Analogical Transfer from Interaction with a Simulated Physical System

Samuel B. Day (day9@indiana.edu) **Robert L. Goldstone (rgoldsto@indiana.edu)**

Department of Psychology, 1101 East 10th Street
Bloomington, IN 47405 USA

Abstract

In two studies, we find that participants are able to transfer strategies learned while interacting with a simulated physical system to a dissimilar and less perceptually-concrete domain. Interestingly, performance on the transfer task was completely unrelated to explicit knowledge of the structural correspondences between the systems. We suggest that direct interaction with a concrete system may lead to a kind of procedural knowledge that provides a good basis for analogical transfer.

Introduction

There is no question that analogical reasoning is a powerful tool for learning. It allows us to look past the simple surface details of a situation, and to focus instead on underlying structure—how the components of a system fit together and relate to one another. In so doing, it allows us to make structurally-sound inferences about new situations, and provides the opportunity to draw on our wealth of existing knowledge. These processes may occur in any kind of situation. Practice problems from mathematics and physics classes are often solved by seeking out prior examples that share the same principles, even if the specific objects and situations described are concretely very different. On a larger scale, there are many stories of important scientific progress relying on apt analogies, such as Rutherford's model of the atom developing from an analogy with the structure of the solar system.

However, research has repeatedly shown that people have a very difficult time taking advantage of this tool. Unless the structural commonalities are somehow pointed out to them, people generally fail to notice that two situations from different domains are analogous. For example, Gick & Holyoak (1980, 1983) provided participants with a concrete example of a problem being solved with a *convergence* strategy, in which several small forces converged at a single location, and summed to produce a large effect. When the participants were subsequently asked to solve an analogous problem from a different domain, however, they were very unlikely to spontaneously recognize the relevance of the prior example, and therefore failed to transfer the solution strategy. The problem was not with the soundness of the analogy itself—when given a hint to think about the prior example, participants were quite good at making use of the relevant strategy. Rather, the issue seemed to be their ability to spontaneously see the connection between the episodes. This general pattern has been shown repeatedly across a wide range of materials.

What underlies this difficulty? One factor that is often cited is the concrete content of the episodes themselves. Although the terminology may vary somewhat, most research on analogical transfer has distinguished between the "deep," abstract, structural aspects of an episode, which are directly relevant for transfer, and the superficial "surface" content, which includes the concrete, domain-specific details of a particular example. For instance, in Rutherford's model of the atom, the abstract structure of multiple entities that revolve around a more massive core is relevant for analogical mapping, while details such as the color and the temperature of the sun are considered irrelevant "surface" features.

One way in which concrete information might impair analogical transfer is simply through competition with the relevant abstract structure. A consistent finding in the literature is that people are very likely to be reminded of a prior episode if it shares concrete features with a current situation, while remindings solely due to shared abstract structure are much more rare (e.g., Gentner, Rattermann & Forbus, 1993; Ross, 1984). Furthermore, even when an appropriate analog has been retrieved from memory, its application to a current problem can be impaired if the concrete features of the entities involved mismatch. For instance, Ross (1987, 1989) reported that superficial similarity between objects in two mathematical problems could reduce transfer performance if those objects played different roles in the two problems.

There is also evidence suggesting that reducing concreteness may facilitate abstract understanding and improve reminding and transfer, at least in some situations. For example, Clement, Mawby & Giles (1994) found that analogical retrieval was improved substantially when the situations were described with very abstract, domain-general terms, rather than more concrete and specific terminology. Goldstone & Sakamoto (2003) found that for participants who performed more poorly in general, the use of a less concrete training task significantly increased transfer. And one of the most robust methods for improving analogical transfer—asking participants to compare two potential analogs before solving a new problem—is presumed to succeed due to the creation of a more abstract, less concrete representation of their common structure (e.g., Gick & Holyoak, 1983; Gentner, Loewenstein & Thompson, 2003).

