

Tracking Relations: The Effects of Visual Attention on Relational Recognition

Katherine A Livins (klivins@ucmerced.edu)

University of California at Merced, Department of Cognitive and Information Sciences, 5200 North Lake Road, Merced, CA, 95343

Leonidas A.A. Doulas (ldoulas@staffmail.ed.ac.uk)

University of Edinburgh, School of Philosophy, Psychology, and Language Sciences, Dugald Stewart Building, 3 Charles St., Edinburgh, EH8 9AD

Michael J. Spivey (spivey@ucmerced.edu)

University of California at Merced, Department of Cognitive and Information Sciences, 5200 North Lake Road, Merced, CA, 95343

Abstract

Relational recognition is the process by which relational representations get recognized (i.e., representations that specify an actor and a patient, and are role sensitive). This process is currently poorly understood, but is an important aspect of relational cognition (Livins & Doulas, 2014). This paper presents two experiments that investigate the degree to which visuospatial factors influence it. The first is an exploratory eye-tracking study that shows that first fixations are correlated with what object gets bound to the actor role, while the second uses priming to show that such fixations can alter which relation is recognized.

Key Words: relational recognition, relational reasoning, embodiment, eye-tracking, priming

The statements “the author enjoyed writing the paper” and “the chef enjoyed cooking the meal” have something in common—they both rely on *relational representations*. Strictly speaking, relations are functions that assign some truth-value to an ordered k-tuple (see Gentner, 1989), which boils down to saying that relations can take arguments, and that the order in which those arguments are specified is important. So, one can say, “I enjoy writing papers”, or “I enjoy cooking meals”, but saying that those things “enjoy me” is simply confusing.

Relational representations are powerful: their focus on roles means that objects that look nothing alike can be compared based on what they are *doing*. Thus, while the two aforementioned statements might involve completely different elements (authors and chefs do not *look* alike, and hopefully neither do papers and meals), they can be understood and compared based on a common relational structure, and the roles that such a structure affords (i.e., creators and their creations). Thus, relational cognition (i.e., cognition that works with these sorts of representations) has been implicated in a number of cognitive functions. For instance, analogy-making (Gentner, 1983; Doulas & Hummel, 2005), inductive reasoning (Hummel & Holyoak, 2003), and many forms of language-use (e.g., Gentner & Namy, 2006) have all been shown to rely heavily on relational processing. As a result, it is perhaps unsurprising that great effort has been spent attempting to account for it.

While these efforts differ in their details, they have converged on a general set of steps that occur during many

types of relational processing. These steps include access, mapping, transfer, and evaluation (e.g., see Holyoak, Gentner, & Kokinov, 2001). Broadly, access involves retrieving a source analog from long-term memory given a particular target (Hummel & Holyoak, 1997), mapping involves finding structural correspondences between that source and target (Hummel & Holyoak, 1997), transfer allows that mapping to be used to draw inferences by applying information about the base analog to the target, (Spellman & Holyoak, 1996), and evaluation involves adapting those inferences for the constraints and requirements of the problem at hand (Holyoak, Gentner & Kokinov, 2001). Learning is also sometimes included as final step in which new information, categories, and schemas can be added to memory based on the completed analogy (Holyoak, Gentner & Kokinov, 2001).

However, these accounts miss an important and fundamental step. Livins and Doulas (2014) pointed out that while they provide a detailed account of how one may reason about relations, they are silent as to how relations are recognized in the first place. *Relational recognition* has had limited study dedicated to it, but existing research suggests that it can be quite challenging. For instance, Gick and Holyoak (1980, 1983) pointed out that people often fail to notice relations unless they are explicitly directed to do so and that, if people fail to do so, then the entire reasoning process may never get off the ground (also see Doulas, et al., 2008). As a result, it seems important to account for relational recognition.

While little is known about how relational recognition works, research has begun by studying the types of factors that might influence it. For instance, Livins & Doulas (2014) found that, unlike mapping, relational recognition may be improved by increased amounts of relational complexity and an integrated relational structure. Likewise, while evidence is currently limited, there are empirical indications that visuospatial factors might also be important.

