

Shaping Relational Category-Learning With Visuospatial Priming

Katherine A. Livins & Michael J. Spivey (klivins; spivey@ucmerced.edu)

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

Leonidas A.A. Doumas (ldoumas@staffmail.ed.ac.uk)

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

Abstract

While relational reasoning has been described as a process at the heart of human cognition, the degree to which relational representations can be primed remains an open debate. This paper will present a category-learning experiment that shows that the learning of spatial relations (above and below) can be primed using a subtle visuospatial stimulus that may capture exogenous attention to produce saccades.

Keywords: relational reasoning, category learning, visuospatial priming, embodiment

Relational representations specify how two things are related to each other based on the roles that they play, rather than on the features that they possess. For example, if you are told that “the monkey hangs from the chair”, then you know that there is a hanging relationship such that the monkey is the actor (i.e., the hanger) and the chair is the patient (i.e., the hung-from thing). You might even infer new information about the elements engaged in the relation by virtue of knowing something about the relation’s roles. So, if you know that hung-from things are typically strong, then you may infer that this chair is strong too, despite the fact that strength is not an inherent feature of being a chair (e.g., folding chairs are notorious for collapsing). Thus, relations provide the opportunity for powerful inferences and generalizations.

Relations are also pervasive. They not only provide the foundation for cognitive processes ranging from analogy-making (Gentner, 1983; Doumas & Hummel, 2005), to inductive generalization (Hummel & Holyoak, 2003), but they have even been related to a plethora of processes that, at first glance, may seem non-relational such as linguistic processing (Gentner & Namy, 2006), and even social cognition (Spellman & Holyoak, 1992). As a result, there has been a field-wide interest in how relations function. To date, the majority of this work has focused on providing accounts of the mechanisms that allow for structural alignment between relations based on shared roles, the transfer of information from one element to another based on those alignments, and the types of mental representations that may be necessary for this type of processing (e.g., Falkenhainer, & Gentner, 1986; Hummel & Holyoak, 2003; Doumas, Hummel, & Sandhofer, 2008).

However, the extensive literature on these topics does not mean that research into relational reasoning is complete. For instance, an open debate is whether, and to what degree, relations can be primed.

Two competing answers to this question have emerged. First, it has been argued that relational reasoning is nothing but priming, and so relational priming must be extremely common. Leech, Mereschal, and Cooper (2008) have championed this position by arguing that relations are learned as patterns, and that reasoning about relations is nothing more than exploiting those associations. For example, they have suggested that A:B::C:D problems may be solved with nothing more than associations between the given concepts. Thus, when one is told that *puppy is to dog as kitten is to something*, the concept “is the offspring of” is primed, which then biases one to respond *cat*. As a result, priming relations is simply a matter of priming a relevant context.

However, Leech et al.’s position has been highly criticized. Most problematically, it seems unable to account for the types of behavior that are characteristic of adult performance. This problem is highlighted in the simple recurrent network that they built to instantiate their account. Doumas and Richland (2008) point out that the model is unable to integrate multiple relations during analogy-making because it functions based on associations alone; however, human adults perform this sort of integration regularly. Likewise, French (2008), and Holyoak and Hummel (2008) point out that far reaching analogies (that share few semantic characteristics) would also be beyond the model’s capabilities, despite the fact that human adults regularly make these sorts of analogies. As a result, the model can only be successful on specific (carefully designed) relational problems and therefore cannot be thought of as a more generalized explanation of reasoning across contexts and content types.

In light of these criticisms, it is unsurprising that a second position has been suggested. In short, it claims that relational reasoning is primable, but that it is rare. Spellman et al. (2001) demonstrated just how rare with an experiment that attempted to prime relational concepts on a lexical decision task. It was found that,

while priming did occur, that it was profoundly limited by the experiment's instructions to attend to relations within pairs of words and across pairs of words.

Ultimately, the degree to which relations can be primed remains questionable: while it may be insufficient to say that all relational reasoning is equivalent to priming, it is almost certainly inaccurate to say that all relational concepts cannot be primed at all. In fact, some accounts of relational reasoning even make it reasonable to expect that relational concepts can be primed more easily than the Spellman et al. study suggests. For instance, the DORA model proposed by Doumas, Hummel, and Sandhoffer (2008), posits that relational representations are learned from exemplars experienced in one's environment; those representations are made more abstract and structured through a refinement process that occurs from exposure to many exemplars. However, some features remain integral to a relation's representation, and so it seems likely that if those features could be accessed, then they could be exploited for the purposes of priming.