Together, these findings seem to suggest that concrete information represents a clear impediment to transfer between dissimilar situations, largely by overshadowing the relevant abstract structures. On the other hand, many researchers have suggested an important relationship between low-level perceptual processes and high-level, abstract representations. For example, Goldstone & Barsalou (1998) argue that most of our abstract conceptual abilities are ultimately grounded in perceptual processes (see Barsalou, 1999, for a more extreme version of this view). Others have suggested that many specific abstract concepts, such as time (Boroditsky, 2000) and mathematics (Lakoff & Nuñez, 2000), are conceptualized through analogies to more concrete knowledge.

Underlying much of this thinking is the belief that human cognition is largely designed to perceive and act on the physical world, and that our more abstract conceptual abilities are built on this foundational architecture. As Clark (1998) puts it, “Biological brains are first and foremost the control systems for biological bodies.” This suggests a very different set of assumptions about the role that concreteness might play in knowledge transfer. Specifically, it suggests that knowledge acquired via interaction with highly perceptual systems that follow the laws of the physical world—arguably the most “concrete” kind of knowledge we possess—might in fact represent an ideal base for transfer to more abstract domains.

In the current studies, we explore this possibility. In the course of interacting with a computer-simulated physical system, participants in our studies learn strategies for achieving specific goals within that system. We examine the extent to which they are then able to transfer these learned strategies to a dissimilar, highly abstract task.

A related question involves the ways in which any such transfer might differ from the kinds of analogical processes that are usually studied. Our representations of interactions with the world are generally more procedural and less explicit than the kinds of representations involved in studies of analogy, which tend to be very explicit and text-based. In Experiment 2, we examine participants’ explicit awareness of the analogical correspondences between the two tasks.

Experiments

Both experiments examined whether participants would be able to take a strategy developed through interaction with a physical system and transfer it to a more abstract, dissimilar domain. Specifically, all participants (in the experimental conditions) first interacted with a simulation involving the oscillating motion of a ball that was suspended between two elastic bands. Although the general operations of this system were consistent with participants’ naïve physical theories, the actual strategies necessary to elicit a desired behavior from the system were often less than intuitive, and generally required a fair amount of trial and error. Next, all individuals participated in an ostensibly unrelated task, which involved regulating the population of city. While this second simulation differed considerably from the first, both in terms of its content and its visual display, the system was governed by the same underlying principles as the first task.

In both tasks, participants were asked to accomplish a specific goal, which required the development of an appropriate strategy. Our primary manipulation was in the relationship between the goals for the two tasks. For some participants, the two goals were analogous, and thus required analogous strategies to achieve. For other participants, the two goals were structurally dissimilar, with each requiring a unique strategy. If participants are in fact able to transfer their learning from the interaction with the physical system to the dissimilar and less perceptually-concrete population task, we should find facilitation for those participants with consistent, analogous goals.

Experiment 1

Participants. 63 Indiana University undergraduates participated in this study for partial course credit.

Materials and Design. In the first task, participants interacted with a simulation of a physical system, a ball that was suspended between two elastic bands that were attached on either side. The bands stretched horizontally in either direction, and each was attached to a stationary pin (see Figure 1). The motion of the ball was fairly realistic, and was easy to grasp intuitively, with greater distance from either pin leading to more “stretching” of the band and greater force pulling toward that pin. Computationally, the ball moved according to some simple physical rules—its natural tendency was to continue along a constant vector, which could be altered by accelerating forces from the two bands. This acceleration increased as a linear function of length of the bands, reflecting a stronger force from each band as it was stretched farther. (In order to increase the realism of the display, the width of each elastic band decreased as its length increased, simulating its tension.). Thus, motion away from one of the stationary pins would increase this distance, causing increasing acceleration toward that pin. This would result in the ball slowing down and eventually moving back toward the pin. However, the existence of two opposing forces meant that the ball’s position would tend to oscillate: movement toward either pin would also tend to be movement *away* from the opposite pin, thus increasing acceleration in that opposing direction. For simplicity, neither gravity nor friction was included in the model, resulting in perpetual motion.

