Cooperrider, K., Gentner, D., & Goldin-Meadow, S. (2016). Gesture reveals spatial analogies during complex relational reasoning. To appear in D. Grodner, D. Mirman, A. Papafragou, J. Trueswell, J. Novick, S. Arunachalam, S. Christie, & C. Norris (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Gesture reveals spatial analogies during complex relational reasoning

Kensy Cooperrider (kensy@uchicago.edu)
Department of Psychology, 5848 S. University Avenue
Chicago, IL 60637 USA

Dedre Gentner (gentner@northwestern.edu)
Department of Psychology, 2029 Sheridan Road
Evanston, IL 60208 USA

Susan Goldin-Meadow (sgm@uchicago.edu)
Department of Psychology, 5848 S. University Avenue
Chicago, IL 60637 USA

Abstract

How do people think about complex relational phenomena like the behavior of the stock market? Here we hypothesize that people reason about such phenomena in part by creating spatial analogies, and we explore this possibility by examining people's spontaneous gestures. Participants read a written lesson describing positive and negative feedback systems and then explained the key differences between them. Though the lesson was highly abstract and free of concrete imagery, participants produced spatial gestures in abundance during their explanations. These spatial gestures, despite being fundamentally abstract, showed clear regularities and often built off of each other to form larger spatial models of relational structure—that is, spatial analogies. Importantly, the spatial richness and systematicity revealed in participants' gestures was largely divorced from spatial language. These results provide evidence for the spontaneous use of spatial analogy during complex relational reasoning.

Keywords: analogy; relational reasoning; gesture; complex systems; spatial cognition

Introduction

Ecosystems in flux. Seesawing financial markets. Shifting climate patterns. What these diverse phenomena have in common is that they are all examples of complex relational systems: they involve multiple causal factors that change over time and bring about changes to other factors in the system. Such systems underlie phenomena throughout the natural and social world, in all domains and at all scales. Yet, despite the ubiquity and importance of these systems, much remains to be learned about the cognitive processes involved in understanding them.

Current evidence suggests that complex relational reasoning presents challenges even for adults. For example, undergraduates have considerable difficulty detecting higher-order causal patterns such as *positive feedback* and *negative feedback,* the focus of the present paper (Rottman, Gentner, & Goldwater, 2012). Expertise in identifying such patterns does develop, either through exposure to the same patterns across a range of domains (Rottman, Gentner, & Goldwater, 2012) or through a scaffolded process of comparing examples (Goldwater & Gentner, 2015). An interesting open question, however, concerns the nature of the representations that people form as they develop such abilities. What are these representations like and how do they differ from people's representations of other kinds of systems?

Possible clues may come from research on how people understand systems more generally. Much previous research has investigated how people understand mechanical processes with multiple causal components, such as sets of gears and pulleys. A major finding of this line of work is that people often develop mental models of the system that are visuospatial in nature (Hegarty, 2004). One line of evidence for the visuospatial character of these models is that when reasoning about such systems, people often produce diagrams (Forbus, Usher, Lovett, & Wetzel, 2011; Novick, 2001; Tversky, 2011) or gestures (e.g. Schwarz & Black, 1996; Nathan & Martinez, 2015). Based on such observations, it seems plausible that people develop mental models of other types of complex systems, such as the causal patterns under consideration here. However, there is a crucial difference between mechanical systems and positive and negative feedback systems. Positive feedback and negative feedback are consummate abstractions. They are relational patterns that may sometimes be instantiated in mechanical or concretely spatial systems—e.g. a flush valve toilet is an example of a negative feedback system—but their relational essence transcends any one concrete instantiation. It might thus seem unhelpful, or even counterproductive, to recruit visuospatial reasoning processes when thinking about such pure abstractions.

