THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Shaping Relations: Exploiting Relational Features For Visuospatial Priming

Citation for published version:

Livins, KA, Doumas, L & Spivey, M 2016, 'Shaping Relations: Exploiting Relational Features For Visuospatial Priming', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 42, no. 1, pp. 127-139. <https://doi.org/10.1037/xlm0000149>

Digital Object Identifier (DOI):

[10.1037/xlm0000149](https://doi.org/10.1037/xlm0000149)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Experimental Psychology: Learning, Memory, and Cognition

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Version accepted in *JEP:LMC*. Author version.
Please do not cite.

**Shaping Relations:
Exploiting Relational Features For Visuospatial Priming**

**Submitted to The Journal of Experimental Psychology:
Learning, Memory and Cognition**

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 (alex.doulas@ed.ac.uk)

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

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

While relational reasoning has been described as a process at the heart of human cognition, the exact character of relational representations remains an open debate. Symbolic-connectionist models of relational cognition suggest that relations are structured representations, but that they are ultimately grounded in feature sets; thus, they predict that activating those features can affect the trajectory of the relational reasoning process. The present work points out that such models do not necessarily specify what those features are though, and endeavors to show that spatial information is likely a part of it. To this end, it presents two experiments that used visuospatial priming to affect the course of relational reasoning. The first is a relational category-learning experiment in which this type of priming was shown to affect which spatial relation was learned when multiple were possible. The second used crossmapping analogy problems, paired with this same type of priming, to show that visuospatial cues can make participants more likely to map analogs based on relational roles, even with short presentation times.

Keywords: relational reasoning, symbolic-connectionism, analogy, visuospatial priming

Relational reasoning is pervasive. It has been implicated in a diverse set of processes including analogy-making (Gentner, 1983; Doumas & Hummel, 2005), inductive generalization (Hummel & Holyoak, 2003), linguistic processing (Gentner & Namy, 2006) and even social cognition (Spellman & Holyoak, 1992). As a result, there has been field-wide interest in how relational cognition functions and how relations may be represented in a cognitive system (e.g., Falkenhainer, & Gentner, 1986; Hummel & Holyoak, 2003; Doumas, Hummel, & Sandhofer, 2008).

Most accounts of relational reasoning treat relations as symbolic structures, which are represented as wholes. For example, relations are often coded with propositional notation (e.g., *predicate(argument, argument)*), and are typically described in terms of the structures that they afford more than the features that they possess (for example, as *k*-place relations, more than being “about” certain topics). Thus, while *chases*(cat, dog) may be described in terms of “a cat chasing a dog”, where the cat is the actor and the dog the patient, few models specify what it means to be in a chasing relationship, let alone what it means to be a chaser or a chased thing.

Gentner’s (1983) Structure-Mapping Theory and its computational instantiation, the Structure Mapping Engine (Falkenhainer, Forbus, & Gentner, 1989), exemplify this type of approach. They represent knowledge as a set of propositional networks and relations as predicates within them. Relational reasoning proceeds by matching tree structures using production rules (e.g., “if x, then y”). Thus, “chasing” is simply represented as “chasing”, and so, while features or properties can be *connected* to a relation within the network, there are no *inherent* features to it beyond that label. Thus, while these sorts of representations are extremely powerful, and excellent for completing complex analogical

mappings, it is difficult to imagine how semantic content could be an *inherent* part of a relation's representation (i.e., how a relation might be composed of such content). As a result, many opponents of the view have pointed out the symbolic approach struggles to offer an explanation of how relational content might be learned in the first place (e.g., see Leech, Mareschal, & Cooper, 2008; O'Reilly, Busby, & Soto, 2003).

As a reaction to this stance, some other approaches have argued against symbolic representation entirely. For instance, Leech, Mareschal, and Cooper (2008) suggested that relations are represented as associations in a simple recurrent network. According to this model, given specific objects as input, context objects prime particular association states that allow the model to produce transformed outputs. For example, imagine that the model was trained on the input “bread” in the context of “knife” to produce “cut-bread” as output, and that it was then also trained on “lemon” in the context of “knife” to produce “cut-lemon” as output. If the model were then presented with an input-output pairing on “bread” and “cut-bread”, it would settle on a context state such that when “lemon” was presented as input, the model would produce “cut-lemon” as output. By this account, the effects of relational reasoning might be nothing more than very clever priming.

This account does bypass some of the problems associated with symbolic accounts—namely, it does not have to deal with the problem of learning structured representations of relations because it does not include structured representations at all. However, it has been widely criticized for being unable to account for the types of behavior that are characteristic of both child and adult relational reasoning. For example, Doumas and Richland (2008) point out that because the model has no way of temporarily binding

objects to relational roles, that it would struggle to integrate multiple relations at a time. That is, it is unclear how the model would deal with mapping *chases*(x,y) and *chases*(y,z) to *follows*(a,b) and *follows*(b,c), even though humans over the age of five routinely solve these sorts of problems (e.g., Richland, Morrison, & Holyoak, 2008). Likewise, French (2008) and Holyoak and Hummel (2008) point out adult humans often make far-reaching analogies across novel analogs that share few semantic characteristics; however, such an ability would be beyond the model's capabilities due to its reliance on heavy training and rigidly structured problem sets. They argue that the model can only be successful on specific (carefully designed) relational problems and so should not be thought of as a generalized explanation of relational reasoning across contexts and content types.

Standing in more of a middle ground, symbolic-connectionist accounts of relational reasoning posit that relations must possess both features *and* structure. For instance, the DORA model, proposed by Doumas, Hummel and Sandhofer (2008), suggests that structured representations of relations are learned from unstructured feature vectors, and are eventually realized (at least in part) by sets of feature nodes firing in particular temporal patterns. Specifically, DORA (and its predecessor LISA; Hummel & Holyoak, 1997, 2003), posits that relational representations are coded across layers of nodes. In the bottom layer a set of distributed features encode objects and relational roles in a distributed fashion. One layer up, localist nodes combine sets of these features to represent particular objects and relational roles. Those roles are then temporarily bound to objects to create more complex relational structures. For example, features such as “feeling”, “happy” and “vulnerable”, may combine to represent *lover*, while “desired”, “beautiful” and “content” may combine to create *loved*; *lover* and *loved* may then be

bound to objects such as “John” and “Mary” (i.e., *lover*(John) + *beloved*(Mary)), which may ultimately combine to create a propositional structure such as “John loves Mary” (or, *loves*(John, Mary)).¹

Ultimately, these models account for a wide range of phenomena from developmental and adult cognition (see e.g., Dumas & Hummel, 2010, 2013; Dumas et al., 2008; Hummel & Holyoak, 1997, 2003; Lim, Dumas, & Sinnett, 2012, 2014; Morrison, Dumas, & Richland, 2011; Morrison et al., 2005; Sandhofer & Dumas, 2008; Son, Dumas, & Goldstone, 2011). Interestingly, they also provide a potential account of relational priming. Generally, priming is understood as a process wherein exposure to a piece of information facilitates later use of that information, or of a related concept (see e.g., Schunn & Dunbar, 1996). Despite the fact that relational cognition is typically described as a structural endeavor (e.g., Gentner, 1983), relational priming has often been explored by exploiting semantic links between a prime and a primed relation (e.g., Schunn & Dunbar, 1996; Spellman, Holyoak, & Morrison, 2001). Symbolic-connectionist models like DORA represent relations as structures that take arguments, but also as distributed patterns of features. As a consequence, priming might be a matter of activating some subset of those features and allowing activation to spread.