How can we access these features though? While there is no obvious answer to this question, embodiment researchers have invented a number of priming paradigms that may be useful (especially given that DORA suggests that relational features are learned at least in part, by experiences in the world).

Eye movements have played a central role in many of these paradigms, likely because they have a low threshold of activation, and are generally resistant to a participant's strategic plans (Spivey, Richardson, & Dale, 2009). For example, Grant and Spivey (2003) looked at the effects of eye movements on the processing of insight problems. Specifically, they had participants attempt to solve the Dunker Radiation problem (Dunker, 1945) while wearing an eye-tracker; as participants worked on the problem they were allowed to look around a potentially useful diagram found that the individuals who successfully solved the problem spent a greater amount of time looking at particular regions of the diagram. Grant and Spivey expected that, if visual attention is at all related to problem solving, then drawing peoples' attention to those regions should increase successful response rates. As a result, they completed a subsequent experiment that did exactly that, and it was found that their expectations were correct. Therefore, it was established that visual looking patterns may be linked to high-level problem-solving.

Thomas and Lleras (2007) used Grant and Spivey's results to further explore the importance of eye movements on reasoning. They asked whether eye movements were important to the reasoning process, or whether visual attention was sufficient for increased success rates, regardless of the ocular movement pattern. They had participants complete two tasks: a visual tracking task, and then a problem-solving task

(the Dunker problem). The visual tracking task involved eye movements around the Grant and Spivey diagram—however only one group of participants moved their eyes in a way consistent with the problem's correct response, while the others simply moved their eyes around the areas of the diagram that Grant and Spivey found to be important. Consistent with the claim that eye movements can prime a correct response, it was found that the "embodied solution" group showed the greatest success rate.

The work by Grant and Spivey (2003) and Thomas and Lleras (2007) is interesting because it suggests a link between ocular movements and complex reasoning—a type of complex reasoning very similar to that involved in relational cognition. Thus, it seems reasonable to suppose that eye movement may then be capable of priming relational cognition as well. However, it is vague to simply say that eye movements may be useful, and consideration of *how* they might be useful is necessary.

The answer may lay in a study conducted by Richardson et al. (2001), which explored whether there are consistencies in the visuospatial imagery associated with action verbs across individuals. After selecting a variety of verbs, they presented participants with sentences involving each verb, along with four pictures of a circle and a square in different spatial alignments; participants were asked which alignment best represented the verb. There was significant consistency across participants, suggesting that individuals may share spatial schemas about those verbs.

Richardson et al. continued to explore this possibility by having another set of participants freely draw representations of the same set of verbs using circles, squares, and arrows of varying sizes. While there was a greater amount of variance across the drawings of abstract verbs (e.g., "*tempted*") than concrete ones (e.g., "*lifted*"), it was generally found that there was still a significant amount of consistency across participants with regard to the angle at which the shapes were placed. For example, "*argued with*" was consistently drawn with a horizontal alignment, while "*respected*" was consistently drawn with a vertical alignment. As a result, Richardson et al. argued that it is likely that spatial traces are part of the representations of the given verbs—a possibility that was supported again in a second paper (Richardson et al., 2003) that used a memory-recall task and found a similar trend.

While verbs are not synonymous with relations, Richardson et al. used verbs that are inherently relational: each specified an actor and a patient (e.g., *pointed at*, *pushed*, *lifted*, and *argued with*) and so their findings suggest that at least some relations have visuospatial features associated with them.

Ultimately then, there exists both computational and empirical evidence to suggest that relational

concepts have ties to bodily experiences. Furthermore, embodiment research suggests that priming paradigms that involve eye movements may be particularly useful for some types higher-cognitive functioning, and that at least some relations may be specifically sensitive to visuospatial manipulations. Thus, the experiment presented below attempts to prime relational cognition using a visuospatial prime.

Experiment

The objective of the experiment was to determine whether it is possible to prime relational category-learning: we employed a pictorial category-learning task to determine whether simple, spatial relations (above and below) can be primed using a subtle visuospatial prime that captures exogenous attention to produce saccades.