Participants were first asked to take a few minutes to explore and familiarize themselves with the behavior of the system. Participants were able to click on the ball and drag it to any position within the display (both horizontally and vertically), and then observe the motion that resulted when the ball was released. Participants proceeded at their own pace, and were allowed to interact with the simulation for as long as they wished.

The experiment then proceeded to the training phase of the task. In this phase, participants were no longer able to drag the ball to a new location. However, they were given a new way to manipulate its behavior: a fan that blew directly rightward across the ball, introducing an additional force in that direction. Since all potential forces acting on the ball in this phase were horizontal (the ball’s starting position was between the two pins), the ball’s path in this phase was constrained along a single horizontal line. The fact that the tension from the two bands was symmetrical meant that all motion was essentially an oscillation around the midpoint between the two pins.

Each participant was asked to accomplish one of two goals: either to cause the ball to reach the pin on the far right, or to cause the ball to “stabilize” directly between the two pins. In other words, each participant needed to consistently increase or decrease the amplitude of the ball’s oscillations through the appropriate use of the fan. In the described system, when the ball is traveling to the right, the force from the fan will add to its net velocity, causing it to travel slightly farther in that direction (and thus increasing its amplitude). When the ball is traveling to the left,

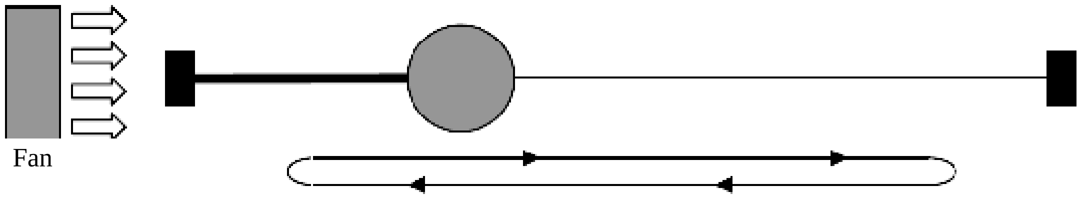

Figure 1: Schematic of training task. The task simulated a ball suspended between two elastic bands. During familiarization, participants could move the ball both vertically and horizontally. During training, however, all forces on the ball were horizontal, leading to oscillating horizontal motion. Participants attempted either to move the ball to the far right pin, or to stabilize the ball in the middle, through appropriate use of force from the fan.

however, the force from the fan will oppose its velocity, slowing the ball's movement in that direction and thus decreasing its amplitude. Constant application of the fan would thus lead to no net change in the ball's amplitude, and would not help in achieving either goal. The optimal strategies are therefore to activate the fan only during the *rightward* part of its oscillation to increase its amplitude and reach the far pin, or only activate it during the *leftward* part of its oscillation to decrease its amplitude and allow it to stabilize in the middle (participants were not informed of these strategies). Participants were required to complete this task seven times. Upon completion, they were told to ask the experimenter to start the next, ostensibly unrelated part of the session. (A Flash implementation of the basic ball task is available online at: cognitrn.psych.indiana.edu/complexsims/Oscillatingball.html).

Participants were told that the next task involved a computer simulation of how a city's population could vary over time, and how it could be influenced by media advertisement. According to the instructions, the city in question was large enough to comfortably hold 500,000 residents. If there were fewer people than this, the city would become more attractive to outsiders because of abundant housing and low traffic congestion. With more than half a million residents, the city would become less attractive because of crowding, crime, and expense. Thus, the city's appeal would increase whenever the population was below this optimum level, and would decrease when the population was above this ideal. Furthermore, the amount of the change in appeal would be greater as the distance from 500,000 increased. Unlike the previous task, the interface for the population simulation was entirely textual, and proceeded in discrete time steps rather real (continuous) time. At each time step, participants were given numerical values for the city's population, its current appeal (which could be positive or negative), and the change in its appeal from the previous time step (which could also be positive or negative). This information was presented in a scrolling text display, which also allowed the information from the previous four steps to remain visible on the screen.