¹ In LISA binding is coded by firing bound roles and fillers in synchrony (e.g., to bind *lover* to John and *beloved* to Mary, the units for *lover* and John fire together, followed by the units for *beloved* and Mary). In DORA, binding information is carried by systematic asynchrony of firing. During learning roles are bound to their fillers by close proximity of firing, with bound roles and fillers firing in sequence. So, to bind *lover* to John and *beloved* to Mary, the units for *lover* fire, followed by the units for John, then the units for *beloved* fire, followed by the units for Mary. For other processes, the asynchrony may be established at the level of role-bindings (such that bound roles and fillers fire together, as in LISA), or even whole propositions (see Dumas et al., 2008 for details).

This account makes the prediction that the specific features of a relation, and their instantiations (i.e., the ways that they are depicted) matter in a priming context—in other words, if one does not prime the correct features, then priming should have little to no effect. For example, if “*chases*” has features like “running”, “horizontality”, and “movement”, then priming a horizontal directionality should have a positive priming effect, while priming a vertical directionality should have little-to-no effect, or perhaps even a negative effect.

This prediction may explain some conflicting experimental literature. On one hand, Spellman, Holyoak and Morrison (2001) found evidence to suggest that relational priming is possible, but that it is rare and requires explicit instruction and an ideal context. They demonstrated this point with an experiment that attempted to prime relational concepts on a lexical decision task. The task required participants to view pairs of letter strings then decide whether the strings were English words or not. String pairs occasionally included pairs of items that exemplified a given relation (e.g., “bird” and “nest”, which are typically associated through a “lives-in” relationship). Later word pairs could exemplify the same relations embodied in those earlier relational pairs (e.g., “bear” and “den”, which also typically exemplify the “lives-in” relation); it was expected that if relations can be primed, then participants should have been faster to classify those later word pairs that exemplified a previously seen relationship. However, a priming effect only occurred when participants were explicitly told to pay attention, not only to the relationship between the words within each pair, but also the relationships between pairs. In other words, it was only observed when similar structures and items were explicitly

highlighted, and so priming was possible, but difficult to achieve and limited by the experiment's instructions.

On the other hand, Bassok et al. (2008) explored the relationship between semantic and arithmetic relations and found evidence to suggest there is an obligatory activation of addition facts (e.g., $2+6=8$) when problems are paired with semantically aligned word pairs (e.g., “tulips-daisies”, which additively create “flowers”). They argued that their findings relied on priming that did not involve explicit instructions (like those required in the study performed by Spellman et al.), and so it must have occurred automatically.

Bassok et al. (2008) suggested that the difference in priming effects might have resulted from methodological differences (i.e., the two experiments used different relations, different experimental paradigms, different instructions, etc.). It is also possible that the primes selected by Bassok et al. simply activated more appropriate feature sets for the relations that they were attempting to prime. Both explanations seem reasonable, especially since they boil down to the claim that the effectiveness of relational priming will depend on the prime, the task, and the stimuli.

If this explanation correctly describes the incongruences in the literature, then priming becomes an issue of content (i.e., we must determine how to activate a given relation's features). However, this issue exposes a deficit in even the existing symbolic-connectionist models: they do not currently explicate which features, or even which types of features, may compose which relations. As a result, they cannot fully specify the effects of activating those features, and so they cannot predict the specific effects of relational priming. Thus, this paper will explore this issue of content, and present two experiments that suggest that perceptual information, specifically visuospatial

information, can have a profound effect on the trajectory of relational reasoning. It will ultimately be argued that this type of information may contribute to relational feature sets.

Our project begins with what is already known about relational content. First, Gentner (1978) pointed out that relational verbs carry different meanings from nouns, because such verbs convey relationships among entities, rather than information about concrete things. Additionally, Gentner (1981) pointed out that many verbs are semantically related, and that such semantics can influence processes such as memory recall. Thus, some set of relational verbs might share enough semantic content to be susceptible to a similar *type* of priming.

Thinking about *image schemas* might help to hone in on this content. Image schemas are generally thought of as primitive structures, which are inherently part of a concept and are derived through culture and worldly interaction (Dodge & Lakoff, 2005; Lakoff & Johnson, 1980; Mandler, 1992). For instance, lifting might have an inherently vertical schema because instances of lifting generally involve a vertical movement. Work in psycholinguistics has been inspired by this idea, and has shown that perceptually-coded visuospatial information (such as vertical or horizontal directionality) may be a natural part of the meanings of certain verbs. For example, Richardson, Spivey, Barsalou and McRae (2003) used both an attention task and a memory-recall task to show that presenting visual stimuli in orientations (vertical or horizontal) that were consistent or inconsistent with the orientation of a verb's meaning affected the speed with which participants completed the task (see also Bergen, Matlock, Lindsay, & Narayanan, 2007). These results suggest that visuospatial alignments (or image schemas) not only affect

how people represent action verbs, but that they also affect how those verbs are processed in tasks where space is functionally irrelevant (also see Toskos et al., 2004 for a similar effect on verb memorization). While not all verbs are relational, Richardson et al. used verbs that are, and so each specified an actor and a patient (e.g., *pointed at*, *pushed*, *lifted*, and *argued with*). As a result, these findings suggest that at least some relations may have visuospatial features associated with them that can be primed by sensorimotor processes such as visual attention and eye movements.

Interestingly, other work has exploited similar perceptual and spatial primes in other domains and found it rather effective. For example, Pedone et al. (2001) showed that one's ability to solve the Duncker Radiation Problem (Duncker, 1945) could be impacted, not only by diagrammatic differences in tasks preceding the problem's presentation (e.g., see Gick & Holyoak, 1983), but also by animating those diagrams. For example, they found that an initial task involving animated converging arrows was followed by a greater chance of finding the solution to the Duncker Radiation Problem since that problem also involves the concept of convergence. They suggested that such animations could alter diagram interpretation and increase the likelihood that it was noticed as a useful source analogue. This may be thought of as automatic priming since the first task was seemingly unrelated to the Duncker Radiation Problem, and yet changes in its animation altered success rates.