At the same time, a separate line of research has investigated how people recruit space when talking and thinking about purely abstract ideas. This tendency can be seen in everyday language, for instance in the spatial words and grammatical structures people draw on to talk about time (e.g. Clark, 1973; Traugott, 1978), or in the extension of spatial prepositions to describe abstract relations of other kinds (Jamrozik & Gentner, 2015). In fact, evidence has now accumulated that this is not just a linguistic phenomenon—people use spatial representations when reasoning online about abstract concepts, whether or not language is involved (Boroditsky, 2001; Casasanto & Bottini, 2014). One clear source of evidence for the use of

space in abstract reasoning comes from the gestures people produce (Cienki, 1998). To date, the best-studied cases of abstract spatial gesture have involved relationally simple concepts, such as the representation of a temporal sequence as a line (Cooperrider, Núnez, & Sweetser, 2014). Nonetheless, such findings raise the intriguing possibility that people might create more complex spatial structures in gesture to represent more complex relational structures.

The above observations lead us to the following hypothesis about how people reason about complex relational patterns like positive and negative feedback: they may do so, at least in part, by creating abstract spatial models of the relational structures involved—that is, spatial analogies. Furthermore, if this hypothesis is correct, then gesture should provide a powerful window onto this phenomenon. Gesture is well suited to the expression of spatial ideas (Alibali, 2005), and it has been shown to reveal implicit aspects of understanding that people have difficulty verbalizing (Goldin-Meadow, 2003; Broaders, Cook, Mitchell, & Goldin-Meadow, 2007). Moreover, the spatial information revealed in people's abstract gestures often goes beyond what is found in the language co-produced with those gestures (Cienki, 1998).

In the present study, we explore this spatial analogy hypothesis by having people read a lesson contrasting two types of complex relational patterns—positive and negative feedback—and then explain the key differences between them. The most interesting possibility is that gesture might reveal spatial analogies—that is, systematic spatial models of relational patterns that are not inherently spatial. We also considered other possible outcomes, however. For one, people might not spatialize much of anything in their explanations. After all, gesture is thought to stem from vivid visuospatial or motoric imagery (e.g. Hostetter & Alibali, 2008), which our lesson lacks. Another possibility is that people might spatialize in gesture, but in a piecemeal fashion. That is, they may occasionally produce abstract spatial gestures (e.g. an upward gesture when describing "increasing") but these gestures will not cohere into a larger model.

Methods

Participants

23 adults from the University of Chicago community participated for course credit or cash. Four participants were excluded from the analyses: three because their gestures were largely occluded on the video; one for producing no gestures at all. In all, data from 19 participants (10 female; mean age = 20.8 years) are reported in the analyses.

Materials and procedure

After giving consent to participate and to be videotaped, participants carried out a series of activities that served both to familiarize them with causal systems and to assess their understanding of them. First, participants completed an adaptation of the Ambiguous Sorting Task (AST) used previously by Rottman, Gentner, & Goldwater (2012) and Goldwater & Gentner (2015). In this type of sorting task, participants are given a set of vignettes printed on index cards and are asked to sort them into categories. Each vignette is an example of one of several types of causal systems (e.g. positive feedback) instantiated in one of several domains (e.g. economics). Participants are also given *seed cards*—vignettes just like those that need to be sorted but which serve as anchors for the categories to be used. A key feature of the task is that the seed cards leave the relevant categories open to interpretation: a participant may categorize the vignettes according to the type of causal system described or, more superficially, by the domain in which that system is couched. In the adaptation of the AST used here, participants were presented with three seed cards: a first unlabeled card describing the phenomenon of stock market bubbles (a positive feedback system), a second unlabeled card describing predator-prey relationships (a negative feedback system), and a third card simply labeled 'other.' Participants were then given 11 new vignettes and were given 5 minutes to sort them.

After the sorting was complete, the experimenter removed the materials and prompted the participant to explain the main difference between the different categories involved in the sorting task. This phase is the *pre-lesson explanation*. Together the sorting task and the pre-lesson explanation serve to familiarize participants with causal systems, which they will go on to learn more about and explain.