It is important to note this experiment was designed with two assumptions. First, category-learning can be relational. This assumption is based on Gentner and Kurtz (2005) who pointed out that, while not all categories are relational, some are. Specifically, relational categories define membership based on some common relational structure instead of member features. For example, *occluders* make up a relational category since they are not defined by their features, but rather by how an object stands in relation to other objects. Category-learning tasks that involve these sorts of categories require the same sorts of mechanisms that underlie analogy-making, mapping, schema-induction, etc. Thus, if it is possible to prime category-learning on a relational category, then it will be possible to claim that relational concepts can be primed.

Secondly, we assume that when someone is presented with a relationally ambiguous exemplar that simultaneously represents a value on two different relations, but where learning one is sufficient for task completion (like deciding whether the exemplar is part of a category), that only one will be learned. The reason for this expectation is that relational reasoning is an explicit process that taxes working memory—the more relations that one entertains, the more working memory is taxed (Doumas et al. 2008). However, working memory is limited, and so people will typically stop working when they have a sufficient answer.

These assumptions are important because this experiment required participants to learn a relational rule in order to decipher category membership. However, every exemplar had two relations present simultaneously, and priming was designed to affect which relation was learned.

Participants: Participants were 105 undergraduate students from the University of California, Merced. They were recruited through a participant pool and received course credit for participation. All participants

had normal to corrected-to-normal vision. Thirteen of these participants were not included in the final statistical analyses because of a lack of rule learning, however they were used to calculate the sample's overall ability to complete the task.

Design: Participants were assigned to one of three groups: a control group that received no prime, a prime-with-vertical-movement group, or a prime-with-horizontal-movement group. All participants began in the same way: they were seated at a computer with a 2560 by 1440 pixel monitor, which showed stimuli presented in an experiment space of 1440 by 900 pixels.

They were told that they were going to see pairs of shapes, and that each pair was positioned according to a rule—they were also told that they were not going to be told what the rule was. Given that this was a feedback-learning paradigm, they were instructed to determine the rule by trial-and-error using the feedback provided each time they entered an answer. These instructions were provided both verbally and in text.

If the participants were in a priming group, they were also told that they may occasionally see “blinking dots”, and that the dots were due to a slow computer attempting to generate the stimuli. In fact, the aforementioned “dots” were the prime, and they could be presented in either a vertical or horizontal fashion. In both cases the “dots” were 130 pixel-large white circles, with another 2 pixels of black outline around them (totaling 15 pixels in size). If the participant was in the horizontal prime condition, then they appeared horizontally aligned and half way down the screen on the y-axis; if participants were in the vertical prime condition, then they appeared vertically aligned half way across the screen on the x-axis. In both cases, the circles were spaced 540 pixels away from each other, spread out around the center in the specified direction (horizontal or vertical). One dot would blink on for 500 ms, then blink off; there would be a 100 ms delay, and then the other would blink on for 500ms. This cycle iterated five times for the initial prime.

Note, participants were not told to watch the dots. However, participants were left alone with no distractions. Thus, while we cannot confirm that they visually tracked the dots, it was expected that the prime captured their exogenous attention.

Participants then began the “training phase” of the task. During this phase, participants would see a fixation cross, which would appear for 1500 ms, then an exemplar. The exemplar categories were created using simple shapes (circles and squares) and their relative placement on the x and y-axes. More specifically, every exemplar showed two shapes, where one shape occluded the other; the specific shapes were selected at random at the beginning of each trial, creating non-predictive shape selections such that each trial could contain two circles, two squares, or one of

each. Every occluder took a value on two different relations: it could be to the left or right of the occluded shape, and it could be above or below it. Thus, every exemplar could be categorized as an “A” if the occluder was above the occluded shape, a “B” if the occluder was below the occluded shape, a “C” if it was to the left of the occluded shape, and a “D” if it was to the right of the occluded shape (see Figure 1).

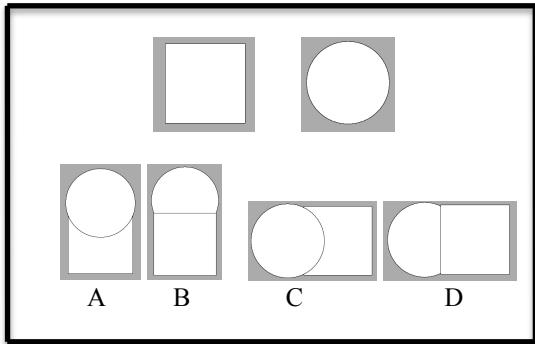


Figure 1: An example of how two shapes could combine to create exemplars that would be classified as an “A”, a “B”, a “C”, and a “D”.