Grant and Spivey (2003) furthered this work when they found that the problem's solution of converging lasers could also be primed by inducing a converging pattern of attention and eye movements over a diagram of the problem (also see Thomas and Lleras, 2007). Their results suggested that the probability of producing a spontaneous solution

(without the use of an explicit analogy) could be affected, not just with symbolic content (such as arrows and animations as primes), but also by one's sensorimotor interaction with a diagram corresponding to the problem itself. If visual patterns can prime the visuospatial solution to a famously difficult insight problem, perhaps they can also prime relational representations of visuospatial relationships.

We argue that this research by Richardson et al. (2003), Pedone et al. (2001), and by Grant and Spivey (2003), is suggestive that visuospatial information might be important to the representational structures of relations, and that exploiting it for the purposes of priming might be useful. However, their tasks are not inherently relational enough to demonstrate that visuospatial information can affect the relational reasoning process (i.e., that the activation of such features can have a robust effect on structural tasks). It is important to note that the qualities of relational tasks are contentious (e.g., see Penn, Holyoak, & Povinelli, 2008), however existing literature suggests that relational representations must have a number of qualities that can be mapped to a task. Specifically, it has been argued that a relation must be represented explicitly, such that it can take novel arguments to which it can be dynamically bound (Doumas et al., 2008; Doumas & Hummel, 2005). These qualities mean that within an experimental context, reasoners must be able to show that they are reasoning about the roles that objects play rather than object properties, and that those objects (and their properties) can change. For example, crossmapping analogy problems use the same objects in different roles, and require a reasoner to ignore those statistical regularities in favor of role-based properties (e.g., an analogy task involving *chases*(dog, cat) and *chases*(cat, dog) should produce a mapping between the first dog and the second cat, and not between the two dogs). By

extension, a relational task should also allow the reasoner to demonstrate flexibility, such that if the objects change (i.e., they are replaced with other objects) the reasoner will still recognize the relation. Experimentally, this can be represented across trials or exemplars.

Thus, the following experiments draw on the visuospatial manipulations employed in Grant and Spivey (2003) and Richardson et al. (2003) and apply them to reasoning tasks that explicitly require the use and manipulation of relational representations. Given that the degree to which relations can be primed with these sorts of methods is currently unknown, the first experiment begins by trying to prime simple, obviously spatial relations (i.e., *above/below* and *left-of/right-of*), while the second attempts to prime more abstract relational verbs (e.g., *chases*, *drops*, etc.) using analogical crossmapping problems. In other words, our overall interest is in whether visuospatial priming can have an effect on the relational reasoning process and, if it can, whether it can equally affect relations with varying degrees of abstractness. It is important to note that if the answer to both questions is affirmative, then the data from these experiments will support the prediction that relational representations not only have featural content that can be exploited, but that such content may include some amount of spatial information.

Experiment 1

This experiment employed a pictorial relational category-learning task to determine whether simple, spatial relations can be primed by a subtle visuospatial prime that may capture exogenous attention. It did so by using relationally ambiguous exemplars that simultaneously belonged to two unique relational categories, where learning either category would suffice for successful classification. Visuospatial priming was congruent

with one category and incongruent with another, and priming was designed to affect which category was learned.

The task required participants to learn a relational category over the course of multiple exemplars. The exemplars used two-dimensional shapes positioned such that one shape always occluded the other. The categories were defined by the occluding shape's relative location to the occluded shape on the x- and y-axes (i.e., whether the occluder was to the left or right of the occluded shape, and whether it was above or below the occluded shape). However, it is important to note that while the exemplars involved shapes, the object attributes of those shapes were non-predictive of category membership – only the location of the occluding shape denoted membership.

The fact that the specific shapes were not predictive of category membership means that our paradigm meets the specified criteria for a relational task. Gentner and Kurtz (2005) pointed out that while not all categories are relational, some are. Specifically, relational categories define membership based on some common relational structure instead of the object attributes exhibited by members. 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. In other words, relational categories are not dependent on specific objects, but on the roles that objects play; it then follows that thinking about them involves using the same cognitive mechanisms as other types of relational cognition. Thus if it is possible to prime category learning on a relational category-learning task, then it is possible to claim that relational reasoning can be primed in a more general sense.

Participants

Participants were 106 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. Fourteen of these participants were not included in the final statistical analyses because of a lack of any rule learning, however they were used to calculate the sample's overall ability to complete the task.

Categories

As previously mentioned, categories were created using circles and squares and their relative placement on the x- and y-axes. More specifically, every exemplar showed two shapes, where one occluded the other; the specific shapes were selected at random at the beginning of each trial such that each trial could contain two circles, two squares, or one of each, and one shape always occluded the other. A pair of shapes was thought of as an “above” configuration if the occluder was above the occluded shape, a “below” configuration if the occluder was below the occluded shape, a “left-of” configuration if it was to the left, and a “right-of” configuration if it was to the right (see Figure 1).

[FIGURE 1 ABOUT HERE]

Every shape-pair simultaneously took a value on both the “left-of/right-of” relation and on the “above/below” relation, thus creating relationally ambiguous stimuli. As a result, stimuli could depict an “above/left-of” configuration that depicted an occluder above and to the left of the occluded shape, a “below/right-of” configuration that depicted

an occluder to the bottom and to the right of the occluded shape, an “above/right-of” configuration that depicted an occluder to the top and to the right of the occluded shape, or a “below/left-of” configuration that depicted an occluder to the bottom and to the left of the occluded shape (see Figure 2).

[FIGURE 2 ABOUT HERE]

It is important to note that this experiment worked under the expectation 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 was 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 should typically stop working when they have a sufficient answer.

Priming

The primes were made up of white circles with a black outline that were 150-pixels in size. The circles were presented in either a vertical or horizontal fashion. If the prime was a horizontal prime, then those circles appeared horizontally aligned along the middle of the screen; if the prime was a vertical prime, then those circles appeared vertically aligned along the middle of the screen. In both cases, the circles were spaced 540 pixels away from each other, spread out around the center in the specified direction.

Priming would begin with one circle blinking on for 500 ms, then blinking off. There would then be a 100 ms delay, then the other circle would blink on for 500 ms on the opposite half of the screen before also blinking off. Priming proceeded by cycling back and forth between those circles in this way (see Figure 3). The vertical prime was designed such that tracking the circles would require vertical saccades and therefore prime the “above/below” relation, while the horizontal prime would require horizontal saccades and therefore prime the “left-of/right-of” relation. It is important to note that participants were not told to watch the circles. However, participants were left alone with no distractions. Thus, while we cannot confirm that they visually tracked the circles, it was expected that the visuospatial prime might capture their exogenous visual attention.