Next, participants were given a one-page written lesson ('Causal Systems Lesson') explaining the differences between positive and negative feedback systems (though without using those labels). The lesson was grounded in the seed cards used in the sorting task. It explained how the stock market vignette exemplifies one type of causal system and how the predator-prey vignette exemplifies a different type. The lesson also moved beyond the particular examples, characterizing in more abstract terms how each type of system involves different relationships between causal factors. Importantly, the lesson used no concrete spatial imagery and very little spatial language. Participants were instructed to study it for 3 minutes and were told that they would later be asked to explain it to another participant.

When the 3 minutes were up, the experimenter removed the lesson and brought in the other participant (who was actually a confederate). The experimenter then prompted the participant as follows: "Please explain the lesson you just read. Go into as much detail as possible, but focus on the differences between the two types of causal systems." The instructions made no mention of gesture. This phase is the *post-lesson explanation*, and it is the focus of our analyses.

Analysis

Participants' performance on the sorting task was analyzed but is not discussed in the present report. Videos of participants' pre- and post-lesson explanations were transcribed and analyzed using ELAN video annotation software (https://tla.mpi.nl/tools/tla-tools/elan/). The gesture

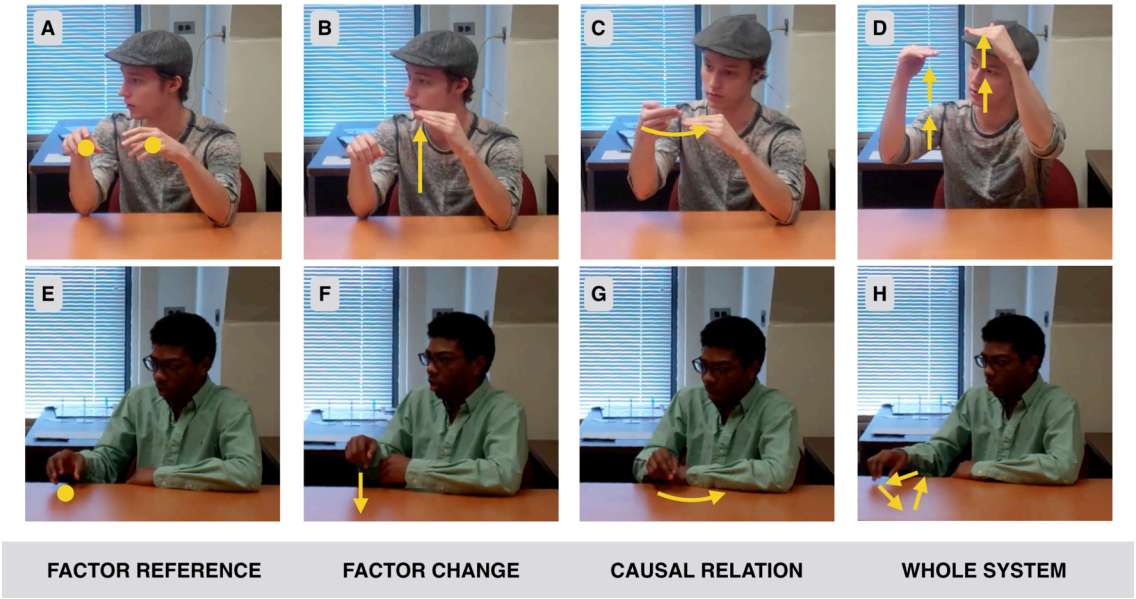

Figure 1: Examples of the different gesture types, taken from two participants' explanations. *Factor reference* gestures (A, E) represent the factors as locations in space, depicted by the yellow circles. *Factor change* gestures (B,F) represent increases and decreases as movements, depicted by the straight yellow arrows. *Causal relation* gestures (C,G) represent causation as movement, depicted by the curved arrows. *Whole system* gestures (D, H) represent the behavior of the system as a whole and often involve multiple movement phases, as depicted by the multiple arrows.

analyses reported here focus on participants' post-lesson explanations.

