11 sec). There was a significant difference ($p < .01$) between the integrated exemplars and the multi-relational exemplars (see Figure 7).

Discussion: This experiment found meaningful between-condition reaction time differences, which were consistent with the findings found in the literature discussed earlier in this paper. Thus, stimuli with greater complexity, integrated or not, were mapped more slowly than stimuli with lower levels of complexity. These results suggest that the findings in experiment one were not due to issues or idiosyncrasies with the stimuli, but instead represent a meaningful difference between how complexity interacts with relational recognition and mapping.

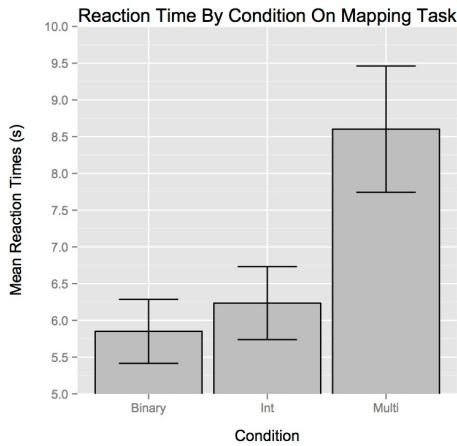


Figure 7: Results from Experiment 2. Error bars represent two standard errors.

Overall Discussion

With these results in mind it seems probable that relational recognition is more sensitive to features than to structure. As a result, it is likely that recognition shares some functional capacities with the access stage of analogy-making, though of course, this is an initial investigation and this relationship should be studied in more detail. Interestingly though, relational recognition seems particularly sensitive to the ratio of relational features to element features, as indicated by participants' performance on the integrated exemplars. Future research could also determine whether this is idiosyncratic to recognition, or whether access shares this sensitivity.

Furthermore, it now seems insufficient to say that relational complexity decreases relational performance, carte blanche. Contrary to the existing evidence on mapping tasks, there recognition (which is a necessary part of relational reasoning) seems to benefit from more complex exemplars. Future research could also delve into whether there are contexts or problem types for which the boost to recognition is more beneficial than the decrement to mapping.

Finally, this research suggests that current models of analogy erroneously take relational recognition for granted. These results suggest that the recognition process functions under unique constraints, and needs to be accounted for if the relational reasoning process is to be described as a whole.

Ultimately this research opens the door to more questions. We admit that we chose a somewhat arbitrary starting point based on trends in the existing literature and reason. Thus, our goal was not to provide all the answers about relational recognition, but to point out a deficit in the current literature and to start a discussion which may lead to those answers.

References

Cho, S., Holyoak, K.J., & Cannon, T.D. (2007). Analogical reasoning in working memory: recourses shared among relational integration, interference resolution, and maintenance. *Memory & Cognition*, 35(6), 1445-1455.

Doumas, L.A.A., Hummel, J.E., & Sandhofer, C.M. (2008). A theory of the discovery and prediction of relational concepts. *Psychological Review*, 115, 1-43.

Forbus, K.D., Gentner, D., Law, K. (1994). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19, 141-205.

Gentner, D. (1989). "The mechanisms of analogical learning" In S. Vossenbou & A. Ortony, *Similarity and analogical reasoning*. Cambridge University Press: New York, 199-233.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195-235). Cambridge: MIT Press.

Gentner, D., Rattermann, M.J., & Forbus, K.D. (1993). Similarity in transfer: Separating Retrievability from Inferential Soundness. *Cognitive Psychology*, 25, 524-575.

Gick, M.L., & Holyoak, K.J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.

Gick, M.L., & Holyoak, K.J. (1980). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.

Halford, G.S., Baker, R., McCredden, J.E., & Bain, J.D. (2005). How many variables can humans process? *Psychological Science*, 16(1), 70-76.

Halford, G.S., Wilson, W.H. & Phillips, W. (1998) Processing capacity defined by relational complexity: Implications for comparative, developmental and cognitive psychology. *Behavioral Brain Sciences*, 21(6), 803-831.

Halford, G.S., Wilson, W.H. & Phillips, S., (2010). Relational knowledge: The foundation of higher cognition. *Trends in Cognitive Sciences*, 14(11), 497-505.

Halford, G.S., Andrews, G., Wilson, W.H., & Phillips, S. (2012). Computational Models of Relational Processes in Cognitive Development, *Cognitive Development*, 27, 481-499.

Holyoak, K.J., Gentner, D., & Kokinov, B.N. (2001). Introduction: The place of analogy in cognition. In D. Gentner, K.J. Holyoak, and B.N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 1-19). Cambridge: MIT Press.

Hummel, J.E., & Holyoak, K.J. (2003). Relational reasoning in a neurally-plausible cognitive architecture: An overview of the LISA project. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society*, 10, 58-75.

Kokinov, B., French, R.B. (2002). Computational models of analogy-making. In L. Nadel (Ed), *Encyclopedia of Cognitive Science* (pp. 113-118). London: Nature Publishing Group.

Richland, L.E., Morrison, R.G., & Holyoak, K.J. (2006). Children's development of analogical reasoning: Insights from scene analogy problems. *Journal of Experimental Child Psychology*, 94, 249-271.

Thagard, P., & Holyoak, K.J. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, 46, 259-310.

Viskontas, I.V., Morrison, R.G., Holyoak, K.J., Hummel, J.E., & Knowlton, B.J. (2004). Relational integration, inhibition, and analogical reasoning in older adults. *Psychology and Aging*, 19(4), 581-591.

Waltz, J.A., Knowlton, B.J., Holyoak, K.J., Boone, K.B., Mishkin, F.S., de Menezes Santos, M., Thomas, C.R., & Miller, B.L. (1999). A system for relational reasoning in human prefrontal cortex. *Psychological Science*, 10(2), 119-125.