[INSERT FIGURE 3 ABOUT HERE]

Procedure

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

They were told that they would see pairs of shapes, and that each pair would be positioned according to a “rule”—they were also told that they would not be told the rule. 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 an answer was entered.

Participants in the priming conditions were also told that they might occasionally see “blinking dots”, which were just the computer attempting to generate the next set of stimuli.

Participants began with a “training phase” of the task. If participants were in a priming condition, this phase began with five iterations of priming in the condition-appropriate direction. Priming was repeated after every five trials.

During this phase, all participants saw a fixation cross for 1500 ms, then an exemplar. The training phase randomly selected a pair of “training rules” in order to conflate a relative location on the horizontal axis with a relative location on a vertical axis. Thus the training phase would include “above/left-of” and “below/right-of” pairs, or “above/right-of” and “below/left-of” pairs. One pair would be randomly associated with the “*A*” key, and the other to the “*L*” key, however the keys were described as representing when shapes “followed the positioning rule” or “did not follow the positioning rule”. Participants would then press a key for every exemplar, and “*Correct*” or “*Incorrect*” would follow each press.

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 “above/left-of” and “below/right-of”, where “above/left-of” 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, the

visuospatial 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 assortment of exemplars representing the two training rules. For example, the training condition could proceed presenting eight exemplars of “above/left-of” followed by a random sequence of “above/left-of” and “below/right-of”. This training regiment was selected based on Clapper (2009), who claimed that this sort of presentation would increase ease of learning in dichotomous category-learning tasks. The initially presented rule was counterbalanced across participants in each condition.

Once the initial 8 training trials were complete, the experiment began counting each participant’s correct responses. Participants continued to see pairs of shapes (and get feedback) until they learned a rule well enough to correctly classify 10 exemplars in a row; however, the counter reset to zero if the participant answered any trial incorrectly.

When participants reached criterion, they were told that they would continue to see pairs of shapes, but that all feedback as to whether they were correct would stop. The test phase of the experiment then began. If a participant was in a priming condition, priming was stopped.

Participants were then presented with a random order of seven exemplars of each possible variable combination (i.e., “above/left-of”, “above/right-of”, “below/right-of”, and “below/left-of” 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 “left-of/right-of” relation with a

value on the “above/below” 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 she learned the “above/below” or “left-of/right-of” rule.

For example, suppose a participant was trained on “above/left-of” and “below/right-of”, where “above/left-of” was associated with an “*A*” key press, and “below/right-of” was associated with an “*L*” key press. “above/right-of” and “below/left-of” pairs could be used to determine which rule the participant had learned: If presented with an “above/right-of” pairing, then an “*A*” key press would indicate that the participant was classifying the stimulus like an “above/left-of” pair. If “above/left-of” and “above/right-of” pairs were classified in the same way, then the participant must have attended to the “above/below” relation (since “Above” is the common relational value between them). Conversely, an “*L*” key press would indicate that the participant had classified by the “left-of/right-of” rule (see Figure 4).

[FIGURE 4 ABOUT HERE]

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 visuospatial priming and the category-learning task. 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 “left-of/right-of” rule, then they were considered to be horizontal-rule-learners; however if they classified 10 by the “left-of/right-of” rule, and 4 by the “above/below” rule, then they were classified as no-rule-learners. The only exception was in the case of dual-rule learners (i.e., those who were considered to have learned both rules): because the task instructions associated one key with exemplars that “followed the rule” and the other key with exemplars that “did not follow the rule”, dual-learners could produce data looked analogous to participants who learned nothing. As a result, we relied upon the debriefing answers such that participants were considered to have learned both rules if they i) reported having learned both rules, and ii) made no more than three classifications inconsistent with the combined rule reported (i.e., where a novel exemplar was classified a “did not follow the rule” exemplar). Participants who did not learn any rule up to criterion were eliminated from subsequent analysis.

An overall chi-squared showed a significant difference between conditions ($\chi(4)=10.433$, $p<.05$) (see Table 1), suggesting that the priming did have an effect. Interestingly though, the control condition showed a strong bias towards a horizontal rule, and so post-hoc testing showed a significant difference between the control condition and the vertical priming condition ($\chi(2)=8.1591$, $p<.05$), and a difference approaching

significance between the vertical and horizontal priming conditions ($\chi(2)=5.9297, p=.05$), but no difference between the horizontal and control conditions ($\chi(2)=0.8509, p=.65$).

[TABLE 1 ABOUT HERE]

Discussion

The results from this experiment showed that participants were more likely to learn whatever category rule they were primed to learn. Thus, they suggest that visuospatial priming can affect which relational category is learned when multiple are equally possible. They further suggest that relational reasoning can, in general, be primed by the axis along which the movement of visual attention is attracted.

Interestingly though, the data also showed other trends worth discussing. For example a number of rule-learners did not learn the “left-of/right-of” rule, nor the “above/below” rule, but instead learned some combination of the two (thus, we called them “dual rule learners”). One possible explanation (that is potentially contradictory to our original hypothesis) is that they explicitly learned both rules (e.g., they might have learned “up” and “left-of”). Alternatively, they might have learned some rule that conflates the two spatial locations into a single relation (e.g. some version of “up-left”). At this time, it is unclear as to which of these possibilities is the case. It is also unclear as to why such learners were especially prominent in the control condition (see Table 1). Future research may need to determine which possibility is more likely and why. That said, the answer to

this question is not central to our current research question—what is crucial is that such learners did not prioritize the primed relation over other possibilities.

Furthermore, we observed a bias towards horizontal rule learning in the control condition (see Table 1). While this result was also unexpected, it may be explained by existing literature. To the point, Chatterjee et al. (1995, 1999, 2001) found that people have a tendency to describe relational scenes with the actor to the left of the patient. A horizontal rule would allow reasoners to follow this tendency, while the vertical rule would explicitly violate it. Notably, this result is congruent with the idea that visuospatial factors are important for relational processing, and suggests that (primed or not) they shape what relation is recognized and subsequently learned.

That said, we recognize that the generalizability of these findings are limited: relations like *left-of* and *above* are reasonably simple relations that have an easily imaginable relationship with horizontality and verticality (respectively). However, some other relations seem more complex and less obviously related to spatial alignments. For instance, *chases* not only involves a relative location between an actor and a patient, but also movement, and the actor's intention behind that movement (i.e., that it is trying to catch the patient). As a result, it seems necessary to address whether more complex relations are primable to the same degree. Experiment 2 investigates this issue.