A first step in the analysis was to identify all gestures in the explanations that were "representational" (e.g. Chu et al., 2013). Representational gestures depict some property of a referent (commonly called "iconic" or "metaphoric" gestures), or point to a referent's location (commonly called "deictic" gestures). In the present data, the representational gestures were abstract in nature—that is, they used location, movement, and spatial arrangement to depict ideas that themselves had no concrete location, did not actually move, and had no visible spatial properties. Once representational gestures were identified, they were then categorized into gesture types (see below). Reliability was assessed by having a second coder categorize the representational gestures in 5 randomly selected explanations (26% of the data). The coders agreed 83% (N=220) of the time in whether a gesture fit into the categorization system (i.e. belonged to one of the four categories). For those gestures that both coders agreed fit into the system, they assigned the gesture to the same category in 85% (N=142) of cases. Finally, we analyzed the gestures' spatial properties and relationships to other gestures in the same explanation, as well as the language that was co-produced with them. Each of these analyses is described in more detail as the results are presented.

Results

Gesture rates and types

Participants produced a mean of 24.12 (SD=4.39) representational gestures per minute speaking. The abstract, textual nature of the Causal Systems Lesson thus did not stand in the way of eliciting representational gestures.

Based on pilot studies involving similar materials, a system was developed for categorizing the recurring ways people gesture to represent elements of feedback systems. First, people locate the factors (e.g. the predator and prey populations in the negative feedback example) by placing their gestures in space or by pointing to locations. These we call *factor reference gestures* (see Fig. 1). Second, people represent changes to the factors (e.g. an increase in the predator population) as movements. These we call *factor change gestures*. Third, people represent causal relations in the system (e.g. how the change in the predator population causes a change in the prey population) as movements, sometimes between previously established locations. These we call *causal relation gestures*. Fourth, people use movements to characterize the behavior of the system as a whole (e.g. the equilibrium that is reached in the predator-prey system). These we call *whole system gestures*. The majority (71%) of participants' representational gestures fell into one of the above four types. However, not all participants produced all four gesture types: 19 (100%) produced factor reference gestures, 18 (95%) produced

factor change gestures, 13 (68%) produced causal relation gestures, and 9 (47%) produced whole system gestures.

Spatial properties

Spatial axes We next analyzed the spatial characteristics of people's gestures, considering each gesture type separately. 96% of factor reference gestures located the factors on the left-right axis, most often with the first-mentioned factor on the left and the second-mentioned factor on the right. Factor change gestures were more variable, with 21% depicting increases and decreases as movements along the left-right axis, 29% along the front-back axis, 26% along the up-down axis, and the rest involving either some combination of these axes or a more complex movement. Causal relation gestures most commonly (75%) depicted causation as movement along the left-right axis. Whole system gestures varied considerably across participants in their use of space and in other qualitative characteristics, but they tended to use multiple movement phases (see Fig. 1, panels D and H) and often involved two hands.

Spatial consistency We next examined how consistent participants' gestures were in their spatial properties over the course of the explanation. To assess this kind of within-participant consistency, we used a measure developed in the study of spatial grammatical devices in signed languages (Senghas & Coppola, 2001). For every gesture, we asked whether it represented a system element (e.g. a particular factor) for the first time or represented a system element that had been previously represented. If the gesture repeated an element, we coded whether it used space in the same way (consistent) or a different way (inconsistent) as the immediately preceding gesture. For this analysis, we focused on factor reference and factor change gestures because causal relation and whole-system gestures, when they occurred, often only occurred once or twice in an explanation. Overall, participants were highly spatially consistent: the majority of factor reference gestures were spatially consistent (mean percentage=87%), as were the majority of factor change gestures (mean percentage=69%).