First, visual cues can boost performance in participants solving problem-solving tasks that involve the recognition of a relational concept. For example, Pedone, et al., (2001) showed that priming the concept of convergence with a schematic animation can improve the likelihood of then solving the Duncker Radiation Problem (a famously

difficult insight problem that relies on the recognition of a “convergence” relationship). Grant and Spivey (2003) further explored how people solve this same problem by presenting a diagram of the problem and recording the eye-movement patterns associated with the generation of successful solutions. They found that moving one’s eyes around the diagram in such a way as to spatially simulate the solution was correlated with success, and that evoking those same patterns with an augmented display improved performance (see also Thomas & Lleras, 2007).

The relational priming literature also hints at the importance of visuospatial factors. For example, Livins, Doumas, and Spivey (2014) showed that ocular movements might affect the learning of spatial relational categories like “next-to” and “above” in the presence of ambiguous stimuli. Specifically, they primed ocular movements congruent with the underlying representational directionality of one of those relations (horizontal for “beside” or vertical for “above/below”); they found that participants were more likely to recognize and learn whatever relational category had been primed. Interestingly though, the manipulation did not affect performance after learning, which suggests that it might have affected the recognition of one relation over the other. Likewise Livins, Doumas, and Spivey (submitted) showed that crossmapping analogy problems are more likely to be solved with a relational mapping instead of a featural mapping if problem-solving is immediately preceded by ocular movements congruent with the relational content in the problem. This result suggests that visuospatial priming may not only alter which relation is recognized or learned, but whether one is recognized at all.

While this literature is suggestive, it is not conclusive: the existing research has hinged on participants’ ability to recognize a relation, but has not tested relational recognition in specific. Thus, a rigorous test of the role of visuospatial factors in relational recognition is necessary, along with a theoretical account of why it might be important.

To date, there has been one existing effort towards this project. Franconeri et al. (2012) looked at the role of vision in the processing of spatial relations (such as “to the left of” or “to the right of”). They argued that the visual system will register such relations, not holistically, but by processing each relationally-relevant object sequentially. This “shift” account predicts that attention and ocular shifts (saccades) between objects help to encode these relations, and that at least one shift is necessary for recognition. For example, if one is looking at a scene with a series of shapes, one might recognize that “one shape is to the left of another” by looking at the left one, then making a saccade to one on the right. Franconeri et al. (2012) used eye-tracking and a relational judgment task to confirm that such movements occurred prior to making relational judgments.

Interestingly, at least one model of relational reasoning predicts that this type of sequential attentional-shift-based processing is necessary for thinking about all relations (not just spatial ones). The DORA model (Doumas et al., 2008) encodes relations across layers of nodes, and uses time to

encode roles and fillers. So, the lowest level is a set of distributed features (much like those found in traditional connectionist networks). One layer up, localist nodes combine sets of those features to code for objects and relational roles, which are then temporarily bound to create more complex relational structures (e.g., see Figure 1). As a result, relations are not actually represented as wholes, but instead as combinations of their roles and fillers. So, for example, *chases* is not represented by a *chases* node, but by the combination of *chaser* and *chased*, which can be temporarily bound to things like *cat* and *dog* to create something like *chaser(dog) + chased(cat)*. Eventually, these role-bindings are combined to create full relational statements like *chased(dog, cat)*. Importantly though DORA binds through temporal asynchrony—in other words, by tracking when units fire and the sequence by which they do so. So, *chased(dog, cat)* would be represented by firing *chaser*, then *dog*, then *chased*, then *cat*, and *chaser* and *dog* would be bound by firing them in immediate temporal proximity. Thus, the model requires the subsequent firing of each relational role—the actor, and then the patient. As a result, the model predicts that one must encode both roles (and the objects that fill them) independently in order to process a relation, and that the order in which things fire is representationally important.

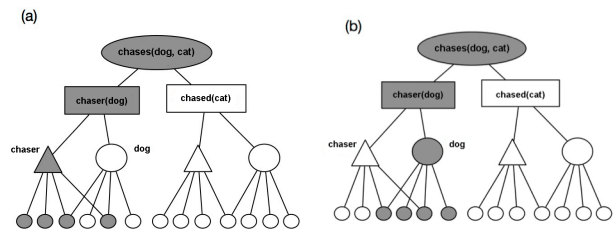


Figure 1: An example of how DORA represents relations. *A* shows the model firing a role, and *B* shows it firing a filler directly after in order to represent *chasing(dog, cat)*.

Thus, it maybe be the case that attention should not just be necessary for recognizing a relation, but also for specifying *which* relation is recognized: if one attends to one object playing one role before another object playing another role, then the former might be designated as the actor, while the later the patient. Visual processing may guide which is attended to first, and so which is fired first, and so which is designated as the actor (thereby affecting relational recognition). This paper will test this possibility.