Experiment 2

The methodology from Experiment 1 was ideal for determining whether relations can be primed when those relations possess obviously different spatial schemas and an equal degree of abstractness. However, one stated goal of this paper is to test whether

visuospatial priming can affect a range of relations, varying in their level of abstractness; as a result, a different methodology must be used.

To this end, Experiment 2 was based on the fact that relational reasoning sometimes requires one to not only overlook object attributes, but also similarities between those attributes. For instance, if one is shown two cups sitting beside each other and asked how they are related, one will need to actively ignore the fact that both objects were cups in order to identify a “beside” relation. However, ignoring object attribute similarities can be difficult, and any number of factors can disrupt the process. For instance brain injury (Waltz et al., 1999), stress (Tohill & Holyoak, 2000), and time pressure (Waltz et al., 2000) have all been shown to increase the probability of focusing on object attribute similarities over relational ones.

Here, time pressure was exploited for the purpose of studying how and to what degree visuospatial priming can affect whether similarities across object attributes can be overlooked in favor of relational ones. The method was inspired by response-deadline studies on similarity judgments, which have suggested that people will make judgments in favor of object attributes at short time scales (somewhere between 700 and 1000 ms), but more relational judgments at longer time scales (Goldstone & Medin, 1994; Gentner & Markman, 1997). Thus, we had participants complete analogy tasks but varied the amount of time that the base analog was presented, and asked whether the amount of time required to make a relational mapping over a mapping based on object attributes (when both were present) could be manipulated with the sort of priming used in Experiment 1.

In order to both prime a variety of relations over the course of the experiment, and to disguise the presence of the priming, participants were told that they were taking part

in a study about multitasking in which they would be constantly switching between two unique tasks—a “ball counting” task and a “find the thing doing the same thing” task. Participants consistently completed one counting trial, which was actually the priming task, followed by one “find the thing doing the same thing” trial, which was actually the analogy task. The goal was to prime each relation with a visuospatial stimulus just before completing an analogy task involving that relation.

While the priming was analogous to that found in Experiment 1, the analogy task was unique and employed pictorial crossmapping analogy problems that were adapted from Richland, Morrison, and Holyoak (2006). Crossmappings are problems in which the objects in the base analog (the relation being mapped from) play different roles in the target analog (the relation being mapped to). For example, the relational statement “the dog chases the cat” specifies two elements (a dog and a cat) involved in a chasing relationship; the statement “the cat chases the dog” specifies the same objects, however those objects are playing opposite roles. As a result, one must ignore similarities involving object attributes in favor of roles in order to reason relationally. Our paradigm displayed problems of this nature, but limited the temporal exposure of the base analog.

The relations used were selected because they were expected to have underlying vertical or horizontal image schemas (as described in Richardson, Spivey, Edelman, & Naples, 2001; Richardson, Spivey, Barsalou, & McRae, 2003; Chatterjee, 2010; Meteyard, & Vigliocco, 2009). We reasoned that if spatial image schemas are part of a relation’s representation (i.e., that they are part of a relation’s features), then it should be possible to exploit them for the purposes of priming. Thus, like in Experiment 1, priming involved a vertical or horizontal visuospatial stimulus, and it was expected that these

primes would access those image schemas in order to promote more accurate relational mappings in congruent priming conditions across time scales.

Participants

Participants were 243 undergraduate students from the University of California, Merced. They were recruited through the school's online participant pool, SONA, and received course credit for participation. All participants had normal to corrected-to-normal vision. Data from 18 of those participants were not included in analysis due to an inability to sufficiently complete the task.

Analogy Stimuli

As previously noted, the stimuli consisted of pictorial scenes adapted from Richland, Morrison, & Holyoak (2006). Each contained six objects dispersed around a black and white, drawn image; all images were 720-by-450 pixels in size and included both living and non-living things. They were presented on a black background. All stimuli were normed by having two experimenters a) count the number of objects in each scene, and b) state the relation that was shown. Full agreement was found.

The experiment used 64 of these scenes depicting 32 different relations. Eight relations were used for training trials, 8 were used as filler items (shown in between target trials), 8 were relations thought to possess horizontal image schemas, and another 8 more were relations expected to possess vertical image schemas (see Table 2). Five of 32 verbs were taken from Richardson et al., (2003), where they had already been normed to show greater than 70% agreement in their image schematic orientations. To expand the set of verbs for this experiment, 27 additional verbs were included. These additional verbs were normed via a simple Mechanical Turk survey that asked 15 participants to

classify the verbs as “horizontal”, “vertical”, or “neutral”. The items used received greater than 70% agreement on their image schematic orientations.

Each relation was instantiated in two different images, creating a base analog (i.e., an image that was to be mapped from) and a target analog (i.e., an image that was to be mapped to). All analogy problems were created such that the base analog was shown in the top half of the screen, while the target analog was shown in the bottom half of the screen. The base analog had one item circled in red, while the target had the numbers 1 through 4 beside different objects, each representing possible answers. In key trials, the enumerated items included a relational match to the circled item, an object attribute match to the circled item, and two distracter items (though note that in filler trials, the enumerated items included a relational match and three distracter items) (see Figure 5 for examples).

[INSERT FIGURE 5 ABOUT HERE]

Priming

Like in Experiment 1, the primes were made up of 150 pixel-large circles with a thin black outline. Each round of priming involved a total of ten circles, a random number of which were colored red, while the rest were white.

The circles blinked on and off one at a time at specified locations on the screen. Again, like in Experiment 1, the circles were positioned across the screen from each other, such that if one looked from one circle to the next, it required a linear eye movement of a specific linearity. In key trials, the locations of the circles required

movements that were congruent or incongruent with the expected image schema of the depicted relation (horizontal or vertical). Filler trials, however, were randomly paired with a priming alignment since they were included only to ensure that participants did not make an explicit connection between the moving dots and the stimuli. Thus, these filler trials were randomly assigned a priming alignment at the beginning of every trial for every participant. They could be vertical, horizontal, or even diagonal.

Procedure

To start, participants were told that the experiment was about multitasking, and that two different tasks would be interleaved trial-for-trial. They were also told that they would always complete one “ball counting” problem, during which they would count the number of red balls shown in a given sequence, then they would switch to a “find the thing doing the same thing” problem where they would identify the item in the target analog that they thought was “doing the same thing” as the item circled in the base analog.

It is important to note that while telling participants to “find the thing doing the same thing” may seem heavy handed, analogy research generally suggests that relational cognition is difficult, and people often do not engage in it unless explicitly directed to do so (e.g., see Gick & Holyoak, 1980, 1983; Spellman, Holyoak, & Morrison, 2001). The goal of the instructions was to give participants a clear understanding of what a “correct answer” might look like. That said, we admit that these instructions may limit the degree to which this study can comment on free relational recognition.