Model integration Finally, we analyzed whether participants' gestures were integrated with a larger spatial model built up over the explanation, or were more piecemeal in nature. Use of model-integrated gestures varied across participants. If, for instance, a participant produces a gesture representing an increase to a factor that incorporates the previously established location of the factor involved, the gesture would be considered *model-integrated*. As an example of a model-integrated gesture, a participant may locate the first factor on the left and then later show an increase in that factor as an upward movement in left space (see Fig. 1, panel B). If, on the other hand, the increase was depicted in neutral space or right space, the gesture would not be considered model-integrated. Similarly, if a participant produces a gesture representing a causal relation between two factors as a movement between the previously established locations of the two factors, the gesture would be considered model-integrated (see Fig. 1, panel C). Alternatively, if the gesture represented causation as a movement in neutral space, it would not be considered model-integrated. Note that factor reference gestures, which serve to establish such locations in the first place, cannot be model-integrated and are excluded from the analysis. Further, we did not expect whole system gestures, which depict a high-level summary of the whole system, to be closely integrated with the detailed causal patterns depicted in the other gestures. Overall, 50% of participants' factor change gestures were model-integrated, as were 72% of their causal relation gestures. 84% of participants produced at least one model-integrated factor change gesture, and 47% produced at least one model-integrated causal relation gesture.

Relationship to language

Finally, we analyzed the relationship between participants' gestures and the language with which they were co-produced. Most often, in 93% of cases, the gestures represented aspects of the system that were simultaneously mentioned in speech. For example, a participant would produce a factor reference gesture while referring in speech to "the first factor" or a factor change gesture while mentioning an "increase." Interestingly, however, gestures sometimes filled in where speech left off—especially when characterizing the behavior of the system as a whole. For example, a speaker describing a positive feedback system said "it's sort of…" trailing off in speech but providing a complex spatial characterization in gesture. These cases may stem from the difficulty of verbalizing the overall system dynamics.

Finally, we investigated how common it was for gestures to be co-produced with overtly *spatial* language. Table 1 provides examples of both spatial and non-spatial language that was co-produced with the different gesture types. Note

Table 1: Examples of spatial and non-spatial language co-produced with gestures

	non-spatial language	spatial language
factor reference	"first factor" "certain variable"	"external variable"
factor change	"increase in" "change"	"rise" "go up"
causal relation	"influences" "causes"	"rebounds" "turns around"
whole system	"self-correcting" "regulate each other"	"negative loop" "building on each other"

that words such as "increase" which have a general definition that is not specifically spatial, were not considered spatial. Strikingly, overall, only 17% of the gestures were co-produced with overtly spatial language. For example, factor reference gestures, though consistently exploiting the left-right axis of space, were not once co-produced with a reference to left or right.

Discussion

We investigated the possibility that people would spontaneously use spatial analogies when reasoning about positive and negative feedback, relational patterns that are complex, widespread, and fundamentally abstract. As a potential window into such hypothesized analogies, we examined the gestures people produced as they tried to articulate the main characteristics of these patterns and the differences between them. Despite the paucity of concrete spatial imagery or language in the lesson we provided, when people explained the patterns, they gestured at strikingly high rates. These gestures did not represent the actual locations or movements of objects—rather, the gestures used space abstractly to represent the different factors in the system, the changes to those factors, the causal relations between them, and the overall dynamics of the systems being described. Over the course of people's explanations, the gestures were highly consistent in where they were placed in space, and they were often integrated into larger spatial models that were built up over time. Finally, the spatial richness we observed in gesture was largely divorced from spatial language, and sometimes divorced from language altogether. In sum, the gestures we observed provided vivid evidence that people draw on spatial analogies during complex relational reasoning, evidence that would have been scarce in a verbal transcript.

One limitation of the present study is that, although the lesson was largely devoid of rich imagistic content, it did include a sprinkling of abstractly spatial words. For example, the phrase "opposite *direction*" was used to describe the change from increasing to decreasing. It remains possible that subjects took these words as cues to build larger spatial models. However, the scarcity of spatial language overall in participants' explanations makes this possibility somewhat doubtful. Nonetheless, further study will be needed to determine whether excluding such words would have a significant impact on the extent to which people create spatial models in gesture.