Participants in this task were also given one of two goals to accomplish: either to cause the population to reach 1,000,000, or to cause the population to stabilize around its optimal value of 500,000. To achieve these goals, participants decided whether or not to introduce media advertisement for the city at each time step, which would temporarily increase its appeal. Participants were required to complete this task three times.

Although the content of the first and second simulations was quite different, the principles governing their operation were essentially identical. The city's population is analogous to the position of the ball at a given point in time. The city's "appeal" therefore maps onto the velocity of the ball, representing the numerical change in population from one time step to the next. Similarly, "change in appeal" is analogous to acceleration, describing the degree to which the *change* in population is increasing or decreasing. Finally, media investment plays the same role as the fan, allowing the participant to add a unidirectional force at any point in time. The behavior of the system is therefore qualitatively the same as that of the ball simulation. The population tends to oscillate around the "midpoint" of 500,000, and adding media advertisement when the population is rising or falling will respectively increase or decrease the amplitude of this oscillation.

The dependent variable in this study was the number of time steps required to complete the population task (averaged across the three attempts). Our primary interest was in whether solution strategies acquired in the first, concretely physical simulation would transfer to the highly abstract (and dissimilar) population task, resulting in shorter average solution times. Specifically, we predicted that solution times in the population task would be shorter when the goals of the two tasks were mutually *consistent* (i.e., Fan both had the goal of maximizing the amplitude, or both had the goal of stabilizing the amplitude) rather than *inconsistent* (e.g., maximizing the ball's location and stabilizing the population, or vice versa). Thus, the study had a 2 (goal consistency) × 2 (population task type) factorial design.

Results and discussion. Participants were required to complete the two simulations within a one-hour experimental session. While most participants completed the ball task within 10 minutes, many found the population simulation quite challenging and failed to finish within the allotted time. Since ability to complete the simulation is obviously a good indicator of how difficult each participant found the task, and since it is directly related to our dependent variable of solution time, these participants were included in the analyses. Each was conservatively given a score of 2000 time steps for their unfinished attempts (slightly less than the longest time that any participant took to complete the task, 2174 steps). 24 participants were assigned this maximum score.

A 2 × 2 (goal consistency × test type) factorial ANOVA revealed a reliable difference between conditions ($F(1, 59) = 10.89$, $p < .001$). A main effect of test type ($F(1, 59) = 28.78$, $p < .001$) reflected the fact that participants took considerably longer to complete the population stabilization task than the population maximizing task (averages of 1324 and 318 time steps, respectively). More relevant to our current interests was the main effect of goal consistency ($F(1, 59) = 4.14$, $p < .05$). Participants required reliably fewer trials to complete the population task when it required achieving a goal that was analogous to that of the training task, and thus required an analogous solution strategy (632 vs. 999 trials). Given that so many participants failed to complete the transfer task, it also makes sense to analyze the results in terms of simple completion rates. Overall, significantly more participants were able to finish the population task when its goal was consistent with the prior task than when it was inconsistent (75% vs. 48%; $\chi^2(1, N = 63) = 4.73$, $p < .05$).

Experiment 2

Experiment 1 provides evidence that representations of perception-action schemes may provide a good source for analogical transfer to more abstract domains. Given the concreteness of the base domain and the dissimilarity between the two simulations, this finding runs somewhat counter to common wisdom about analogical transfer. We propose that the interactive and physical (although simulated) nature of the training task differentiates the current studies from much of the previous literature that has found such poor cross-domain transfer.

This raises some interesting questions about the nature of the transfer that we are observing, and how it might differ from that found in most studies of analogical processing. In fact, one such possible difference is implicit in the design of the study itself. We hypothesized that participants might transfer an abstract solution *strategy*—a specific method of *interacting* with a system—from a very concrete domain. This is consistent with our intuitions about the kind of information that is being acquired. However, it might have been reasonable to predict alternatively that participants would map their knowledge of the entire set of rules governing the system, not just a particular strategy, since the two systems were essentially isomorphic. This should have led to facilitation in all conditions, and thus little or no difference between the groups. Of course, it is possible that *both* of these kinds of transfer are operating. The second experiment uses a control condition to examine this possibility.