Experiment 1

The general objective of this experiment is to determine whether eye-movement patterns can predict relational recognition. It will use a paradigm similar to that found in Gleitman et al. (2007), which used eye-tracking to show that gaze can shape the structure of a sentence used to describe a given scene. For example, it showed participants an image of a dog chasing a man and asked whether it could be

described by statements like, “The man chases the dog” or “The dog flees from the man”. They found that the selected structure was influenced by the item that was looked at first. However, the study looked only at sentence *structure*, along with the actor/patient designations in a given verb. They did not, however, look at whether this designation could change *what verb or relation was represented/identified entirely*. The current work addresses this issue. Thus, it will test whether the first object that one fixates on in a scene can predict not only which object is treated as the actor or patient (e.g., Griffin & Bock, 2000), but also what *relation* is explicitly recognized and identified. First fixations were used because they are a measure of visual attention.

This experiment hinges on the fact that the scenes used, and the real world in general, depict numerous relations at any given time. For example, a picture of a mother feeding a child might depict a feeding relationship, as well as an eating relationship between the child and the food, a sitting relationship between the mother and a chair, and any number of spatial relationships (next to, beside, etc.; see Figure 2). Relational recognition involves, at some level, prioritizing one relation over the others, and so understanding relational recognition will mean asking why that prioritization might occur and what factors might cause it. Thus, the research question asked here is whether initial visual fixation within a scene is one such factor.



Figure 2: An example of a scene that might be described as “feeding”, but which also depicts an “eating” relationship, as well as numerous spatial ones.

Participants: Participants were 58 University of California Merced undergraduates. All were over 18 years of age, had normal vision to corrected-to-normal vision with contacts (no glasses were allowed). The data from two more were collected but excluded due to low eye-tracking locks.

Materials: Stimuli consisted of 21 pictorial scenes adapted from Richland, Morrison, and Holyoak (2006). Each stimulus had six objects dispersed around a white and black, drawn image. All stimuli were 720 by 450 pixels in size and were presented on a black background. The images were centered on a computer screen such that there was a black outline around them, totaling 1440 by 900 pixels in size. The images included both living and non-living elements.

Every image depicted two objects engaged in a primary relational activity (e.g., while one person hugging a dog might be described as being “next-to” it, “hugging” might be a more prominent relation in the scene; see Figure 3). Each stimulus was coded by two experimenters and when the results were compared, 100% agreement was found.

Overall, there were two classifications of relations. First, key relations were chosen because they could be represented as 2-place relations and were amenable to one-word

descriptions that differed depending on which object was bound to which role (i.e., which object was designated as the actor and which was designated as the patient). For example, *chasing(x,y)* might be described as *escaping(y,x)*. These relations were depicted such that two of the image’s primary objects (the ones engaged in the primary relation) began equidistant from the center of the screen on the x-axis (see Figure 3). The full list of these relations can be found in Table 1. Second, filler items were chosen because they were also expressible as two-place relations, but had a more prominent single relation (see Table 2). These relations were not depicted with the prominent relational items in the center of the screen since their inclusion was intended to ensure that participants did not develop trained biases to one location or side of the screen.



Figure 3: An example of a scene that possess multiple relations, but in which one (a “kissing” relationship) might be more prominent than the others.




Figure 4: An example of a key stimulus in which the two relationally-engaged objects begin an equal distance away from the image-center.

Possible Relation Description 1	Possible Relational Description 2	Objects Used In Stimuli
Chasing	Escaping	boy, cat
Talking	Listening	woman1,woman2
Lifting	Hanging	woman, monkey
Hunting	Escaping	man, elephant
Kicking	Covering	boy, dog
Showing	Watching	boy, woman
Dropping	Falling	woman, baby
Pulling	Riding	boy, dog
Eating	Feeding	mother, child
Pushing	Riding	girl, boy

Table 1: Key relations used in Experiments 1 and 2. Each relation afforded multiple relational descriptions.

Stimuli were presented in a random order using the Pygame module. Pygame was interfaced with an EyeLink II (i.e., a binocular eye-track made by SR Research) to collect ocular fixations and saccades. Each stimulus had a small text-box below it so that participants could enter an answer by typing it in and then pressing “Enter”. Possible spatial biases were controlled by flipping the images on their horizontal axes across participants. Thus, half of the participants saw one item on the right hand side of the screen, while the other half saw that same item on the left.