As we have argued, the gestures we observed revealed the spontaneous use of sophisticated spatial analogies—that is, spatial models of relational structure. Spatial analogy is likely a ubiquitous process in human reasoning. Perhaps the best-studied examples to date have involved reasoning about maps and scale models (e.g. Uttal & Wellman, 1989). In such cases, a set of concrete spatial relations in the world is mapped in schematic fashion to some spatial representation of that world. The analogical mapping is thus between one spatial format and another spatial format. By contrast, in the spatial analogies under examination here, the base concept is a purely abstract set of entities and relations—factors, changes, and causation—that is mapped to a set of spatial relations—locations, movements, and movements between locations. Prior work has demonstrated that people are able to understand and reason with spatial analogies of this abstract type (Gattis, 2004), but little work to date has examined whether the spatial analogies are spontaneously created or recruited on the fly. Informal observations, in addition to our own data, suggest that spatial analogies may constitute a powerful strategy in both cognition and communication. The ubiquity of abstract spatial models like Venn diagrams, family trees, and cladograms, for example, hints at the wider utility of spatial analogy in relational reasoning, far beyond our chosen test case of positive and negative feedback patterns (Novick, 1996; Tversky, 2011). Interestingly, the phenomenon of model-integrated gestures we have described also resembles a phenomenon in established signed languages sometimes described as "spatial modulation" (Senghas & Coppola, 2001). In American Sign Language, for example, a verb may be said to be "spatially modulated" if it incorporates spatial information that was previously established for one or more of its arguments. As our data show, hearing gesturers do something very similar under the right circumstances (see also So, Coppola, Licciardello & Goldin-Meadow, 2005).

Analogy is often thought of as an effortful process in which someone, struggling to capture a new idea, alights on an apt comparison. However, empirical work has shown that this formulation is, at least in some cases, misleading: analogical mapping can occur unintentionally, without any effort (Day & Gentner, 2007). We suggest a similar unintentional deployment of analogy may be at work here. Participants very rarely referred to their gestures—or to the spatial information contained therein—explicitly (e.g. "Imagine the system is like this"). Nor did they show signs of engaging in an effortful process of design and development, as might be signaled by restarts or amendments to the spatial structure. Rather, we suggest that participants constructed these spatial models fluidly and more or less unconsciously as they articulated the relational structure they were describing.

A related issue is whether the abstract gestures we observed were helpful to the speaker over and above any role they may have served in communication. Prior work has shown that gesturing can help speakers by reducing cognitive load (Goldin-Meadow, Nusbaum, Kelly, & Wagner, 2001). To our knowledge, though, cognitive benefits of this sort have not been shown for abstract spatial gestures of the type described here. In fact, if spatial analogy is an effortful process, as is sometimes assumed, then producing gestures like those documented here would actually increase cognitive load rather than lightening it. Testing these possibilities is a direction for future work.

Complex relational patterns underlie diverse phenomena across the natural and social worlds. While earlier work has demonstrated the difficulties of reasoning about such patterns, less is known about the kinds of representations

people bring to bear during such reasoning. Here we provide evidence for the role of spontaneous spatial analogy in this kind of reasoning, and for gesture as one means of externalizing such analogies. Though we have barely scratched the surface of this arena, it remains plausible that spatial analogies will prove to be a central ingredient in the human ability to understand complex relational phenomena.

Acknowledgments

Funding for this study was provided by the NSF Spatial Intelligence and Learning Center (SBE 0541957, Gentner and Goldin-Meadow are co-PIs) and NICHD (R01-HD47450 to Goldin-Meadow). We thank Ishara Ruffins and Erin Richard for research assistance.

References

Alibali, M. W. (2005). Gesture in Spatial Cognition: Expressing, Communicating, and Thinking About Spatial Information Information. *Spatial Cognition & Computation*, *5*(4), 307–331.

Boroditsky, L. (2001). Does language shape thought? Mandarin and English speakers' conceptions of time. *Cognitive Psychology*, *43*(1), 1–22.

Broaders, S. C., Cook, S. W., Mitchell, Z., & Goldin-Meadow, S. (2007). Making children gesture brings out implicit knowledge and leads to learning. *Journal of Experimental Psychology: General*, *136*(4), 539–50.