Another intriguing potential difference between our effects and those found in more traditional kinds of transfer studies is the role of explicit knowledge of the correspondences between the tasks. There is general agreement in the literature that analogical reasoning begins with reminding of a prior analogous case, followed by a mapping process, in which correspondences are established between the components of the two representations (e.g., Forbus, Gentner & Law, 1995; Hummel & Holyoak, 1997). Although the issue of explicit awareness is not generally discussed, it seems to be assumed that the processes involved are largely explicit, with individuals having direct access to the output of their analogical processing (though see Day & Gentner, 2007a). However, the information that is transferred in the current studies seems in some ways more related to procedural knowledge (Squire, 1987) than the semantic and episodic knowledge that are the basis of most analogy research. As such, it is possible that it might be similarly resistant to explicit examination. In the current experiment, we examine this possibility by asking participants about their awareness of the relationship between the tasks, both through open-ended questions and through a correspondence-matching task.

Participants. 91 Indiana University undergraduates participated in this study for partial course credit.

Materials and Design. The overall structure of this experiment was very similar to that of Experiment 1, with a few important differences. First, we included a control condition to explore the possibility that exposure to an analogous system was benefiting all participants, regardless of the consistency of the goals between the two tasks. The control task required participants to guide a spacecraft to its home planet through the appropriate placement of an "attractor" in space. The attractor exerted an attractive force on the spacecraft which was inversely proportional to the distance between them; thus, as the craft moved closer to the attractor, its acceleration toward the attractor increased. The spacecraft in question would otherwise follow a set trajectory—participants had no method of moving the craft other than the placement of the attractor. The task therefore required participants to find the particular placement of the attractor that would cause the ship's trajectory to curve until it reached the destination planet. Thus, the control task contained the fundamental mechanical elements of the experimental conditions—motion along a constant vector that is altered by accelerating forces from a stationary point—but the overall structure of the system was quite different. If it is only the procedural strategy that is being transferred, control group performance should be roughly equivalent to the inconsistent goal condition. If, on the other hand, participants are benefiting from their exposure to the analogous structure in general, then we should find that those in the control condition perform more poorly than both of the experimental conditions. The experiment therefore had a 3 (task consistency: analogous with same goal, analogous with different goal, non-analogous control) × 2 (population task type) factorial design.

We also took steps to simplify the population task somewhat. As reported, a significant number of participants in the first experiment failed to complete the task within the one-hour session. While this is interesting from a theoretical perspective, and provides useful data about the relative difficulty of the task for the different experimental groups, it also introduces a fixed ceiling for solution times, and may therefore be obscuring some interesting variance in the data. We therefore added a simple visual display—a line graph, with population on the y-axis and time steps on the x-axis—to both versions of the population task in order to facilitate the tracking of population changes over time.

After the population task, we took two measures to determine participants' explicit understanding of the relationship between the tasks. (These tasks were not administered to the control group, since their tasks were not analogous). First, we asked a series of open-ended questions to assess awareness that there was any relationship between the ball and population tasks. Participants were first asked for their general response to the tasks, and whether they had noticed anything interesting or unusual about them. Next they were asked how similar they found the two tasks, and to describe any similarities they had noticed. Finally, if participants reported that they had found the tasks similar, they were asked to report when in the session they had first noticed this similarity. Next, participants completed a matching task, in which they selected which component from the population task corresponded with a particular component from the ball task. For example, "fan" corresponded to "media investment." The ball task components were presented one at a time, to minimize responses based on a "process of elimination" across the entire set. Six correspondences were matched in total.

Results and discussion. In spite of attempts to simplify the population task, several participants still failed to finish within the allotted time (29 out of 91). As in the first experiment, these participants were given solution times of 2000 time steps.