Primary Relation	Objects Used In Stimuli
Brushing	girl, hair
Cooking	man, food
Fighting	boy1, boy2
Hoisting	girl, monkey
Kissing	girl, dog
Opening	girl, gift
Pouring	boy, water
Reaching	man, baby
Scolding	woman, girl
Towing	tow-truck, car

Table 2: A list of filler items used in Experiment 1

Design: The experiment began with eye-tracker calibration. For this process, each participant was fitted with the head-mounted eye-tracker so that it was securely fastened. They sat approximately 36 inches from a 24-inch flat panel LCD monitor. Cameras were adjusted and focused, and the thresholds for detecting pupils were automatically calibrated. This allowed the experimenter to ensure that the track was not lost at any location on the screen. A nine-point calibration was performed before validation, which ensured that there were no tracking errors. If validation showed minimal error, then the experiment began.

Participants were then told (both verbally and in text) that they needed to type the relational verb that they thought was most prominently depicted in each picture. A single training trial was then given. It began with a fixation cross that was shown for 1000ms, was white, and was centered on screen. A relational image was then shown, which depicted a “playing-with” verb, but was otherwise the same as the rest of the stimuli. Participants were told to type an answer, and then shown their own answer with the possible candidate answer of “playing-with”. Both were shown so to ensure that they understood what a relational verb was. Instructions were then reiterated.

Participants then began the experiment. They worked, self-paced, through all problems (no further instructions were given). Drift-corrects were taken every 5 trials to ensure that the eye-tracking lock was maintained.

Results: Two measures were collected. First, we analyzed participants’ responses. These were in the form of words, and coded based on which object was bound to the actor role (for example, “chasing” would designate the boy in Figure 3 as the actor). For the sake of calculations, one relation was chosen as the default for each image (in every case this default was the relation listed in the first column of Table 1), and responses were coded as 1 for “actor-based” or 0 for “patient-based”. So, for example, “chasing” was considered the default for one image, and so a “chasing” response was coded as “actor-based”, while “escaping” was coded as “patient-based”.

Given that this experiment was exploratory, and that we wanted to determine whether there is a correlation between looking at an item and recognizing a relation where that item is the actor, we had a number of exclusion criteria. First, any non-verb responses were eliminated (e.g.,

“friendship”) since such answers showed a lack of understanding with regard to the task. Likewise, any responses that were either non-relational (e.g., “running”) or unclear with regard to which object was the actor (e.g., “playing”) were eliminated. It is interesting to note that, despite the open nature of the responses, there was a high degree of commonality across answers. For example, for one stimulus “feeding” was provided 44 times, and “eating” was provided 6 times—no other answers were given. Likewise, another stimulus was described as “kicking” in 53 out of 54 valid responses. This result suggests that each image had a “dominant” relation to participants.

Second, ocular attention was tracked. We were specifically interested in the first item of fixation, which was operationalized as the first object within an image’s primary relation that was fixated upon. Analysis began by specifying square “areas of interest” around each object, and then checked whether a fixation was within that area. Like in the case of participant responses, fixations were coded as being “actor-” or “patient-oriented”.

Overall, 352 (approximately 72.43%) of responses matched the item of first fixation (while 134, or approximately 27.57%, did not; see Figure 5). However, because this was a repeated measures design, we used mixed effects logistic regression (see Jager, 2008) to further interpret these results. For this analysis, assuming a dominant relation (we used the first column of Table 1 for this purpose) the actor/patient orientation of the participants’ response was treated as the criterion variable, while the first fixation was treated as a predictor. Given that this experiment used a repeated-measures design, Participant ID (of which there were 56) was also included in the model as a nested factor, along with the image, which was treated as a random factor. The model is described in Table 3, and a likelihood ratio test was used to compare it to a null model; it was found that first fixation made a significant difference ($\chi(1)=3.926, p<.05$).

Predictor	Coef (β)	SE(β)	z	p	Odds Ratio
Intercept	2.1000	0.8983	2.338	0.0194	8.165783
First Fixation	0.7015	0.3310	2.120	0.0340	2.016787

Table 3: The model results from Experiment 1

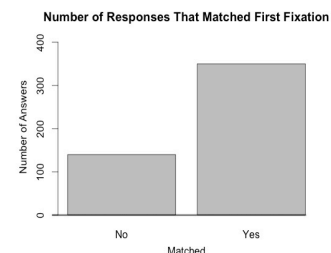


Figure 5: A graphical representation of the overall number of responses that matched the first fixation made within each stimulus.