Combining values on the two relations allowed for the creation of ambiguous stimuli such that every exemplar would simultaneously represent more than one relation. In other words, category membership was specified by the values taken on multiple relations. As a result, A/C pairings could be created to depict an occluder that was above and to the left of the occluded shape, B/D pairings could depict an occluder that was to the bottom and to the right of the occluded shape, A/D pairings that could depict an occluder that was to the top and to right of the occluded shape, and B/C pairings that could depict an occluder to the bottom and to the left of the occluded shape (see Figure 2).

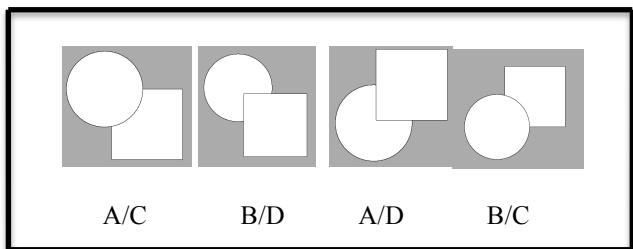


Figure 2: Examples of exemplars that combine a value on the left/right relation with a value on the above/below relation.

The training phase was programmed to randomly select a pair of training rules, which conflated a relative location on the horizontal axis with a relative location on a vertical axis. Thus the training phase would include A/C and B/D pairs, or A/D and B/C pairs. One rule pair would be randomly associated with the “A”

key, and the other to the “L” key. Participants would press a key for every exemplar, and “Correct” or “Incorrect” would appear every time.

Since the values across the two relations were conflated, participants could learn a horizontal rule, a vertical rule, or both rules. For example, if a participant’s training rules were A/C and B/D, where A/C was assigned to the “A” key, then she could learn that “A” needed to be pressed whenever the occluder was to the left of the occluded shape, or she could learn that “A” needed to be pressed whenever the occluder was above the occluded shape, or she could learn that she needed to press “A” whenever the occluder was above and to the left of the occluded shape. As a result, priming was always consistent with one rule, and inconsistent with another rule.

Training began by presenting 8 exemplars of the same training rule, and then switched to random selection between the two available rules for every exemplar after that. So, for example, the training condition could proceed by presenting 8 exemplars of A/C, A/C, A/C, A/C, A/C, A/C, A/C, A/C...[random]. The initially presented rule was counterbalanced across participants in each condition. This training regimen was selected based on Clapper (2009), who claimed that this sort of presentation would increase ease of learning in dichotomous category learning tasks.

The experiment began keeping track the number of correct responses that a participant gave after the initial 8 training trials ended (i.e., when random presentation began). Participants continued to see pairs of shapes, and get feedback until they learned a rule well enough to correctly classify 10 trials in a row. If a participant answered a trial incorrectly, the counter reset to zero and if a participant was in a priming condition, then the prime would reappear after every 5 trials until the criterion was met; however, the priming would only appear for 3 iterations instead of the 5 that were presented at the beginning of the experiment.

Once the participant reached criterion, they were told that they would continue to see pairs of shapes, but that all feedback as to whether they were correct or incorrect would stop. They were also told to continue to use the same rule that they had learned for the remainder of the experiment.

The test phase of the experiment then began. If a participant was in a priming condition, priming was stopped (since the goal of the priming was to affect rule learning, and the rule was learned by this point).

Participants were presented with a random order of seven exemplars of each possible variable combination, such that they would now see A/C, A/D, B/C, and B/D shape alignments. The goal of the test phase was to allow the experimenters to determine the rule that the participant had learned and was then applying, which could be achieved by looking at their responses to novel

alignment combinations: Since training had conflated a value on the beside relation with a value on the above relation in two different ways (each marked by a specific key press), the novel stimuli would contain half of each trained pair. Thus, a response to a novel stimulus would indicate which pair the participant thought the novel pair was like, and therefore whether they learned the “above” or “beside” rule.