The priming task was presented on a 1920-by-1080-pixel sized monitor positioned on the right hand side of the desk, while the analogy task was presented on a

2560-by-1440-pixel monitor on the left. The tasks were presented on different, and differently-sized, screens in order to avoid priming a specific location on the screen where the analogy problems would be shown.

The experiment design was a 2x5 between-subjects factorial design. Participants were randomly assigned to either a congruent or incongruent priming condition and to one of the five analog presentation times. So, for example, if a participant was assigned to the congruent 500 ms condition, and she was completing the “*chasing*” trial (where “*chasing*” is expected to have a horizontal image schema), then she would complete a ball counting task where the circles blinked on and off in a horizontal way, followed by the “*chasing*” analogy problem where the base analog would be displayed for 500 ms. If the next trial involved the “*lifting*” problem (where “*lifting*” is expected to have a vertical image schema), she would then complete a ball counting task where the circles blinked in a vertical way, then complete the “*lifting*” analogy problem where the base analog would again only be shown for 500 ms.

After participants were assigned to a condition, all participants had both tasks explained to them, and then the experimenter guided them through the 8 training trials. The experimenter provided verbal cues to switch tasks during this phase to ensure that the participants stayed in sequence. Cues included “switch computers”, or just “switch” after the initial training trial. The experimenter stopped providing cues all together when the participant was able to switch tasks on his or her own.

When training was complete, participants began the active trials, which were self-paced. The relations with horizontal image schemas, the relations with vertical image schemas, and the filler relations were randomly ordered for each participant.

Results

Key trial performance was measured by the number of crossmapping problems that were correctly answered with the relational mapping (over the featural mapping and other distractor items). This particular measure was used because it targets the degree to which the priming affected the probability of correctly identifying a relational mapping under both time pressure and featural distraction.

A two-way ANOVA showed that priming congruency did have an effect on overall accuracy on key trials ($F(1,224)=47.890$, $p<.01$). Likewise, those in the longer temporal intervals did better than those at shorter temporal intervals ($F(4,224)=3.976$, $p<.01$). There was no interaction between the condition and the presentation times ($F(4,224)=.069$, $p=.991$), however, planned comparisons by condition showed significant differences at all presentation times, with those in the congruent priming conditions doing better than those in the incongruent priming conditions (see Table 3 and Figure 6).

[INSERT TABLE 3 HERE]

[INSERT FIGURE 6 ABOUT HERE]

Another two-way ANOVA was run in order to test the generalizability of our stimuli post-hoc. Specifically, we ran an item analysis common within psycholinguistics, originally outlined by Clark (1973). This analysis involves running a two-way ANOVA where stimuli were treated as the random factor (instead of the participants). It showed similar trends to the reported participant statistics with regard to overall problem accuracy: The temporal presentation of the base was a significant factor ($F(4)=3.531$,

$p < .01$), along with the priming condition ($F(1) = 13.768$, $p < .01$). Again, there was no interaction ($F(4) = .146$, $p = .964$). We also calculated Chronbach's alpha (which is a test of internal consistency) where we pooled responses across times and priming conditions, and tested the reliability of our stimuli. The test suggested extremely high internal consistency ($\alpha = .961$). Given these two tests, we suggest that our stimuli are reliable.

Discussion

The findings of Experiment 2 demonstrate three things: First, they suggest that priming congruency affects the probability of selecting a relational mapping over a featural mapping and distractor items on a crossmapping analogy problem overall. Second, they suggest that longer presentation times of the base analog also affect the probability of selecting that featural mapping overall. And third (and perhaps most importantly), they suggest that priming will significantly increase the number of relational mappings made, not only at longer presentation times, but at shorter ones as well. In other words, visuospatial priming can alter the likelihood that one makes a relational match when one is under time pressure.

Furthermore, because Experiment 2 defined congruency based on the degree to which priming matched the expected image schemas of the targeted relations, it seems possible to claim that visuospatial priming can be used to target those image schemas (or at least, exploit them). It also seems possible to claim that this sort of priming works on a variety of relations with varying levels of explicit spatiality and abstractness (so, not just on spatial relations such as “above”, but also on relational verbs such as “chases”).

That said, this experiment does leave some open questions. Specifically, every relational scene (like the real world) depicts multiple relations simultaneously, however

the priming seems to have affected one over the others. For example, Figure 1A depicts a *chasing* relation in which both parties are also “running”, one element is “in front” of another, etc. It is unclear why one relation was prioritized over the others, however, our results do suggest that such a prioritization occurs. Interestingly, recent research on the correlations between relational recognition and specific eye movements found a similar effect, and suggests that some relations are more likely to be recognized, even in free verbal-description tasks (Livins, Doumas, & Spivey, submitted). An understanding why and how this phenomenon occurs will need to be provided in the future, especially since the factors of relational recognition (i.e., the process by which relations are recognized) are still not well understood (Livins & Doumas, in press).

General Discussion

The paper has had two goals. First, it has explored the relationship between featural content and relational reasoning. This project is important because that relationship is contentious and not well understood. For example, existing computational models seem to either fail to sufficiently account for the sorts of features that relational representations might possess (i.e., in the case of entirely symbolic models), or are so reliant on features that they fail to capture the structural power necessary to account for human relational performance (i.e., in the case of entirely non-symbolic models). Even the more progressive symbolic-connectionist models (like the DORA model, Doumas et al., 2008) that attempt to account for both relational structure and featural content seem deficient in that they currently fail to explicate what features (or even what types of

features) might be part of what relations. As a result, the existing literature is inconclusive.

This paper has furthered this debate by exploiting image schemas for the purpose of affecting reasoning about relational verbs. To the point, Experiment 1 found that priming such schemas could make a particular relation more likely to be recognized and then ultimately learned as a category when two relations are present and equally possible, while Experiment 2 showed that such priming could make the recognition of a relation more likely over featural consistencies and distractors in the base and target analogs. As a result, our data are consistent with predictions that relational features (and not just structure) are important to relational processing (e.g., Doumas et al., 2008; Hummel & Holyoak, 1997, 2003).

It is, however, important to note that these studies do not explicate the exact representational structure of relational features, nor of the mechanisms that implicate them in the reasoning process. For example, they suggest that such features might be implicated in relational recognition but do not specify the process by which it might occur. As a result, the influence that the priming exerted on the relational representations in these studies may be due to an internal representation of orientations as image-schematic features (e.g., Dodge & Lakoff, 2005; Mandler, 1992), or it may be due to a pattern of visual-attentional interaction with the environment that changes the way the critical analogy stimuli are processed in the first place, thereby activating different sets of features (e.g., Chemero, 2011; Spivey, 2007). Future work will need to determine whether one of these influences is more powerful, or whether they come together to create the priming effects observed here. This project will likely need to consider image-

schema defying presentations of relations (for example, while we classified “chasing” as a horizontal verb, it seems quite possible for a cat to chase a squirrel up a tree in a vertical direction) in order to specify exactly how priming changes the trajectory of the reasoning process. Such work can then be used to guide accounts of relational representation and recognition. That said, what our results *do* show is that visuospatial orientation information behaves as an integral part of that process, and that such information can be derived from the image schema literature. In other words, they show that visual and spatial information are factors in the process, and so they will need to be incorporated into any mechanistic account.