Casasanto, D., & Bottini, R. (2014). Spatial language and abstract concepts. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(2), 139–149.

Chu, M., Meyer, A., Foulkes, L., & Kita, S. (2013). Individual Differences in Frequency and Saliency of Speech-Accompanying Gestures: The Role of Cognitive Abilities and Empathy. *Journal of Experimental Psychology. General*.

Cienki, A. (1998). Metaphoric gestures and some of their relations to verbal metaphorical expressions. In J.-P. Koenig (Ed.), *Discourse and cognition: Bridging the gap*. Stanford, California: CSLI.

Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press.

Cooperrider, K., Núñez, R., & Sweetser, E. (2014). The conceptualization of time in gesture. In C. Müller, A. Cienki, E. Fricke, S. Ladewig, D. McNeill, & J. Bressem (Eds.), *Body-Language-Communication (vol. 2)*. New York: Mouton de Gruyter.

Day, S. B., & Gentner, D. (2007). Nonintentional analogical inference in text comprehension. *Memory & Cognition*, *35*(1), 39–49.

Forbus, K., Usher, J., Lovett, A., & Wetzel, J. (2011). CogSketch: Sketch understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science,* 1-19.

Gattis, M. (2004). Mapping relational structure in spatial reasoning. *Cognitive Science*, *28*, 589–610.

Goldin-Meadow, S., Nusbaum, H. C., Kelly, S. D., & Wagner, S. M. (2001). Explaining math: Gesturing lightens the load. *Psychological Science*, *12*, 516–522.

Goldin-Meadow, S. (2003). *Hearing gesture: How our hands help us think*. Cambridge, MA: Harvard U. Press.

Goldwater, M. B., & Gentner, D. (2015). On the acquisition of abstract knowledge: Structural alignment and explication in learning causal system categories. *Cognition*, *137*, 137–153.

Hegarty, M. (2004). Mechanical reasoning by mental simulation. *Trends in Cognitive Sciences*, *8*(6), 280–5.

Hostetter, A. B., & Alibali, M. W. (2008). Visible embodiment: Gestures as simulated action. *Psychonomic Bulletin & Review*, *15*(3), 495–514.

Jamrozik, A., & Gentner, D. (2015). Well-hidden regularities: Abstract Uses of in and on retain an aspect of their spatial meaning. *Cognitive Science*.

Novick, L. R. (2001). Spatial diagrams: Key instruments in the toolbox for thought. In D. L. Medin (Ed.), *The psychology of learning and motivation*, *Vol*. 40. San Diego, CA: Academic Press.

Nathan, M. J., & Martinez, C. V. J. (2015). Gesture as model enactment: the role of gesture in mental model construction and inference making when learning from text. *Learning: Research and Practice*, *1*(1), 4–37.

Rottman, B. M., Gentner, D., & Goldwater, M. B. (2012). Causal systems categories: Differences in novice and expert categorization of causal phenomena. *Cognitive Science*, *36*(5), 919–32.

Schwartz, D. L., & Black, J. B. (1996). Shuttling Between Depictive Models and Abstract Rules: Induction and Fallback. *Cognitive Science*, *20*(4), 457–497.

Senghas, A., & Coppola, M. (2001). Children Creating Language: How Nicaraguan Sign Language Acquired a Spatial Grammar. *Psychological Science*, *12*(4), 323–328.

So, C., Coppola, M., Licciardello, V., & Goldin-Meadow, S. (2005). The seeds of spatial grammar in the manual modality. *Cognitive Science, 29*, 1029-1043.

Traugott, E. C. (1978). On the Expression of Spatio-Temporal Relations in Language. In *Universals of Human Language*. *Volume 3: Word Structure*. Stanford, CA: Stanford.

Tversky, B. (2011). Visualizing thought. *Topics in Cognitive Science*, *3*(3), 499-535.

Uttal, D. H., & Wellman, H. M. (1989). Young children's representation of spatial information acquired from maps. *Developmental Psychology*, *25*(1), 128–138.