Discussion: The results of this study suggest that there exists a relationship between the item that one fixates on first and the item that one designates as a relational actor. As a result, they suggest that fixation is somehow related to what relation that is recognized. However, this study was correlational in nature, and so Experiment 2 will attempt to direct visual attention to different objects in order to determine whether this relationship is causal.

Experiment 2

The objective of this experiment was to determine whether the trajectory of relational recognition may be manipulated by visual attention. Specifically, because of the results of Experiment 1, it will test whether priming the first item of fixation can change what relation is identified. The experiment will be almost identical to Experiment 1, however, it will direct visual attention towards a specific object in each scene at the beginning of every trial.

Participants: Participants were 132 University of California Merced undergraduates that were otherwise analogous to those used in Experiment 1. Four participants were eliminated entirely due to poor eye-tracking locks.

Materials: The materials were the same as those listed in Experiment 1 with one addition. Priming was achieved by exploiting the eye-tracker’s normal calibration process. Specifically, calibration involved a series of small black dots with a white center point that appeared in various places around the screen. It required participants to fixate on the center of those dots and to press “spacebar”. Thus, key trials involved two extra “calibration dots”: one just before an image was shown, and then one 100 to 500 ms after the image appeared (the exact amount of time was randomly generated). A random number of filler trials also had extra “calibration” dots, but the locations of the dots were randomly generated and scattered across the screen.

Design: This experiment proceeded in almost the same way as Experiment 1. However, during initial calibration the experimenter emphasized that she was having trouble getting a lock on the participant and so extra calibration throughout the study might be required.

Two controls were used: First, like in Experiment 1, images were flipped on their horizontal axes for half of the participants. Second, each relationally relevant item was primed for half of the participants. So, for example, if a trial depicted *chases*(boy, cat) (or *escapes*(cat,boy)), then half of the participants were primed to initially fixate on the boy, while the other half were primed to fixate on the cat.

Results: Once again, participant responses and first fixations were tracked. However, the coding system for the responses changed slightly due to the research question. To the point, our goal was to determine whether making someone fixate on a specific object would change the relation given. Thus, we allowed for neutral responses in this experiment (and not just actor or patient based ones, like in the previous experiment). For example, “conversing” was allowed for the “talking” stimulus, despite the fact that conversing is a bidirectional relation. This approach seemed

especially warranted given that the data from the first experiment indicated that most stimuli had a dominant relation that was recognized by most participants (i.e., one object that was typically bound to the actor role), and so looking for deviations seemed worthwhile.

First fixations were tracked and used to eliminate participants. Again, given that our research question was whether changing participants’ first fixations would change the course of the recognition process, we used fixations to ensure that participants actually fixated on the prime. Trials in which a participant initially fixated on a different object were eliminated (this included .03% of all trials).

We then used a mixed effects multinomial logistic regression model to interpret our results. Once again, participants’ answers were treated as the criterion variable, while the prime was treated as the predictor, participant ID was treated as a nested variable, and image as a random variable. The model is described in Table 4, and a likelihood ratio test comparing the model to null showed that priming was a significant factor ($\chi(1)=35.343, p<.01$). Specifically, it showed (again, by odds ratio) that one is 4.25 times more likely to recognize a relation that uses the primed item as an actor. The overall differences in responses by condition can be seen in Figure 5.

Predictor	Coef (β)	SE(β)	z	p	Odds Ratio
Intercept	2.0807	0.8663	2.430	0.0151	8.010303
Primed	1.4462	0.2477	5.839	<0.0001	4.246898

Table 4: The model results from Experiment 2

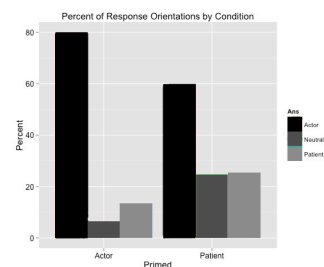


Figure 6: A graphical representation of the overall number of responses that matched the first fixation made within each stimulus.

Discussion: The results of this study suggest that it is, indeed, possible to shape relational recognition by manipulating which item is fixated on first. Thus, the relationship between first fixation and relational recognition is not just correlational, but causal.