The second goal of this paper has been to further the relational priming debate. To the point, some research has suggested that such priming is possible but difficult (e.g., Spellman, Holyoak, & Morrison, 2001), while other work has suggested that it can occur more automatically (e.g., Bassok et al., 2008). While our instructions were heavy-handed (i.e., they told participants to look for a relational match), they did not tell participants that there was a relationship between priming and the task. This is a methodological distinction from Spellman et al. (2001), who told participants to not only look for relationships between words within an analog, but also relationships between analogs (which would be functionally equivalent to telling participants to look for a relationship between the priming and the task). Thus, the data here support the more automatic account by showing that simple, task-irrelevant eye-movements are sufficient for priming.

Furthermore, because these experiments exploited expected image schemas and associations between relations and space, the data also suggests that such automatic

priming will likely be an issue of stimulating the *right* sorts of content—i.e., content that has some meaningful link to the relation. The relations used in these studies had close links to space, and so space was particularly useful; however, it also seems imaginable that other sorts of featural content might be useful. For example, relations such as “*after*” might have a temporal component that could be exploited. While the results presented in this paper do not speak to this possibility, they do suggest a route for future work. In other words, more effort will need to be spent determining what relational content might look like before priming studies can be conducted (e.g., if one wants to prime a relation like “*causes*”, one might first need to spend time determining what features causes might have).

Additionally, while this paper did not explicitly set out to comment on the relationship between complex reasoning and sensory-motor processing (instead it focused more specifically on visual and spatial processing), such a commentary is a natural result of the experimental designs and results. To this point, our primes were visual and spatial in nature, suggesting that such perceptive information can impact the outcomes of complex, relational reasoning. While our data is not sufficient for claiming that symbols are grounded in sensory-motor information, it is *consistent* with some embodied perspectives. For example, Barsalou’s Perceptual Symbol Systems Hypothesis suggests that when one perceives a stimulus that one’s sensory states become stored within memory; when one recalls that stimulus (or one like it) that sensory information is recalled along with the stimulus itself (Barsalou, 1999). The DORA model (Doumas et al., 2008) suggests a similar possibility in the way that it learns relations. To the point, the model suggests that early exemplars and experiences are stored as feature-sets in

memory, and later information is compared to those memories; relations are essentially learned by comparing new information to stored information. While DORA does not specify what types of information might end up getting stored, it remains possible that sensory or spatial information will be included.

Currently, one unanswered question is whether the priming effects observed here are visual- or spatial- specific, or more generalized (for example, will moving one's hands prime chasing in the same way that moving one's eyes does?) To the point, we believe that answering this question might help to specify how sensory-motor information is linked/part of relational processing. For example, if priming is possible across modalities, then it might suggest that a more generalized set of spatial features are encoded during relational learning and that they are either closely linked to, or a part of, relational content. However, if the priming of specific relations is limited to certain modalities (e.g., if chasing is only primable with visuospatial cues), then it might suggest that specific modal information is crucial for their processing. This later possibility would support something like the Perceptual Symbol System perspective.

Thus, future research will need to include both computational and experimental work. The relations used here were both spatial (as in Experiment 1) and relatively simple (even in Experiment 2). Furthermore, their image schemas seem somewhat intuitive, but increasingly abstract and complex relations may be harder (or even impossible) to specify (for example, it seems more difficult to imagine the spatial features of a relation like "*derives*"). The current work does not test the limits of spatial priming in this regard, and future work will need to address this issue by using increasingly abstract and complex relations. Only then will computational models be able to fully account for

the priming effects seen here, and only then will it be possible to more clearly specify the relationship between the perception and relational cognition.

References

Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577-660.

- Bassok, M., Pedigo, S.F., & Oskarsson, A.T. (2008). Priming addition facts with semantic relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(2), 343-352.
- Bergen, B. K., Lindsay, S., Matlock, T., & Narayanan, S. (2007). Spatial and linguistic aspects of visual imagery in sentence comprehension. *Cognitive Science*, 31(5), 733-764.
- Chatterjee, A. et al. (1995) Spatial characteristics of thematic role representation. *Neuropsychologia* 33, 643–648
- Chatterjee, A. Southwood, M.H., & Basilico, D. (1999). Verbs, events and spatial representations. *Neuropsychologia*, 37, 395–402.
- Chatterjee, A., et al., (2001). Language and space: some interactions. *TRENDS in Cognitive Science*, 5(2), 55-61.
- Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition*, 2(1), 79-116.
- Chemero, A. (2011). *Radical embodied cognitive science*. MIT press.
- 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.
- Dodge, E., & Lakoff, G. (2005). Image schemas: From linguistic analysis to neural grounding. In B. Hampe (Ed.), *From perception to meaning: Image schemas in Cognitive Linguistics* (pp. 35-56). Berlin: Mouton de Gruyter.
- Doumas, 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.

- 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., & Hummel, J.E. (2010). Computational models of higher cognition. In K.J. Holyoak, and R.G. Morrison (Eds.) *The Oxford handbook of thinking and reasoning*. Oxford, England: Oxford University Press.
- Doumas, L.A.A., & Hummel, J.E. (2013). Comparison and mapping facilitate relation discovery and prediction. *PLOS One*, 8(6), e63889.
- Doumas, L. A. A. & Richland, L. E. (2008). Developing structured representations. *Behavioral and Brain Sciences*, 31, 384-385.
- Duncker, K. (1945). On Problem Solving. In: *Psychological Monographs*, vol. 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.
- Falkenhainer, K.D., Forbus, K.D., & Gentner, D. (1989). The Structure Mapping Engine: Algorithm and examples. *Artificial Intelligence* 41(1), 1-63.
- 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. (1978). On relational meaning: the acquisition of verb meaning. *Child Development*, 49, 988-998.
- Gentner, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, 13, 56-83.
- 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., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52(1), 45.
- 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. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Goldstone, R. L., & Medin, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(1), 29.
- Grant, E.R., & Spivey, M. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, 14(5), 462-466.
- Holyoak, K.J., & Hummel, J.E. (2008). No way to start a space program: Associationism as a launch pad for analogical mapping. *Behavioral and Brain Sciences*, 31, 388-389.
- 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. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3), 427-466.