For example, suppose that a participant had been trained on A/C and B/D, where A/C had been associated with an “A” key press, and B/D had been associated with an “L” key press. A/D and B/C pairs could be used to determine which rule the participant had learned: If presented with an A/D pairing, then an “A” key press would indicate that the participant was classifying the stimulus like an A/C pair. If A/C and A/D pairs are classified in the same way, then the participant must be attending to the above/below relation (since A is the common relational value between them). Conversely, an “L” key press would indicate that the participant was classifying by the “beside” rule (See Figure 3).

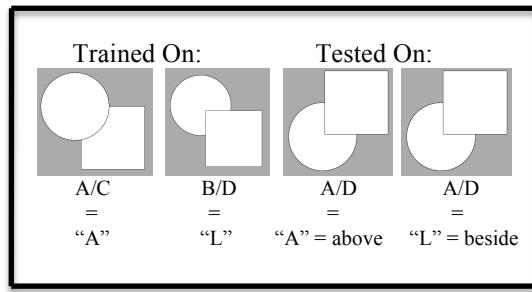


Figure 3: An example of a possible training set with a possible test trial. If trained on A/C and B/D and given A/D as a test trial, an “A” key press would indicate that A/D is being classified in the same way as A/C, while “L” would indicate that it was being classified in the same way as a B/D.

Once testing was complete, participants were debriefed. The experimenter asked them i) what rule they learned, and ii) if they were in a priming condition, what they thought the experiment was about.

Results: No participant made an explicit connection between the priming and the category-learning task when asked the second debriefing question. One participant did respond with “maybe something to do with eye movements” because he admitted to knowing that the affiliated lab conducts eye-tracking work. While he did not make a connection between the prime and the task, his data were eliminated.

With regard to rule learning, participants were considered to have learned a rule if they made no more than 3 inconsistent responses across the 14 novel stimuli during the test trials. For example, if they classified 11 of the novel exemplars by the “horizontal” rule, they were considered to be horizontal-rule-

learners; however if they classified 10 by the “horizontal” rule, and 4 by “vertical” rule, they were classified as no-rule-learners. This criterion means that they were expected to have a 78.6% accuracy rate to be considered as having learned a rule. The only exception was in the case of dual-rule learners (i.e., those that were considered to have learned both rules) since their data would look analogous to participants that learned nothing. As a result, we relied upon the debriefing answers such that participants were considered to have learned both rules if and only if they i) reported having learned both rules, and ii) when they made no more than three classifications inconsistent with that reported rule. Participants that did not learn any rule up to criterion were eliminated from subsequent calculation.

ANOVAs showed no significant difference on how quickly participants classified novel stimuli or on how many training trials were required for learning between conditions. However, a global Chi-Squared did show that a significant number of participants learned the rule that was congruent with the prime that they received ($\chi^2(4)=10.433, p<.05$) (See Table 1 and Figure 4).

	Control Condition	Vertical Priming	Horizontal Priming
Horizontal Rule Learned	13	7	15
Vertical Rule Learned	7	17	9
Both Rules Learned	11	5	8
No Rules Learned	4	5	4

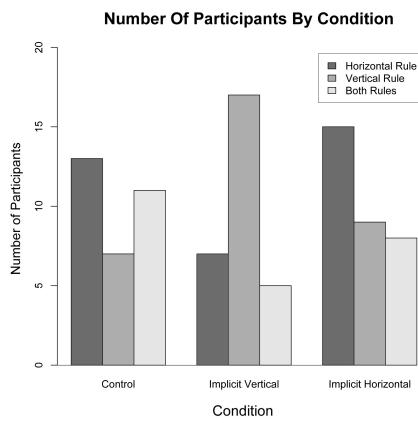


Table 1 (above) and Figure 4 (below): Number of participants that learned each rule, organized by condition.

Discussion: The results of the chi-squared suggest that relational category-learning can be affected by priming. As a result, the greater theoretical claim that relational reasoning can be primed with visuospatial movement

also seems to be supported. Thus, while it may be the case that an overly ambitious claim like “all analogical cognition is priming” is problematic, it also appears that, in the case of spatial relations, relational concepts can be primed with a relatively subtle prime.

That said, this study does raise further questions. First, it now seems important to ask *how* relations are primed. Relational reasoning has traditionally been described as a combination of steps (access, mapping, transfer and evaluation) (Kokinov & French, 2002); this experiment does not comment on which of these stages has been affected (though it seems logical to expect that it was access and/or mapping).