Overall Discussion

When considered together, these studies suggest that relational recognition is not only correlated with one’s visual attention, but that it can also be changed by that attention. Specifically, Experiment 1 showed that the first item that one fixates on when scanning a scene may predict which relation is recognized, while the second suggests that

manipulating that first fixation by directing it at one item or another can make it more likely that one will recognize a relation that designates that object as an actor. As a result they have at least two theoretical implications.

First, very generally, they help to link visual processing to relational processing—at minimum, they suggest that where one looks affects what one attends to, which affects what relation one recognizes, and therefore what relation one reasons about. This is an important step for embodied efforts, which have argued that the body is an important part of cognitive functioning (e.g., see Spivey 2008), but does not detract from the possible importance of symbolic content (which is emphasized in the relational reasoning literature, e.g., see Gentner, 1983; Dumas & Hummel, 2005). That said, an interesting question to ask in the future is whether and how this process feeds back into perceptual processing to create an overall trajectory of reasoning in dynamic, information-rich environments.

Secondly, these results also have implications for debates about mental representation—especially with regard to how relations are represented. DORA and its predecessor LISA (see Hummel & Holyoak, 2003) are unique in the way that they use role-filler bindings and time to create more complex relational structures. An important prediction of those representational structures is that relations are not processed holistically, but in terms of roles, the items that fill them, and the temporal sequence in which they fire. This account has been supported by Franconeri et al. (2012), and extended beyond the realm of spatial relations in this work.

Finally, these studies help contribute to our understanding of relational recognition as a process. Relational recognition is generally not well understood, and research on it is just beginning. Livins and Dumas (2014) suggested that relational complexity is one important factor, and here visual attention can be specified as another. That said, we did not find the correlation between fixation and recognition to be perfect, which suggests that other factors are still yet to be found. One possibility is how linguistically common a relation is over alternative descriptions (e.g., the word “kicking” is used more commonly than the word “cowering”). Future research will look at such factors in order to provide a more cohesive account of relational recognition in general. Ultimately, the goal will be to describe recognition in such a way as to incorporate it into the overall relational-reasoning process.

References

- Dumas, L. A. A. & Hummel, J. E. (2005). Modeling human mental representations: *The Cambridge handbook of thinking and reasoning* (pp. 73-91). Cambridge, UK: Cambridge University Press.
- Dumas, L.A.A., Hummel, J.E., & Sandhofer, C.M. (2008). A theory of the discovery and prediction of relational concepts. *Psychological Review*, 115(1), 1-43.
- Franconeri, S.L., Scimeca, J.M., Roth, J.C., Helseth, S.A., & Kahn, L.E. (2012). Flexible visual processing of spatial relationships. *Cognition*, 122, 210-227.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning*. New York: Cambridge University Press.
- Gentner, D., & Namy, L. L. (2006). Analogical processes in language learning. *Current Directions in Psychological Science*, 15(6), 297-301.
- 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.
- Gleitman, L.R., January, D., Nappa, R., Trueswell, J.C. (2007). On the give and take between event apprehension and utterance formulation. *Journal of Memory and Language*, 57(4), 544-569.
- Grant, E.R., & Spivey, M. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, 14(5), 462-466.
- Griffin, Z. & Bock, K. (2000). What the eyes say about speaking. *Psychological Science*, 11, 274-279.
- Hummel, J.E., & Holyoak, K.J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*, 110, 220-264.
- Hummel, J.E., & Holyoak, K.J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427-466.
- Jaeger, T.F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59(4), 434-446.
- Livins, K.A., & Dumas, L.A.A. (in press). Recognizing relations: What can be learned from considering complexity. *Thinking and Reasoning*.
- Livins, K.A., Dumas, L.A.A., & Spivey, M.J. (submitted). Shaping relations: Exploiting relational features for visuospatial priming
- Livins, K.A., Dumas, L.A.A., & Spivey, M.J. (2014). The effects of visuospatial priming on category learning. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*. Quebec City: Cognitive Science Society.
- Pedone, R., Hummel, J.E., & Holyoak, K.J. (2001). The use of diagrams in analogical problem solving. *Memory & Cognition*, 29, 214-221.
- 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.
- Spellman, B.A., Holyoak, K.J. (1996). Pragmatics in analogical mapping. *Cognitive Psychology*, 31, 307-346.
- Thomas, L.E. & Lleras, A. (2007). Moving eyes and moving thought: On the spatial compatibility between eye movements and cognition. *Psychonomic Bulletin and Review*, 14(4), 663-668.