- Kokinov, B., & French, R.M. (2002). Computational models of analogy-making. In L. Nadel (Ed.), *Encyclopedia of Cognitive Science* (pp. 113-118). London: Nature Publishing Group.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- 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.
- Lim, A., Dumas, L. A., & Sinnett, S. (2014). Supramodal representations in melodic perception. In *36th Annual Conference of the Cognitive Science Society, Quebec, Canada*.
- Lim, A., Dumas, L. A., & Sinnett, S. (2012). Modeling Melodic Perception as Relational Learning Using a Symbolic-Connectionist Architecture (DORA). In *34th Annual Conference of the Cognitive Science Society, Sapporo, Japan*.
- 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). Tracking relations: The effects of visual attention on relational recognition.
- Mandler, J.M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review*, 99(4), 587.
- Meteyard, L., & Vigliocco, G., (2009). Verbs in space: Axis and direction of motion norms for 299 English verbs. *Behavior Research Methods*, (41), 565-574.

- Morrison, R.G., Dumas, L. A., & Richland, L. E. (2011). A computational account of children's analogical reasoning: balancing inhibitory control in working memory and relational representation. *Developmental science*, 14(3), 516-529.
- O'Reilly, R.C., Busby, R.S. & Soto, R. (2003). Three Forms of Binding and their Neural Substrates: Alternatives to Temporal Synchrony. A. Cleeremans (Ed) *The Unity of Consciousness: Binding, Integration, and Dissociation*, 168-192, Oxford: Oxford University Press.
- Pedone, R., Hummel, J.E., & Holyoak, K.J. (2001). The use of diagrams in analogical problem solving. *Memory & Cognition*, 29, 214–221.
- Penn, D.C., Holyoak, K.J., & Povinelle, D.J. (2008). Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31, 109-178
- Richardson, D.C., Spivey, M.J., Edelman, S., & Naples, A.J., (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 (pp. 873–878). Mahwah, 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.
- Sandhofer, C.M., & Dumas, L. A. (2008). Order of presentation effects in learning color categories. *Journal of Cognition and Development*, 9(2), 194-221.
- Schunn, C.D., & Dunbar, K. (1996). Priming, analogy, and awareness in complex reasoning. *Memory & Cognition*, 24(3), 271-284.

- Son, J.Y., Smith, L. B., & Goldstone, R. L. (2011). Connecting instances to promote children's relational reasoning. *Journal of Experimental Child Psychology*, 108(2), 260-277.
- 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.
- Tohill, J.M., & Holyoak, K. J. (2000). The impact of anxiety on analogical reasoning. *Thinking & Reasoning*, 6(1), 27-40.
- Toskos, A.C., Hanania, R., Hockema, S. (2004). Eye scanpaths influence memory for spoken verbs. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. (pp. 1643). Chicago, Illinois.
- 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, 119-125.
- Waltz, J.A., Lau, A., Grewal, S. K., & Holyoak, K. J. (2000). The role of working memory in analogical mapping. *Memory & Cognition*, 28(7), 1205-1212.

Figure Description

Figure 1: An example of how two shapes could combine to create exemplars that had an occluder take a value on the “left-of/right-of” dimension or the “above/below” dimension. In both case the critical category-defining relationship was the placement of the occluder.

Figure 2: Examples of stimuli that combine a value on the “left-of/right-of” relation with a value on the “above/below” relation.

Figure 3: An example of how two priming cycles would progress over time (where time is depicted as movement from left to right).

Figure 4: An example of possible training set and test phase exemplars. Imagine the participant was trained on “above/left-of” and “below/right-of” exemplars, where an “A” key press was paired with “above/left-of”, and an “L” key press was paired with “below/right-of”. If that participant were then shown an “above/right-of” exemplar during the test phase, then an “A” key press would indicate that “above/right-of” was being classified in the same way as an “above/left-of” exemplar, while “L” would indicate that it was being classified in the same way as a “below/right-of” exemplar.

Figure 5: Three stimulus examples from Experiment 2. Figure 5a shows the horizontal verb *chasing* depicted such that *chasing*(cat, mouse) must be mapped to *chasing*(boy, cat). The answer of “boy” would make the relational mapping, while the answer of “cat” would make the featural mapping. Figure 5b shows the filler item *performing-for*, where *performing-for*(ring-master, audience) must be mapped to *performing-for*(boy, audience). Here, the audience looks different across images and no exact featural matches are present. Figure 5c shows the vertical verb *bombing*, depicted such that *bombing*(boy, girl) must be mapped to *bombing*(girl, monkey). Here, the answer “monkey” would make a relational mapping, while the answer “girl” would make the featural mapping.

Figure 6: A graphical representation of the number of relational mappings made in Experiment 2, organized by congruency condition. The x-axis represents the possible amounts of time that the base analog could be displayed, while the y-axis represents the raw number of questions answered. Error bars represent the Standard Errors.

Table 1: The number of participants who learned each rule, organized by priming type.

Table 2: The list of relations used in Experiment 2, organized by type (those with vertical image schemas, those with horizontal image schemas, those with neutral image schemas, and those used for training).

Table 3: Overall accuracy on key trials in Experiment 2, organized by condition. The final column shows the results from planned comparisons (protected t-tests with an alpha level of .01). Stars show comparisons where variances were unequal between groups, and the degrees of freedom were adjusted to correct for it.

Table 1:

	Control	Vertical Prime	Horizontal Prime
Horizontal Rule Learned	13	7	15
Vertical Rule Learned	7	17	9
Both Rules Learned	11	5	8
No Rules Learned	4	6	4

Table 2

Vertical	Horizontal	Fillers	Training
Pouring-on	Chasing	Kissing	Riding
Dropping	Pulling	Playing-with	Talking
Hanging-from	Pushing	Resting-on	Balancing
Hoisting	Kicking	Cooking	Feeding
Lifting	Towing	Cleaning	Sheltering
Reaching-for	Points-at	Driving	Scolding
Bombing	Hunting	Opening	Hitting
Climbing	Gives-to	Performing-for	Brushing

Table 3

	Congruent	Incongruent	Planned Comparison
400	M=10.38 SD=2.64	M=7.85 SD=3.133	t(39)=2.804, p<.01
500	M=10.41 SD=2.26	M=7.77 SD=3.94	t(33.479)=2.723, p<.01*
600	M=11.90 SD=1.92	M=8.91 SD=3.24	t(40)=3.598, p<.01
700	M=11.56 SD=1.92	M=9.13 SD=3.52	t(35.234)=2.991, p<.01*
800	M=12.17 SD=1.74	M=9.68 SD=2.84	t(40.007)=3.715, p<.01*