Secondly, it appears from the distribution of the control condition that there is a horizontal bias for ambiguously horizontal/vertical stimuli. It seems imaginable that this bias could have been due to the horizontal location of the keys used for response (“A” and “L” are in horizontal alignment on a standard keyboard). As a result, inadvertent priming could be a concern when developing relational priming paradigms, and future research may investigate whether this bias changes with a different response mechanism.

Thirdly, we predicted that participants would learn one relational category when two were present if one rule was sufficient for completing the task. This prediction was true for the majority of participants, however, dual-rule learning was somewhat common (especially in the control condition). This trend was likely due to the simplicity of the task, and may disappear if the task were made more complicated and working memory taxed to a greater degree. Future research may also focus on the effects of priming under more complicated tasks in order to explore the relationship between complexity, working memory, and relational priming.

Finally, we must question whether the results of this study would be applicable to all relations. To the point, this experiment used simple spatial relations, however, relations vary in their abstraction levels, and it seems possible that more abstract relations like “ameliorates” may be more difficult to prime (or, perhaps, more or less susceptible to a different kind of prime). Future research may exploit the paradigm presented here, but vary the types of relations used.

Ultimately though, this experiment suggests that priming relations is a complicated issue. It may not be the case that relational reasoning is entirely priming, however, it does appear that at least some relational cognition can be primed more easily than the literature indicates. Reliance upon physical input may help to explore the boundaries of this phenomenon, and help to specify how relations relate to real-world experiences.

References

Clapper, J.P. (2009). Category learning as schema induction. In M.A. Gluck, J.R. Anderson, & S.M. Kosslyn (Eds.), *Memory and Mind: A Festschrift for Gordon H. Bower* (pp. 307-326). New York: Lawrence Erlbaum Associates.

Doumas, L. A. A. & Hummel, J. E. (2005). Modeling human mental representations: What works, what doesn't, and why. In K. J. Holyoak and R. G. Morrison (Eds.), *The Cambridge Handbook of Thinking and Reasoning* (pp. 73-91). Cambridge, UK: Cambridge University Press.

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.

Doumas, L. A. A. & Richland, L. E. (2008). Developing structured representations. *Behavioral and Brain Sciences*, 31, 384-385.

Dunker, K. (1945). On Problem Solving. In: *Psychological Monographs*. 58, 270.

Falkenhainer, K.D., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of the Meeting of the American Association for Artificial Intelligence*, 272-277.

French, R.M. (2008). Relational priming is to analogy-making as one-ball juggling is to seven-ball juggling. *Behavioral and Brain Sciences*, 31, 386-387.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7.2, 155-170.

Gentner, D., & Kurtz, K. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman & P. W. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 151-175). Washington, DC: APA.

Gentner, D., & Namy, L. L. (2006). Analogical processes in language learning. *Current Directions in Psychological Science*, 15(6), 297-301.

Grant, E.R., & Spivey, M. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, 14(5), 462-466.

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.

Hummel, J.E., & Holyoak, K.J. (2008). No way to start a space program: Associationism as a launch pad for analogical mapping. *Behavioral and Brain Sciences*, 31, 388-389.

Leech, R., Mareschal, D., & Cooper, R.P. (2008). Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences*, 31, 357-414.

Richardson, D.C., Spivey, M.J., Edelman, S., & Naples, A.D., (2001). “Language is spatial”: Experimental evidence for image schemas of concrete and abstract verbs. In *Proceedings of the 23rd annual meeting of the cognitive science society* (p.p. 873-878). Mawhah, NJ: Erlbaum.

Richardson, D.C., Spivey, M.J., Barsalou, L.W., McRae, K. (2003). Spatial representations activated during real-time comprehension of verbs. *Cognitive Science*, 27, 767-780.

Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush?: Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, 62, 913-933.

Spellman, B.A., Holyoak, K.J., & Morrison, R.G. (2001). Analogical priming via semantic relations. *Memory and Cognition*, 29, 383-393.

Spivey, M.J. (2007). *Continuity of Mind*. New York: Oxford University Press.

Spivey, M.J., Richardson, D.C., & Dale, R. (2009). The movement of eye and hand as a window into language and cognition. In E. Morsella, J.A. Bargh, & P.M. Gollwitzer (Eds.), *Oxford Handbook of Human Action* (pp. 225-249). New York: Oxford University Press.

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.