A 3 × 2 (goal consistency × test type) factorial ANOVA revealed a reliable difference between conditions ($F(5, 85) = 9.93$, $p < .001$). A main effect of test type ($F(1, 85) = 36.94$, $p < .001$) again reflected the fact that participants took reliably longer to complete the population stabilization task than the population maximizing task. More importantly, we again found a main effect of the relationship between the task goals ($F(2, 85) = 3.57$, $p < .05$). Post hoc analyses using Tukey's HSD procedure revealed reliable differences between the consistent- and inconsistent-goal groups (447 vs. 870 trials; $p < .05$) and between the consistent-goal and the control group (447 vs. 932 trials; $p < .05$), but found no difference between the control and the inconsistent-goal conditions ($p = .97$). A similar pattern emerges from analysis of completion rates: 79% of participants in the consistent-goal group successfully completed the population task, compared with 57% and 58% in the inconsistent and control groups, and the difference between the two analogy groups is significant ($\chi^2(1, N = 58) = 4.38$, $p < .05$).

Experiment 2 therefore replicates the basic finding of Experiment 1, with participants completing the transfer task significantly faster if it required a strategy that was analogous to that of the training task. Additionally, these results show that there was no facilitation for simply interacting with and learning the rules of an analogous system, since the inconsistent-goal group performed no better than the control condition. Rather, the transfer seemed to be in the form of knowledge of particular strategies for interacting with the system, perhaps a type of procedural knowledge. However, that knowledge was clearly in a sufficiently abstract form to allow transfer to a very dissimilar task and domain.

Next, we analyzed participants' responses to the open-ended questions and the correspondence-matching task. Open-ended responses were coded and assessed for whether participants reported noticing any relevant structural commonalities during the course of the tasks. Similarities that were not relevant to the analogous structure were not counted (e.g., "Both involved clicking the mouse button"), nor were similarities that participants reported noticing after completion of the population task. The correspondence-matching tasks produced scores between 0 and 6, reflecting the number of correct matches.

In spite of the added visual display, which made the oscillating movement of the population quite salient, only about 1/3 of the participants in the experimental groups (17 out of 58) reported noticing any structural commonalities between the tasks (based on the open-ended questions). This is consistent with prior findings of poor explicit reminding between dissimilar analogous situations, even when they are in close temporal proximity. Not surprisingly, those participants who noticed that the two tasks were analogous performed better overall on the transfer task ($t = 2.26$, $p < .05$). However, when looking at those who did versus did not recognize the analogy separately, we found similar advantages for consistent over inconsistent-goal conditions, though neither recognition group was statistically reliable on its own: 164 vs. 529 time steps ($t(15) = 1.67$, $p = .11$) for those who recognized some commonalities between tasks, 602 vs. 1007 time steps ($t(39) = 1.53$, $p = .13$) for those who did not (there was no interaction between condition and recognition of commonalities).

Similarly, there was absolutely no correlation between accuracy in correspondences and performance on the transfer task for either condition ($R^2 < .01$, $p = .99$), nor was transfer related to any particular correspondence item. Remarkably, even the most seemingly fundamental correspondences were uncorrelated with transfer performance (by point-biserial analysis), including the mappings between "fan" and "media investment" ($R^2 < .01$, $p > .7$) and between "ball location" and "population" ($R^2 < .01$, $p > .9$). This lack of correlation was not the result of a restricted range in the matching task scores, since average accuracy on those items was roughly 50% (.42). Thus, although recognizing the *existence* of a deeper relationship between the tasks was related to better overall performance, this recognition did not appear to affect levels of analogical transfer (in terms of consistent vs. inconsistent conditions), nor did explicit recognition of the specific structural correspondences between the two systems.

General Discussion

In these two studies, we find evidence for analogical transfer from a simulated concrete physical system. We believe that participants' active participation in the oscillating ball task led to a fairly "visceral" understanding of how to interact with and accomplish certain goals within that system. This procedural knowledge was then able to act as a base for analogical transfer to an abstract and dissimilar domain. Interestingly, although awareness of a deeper connection between the two tasks was associated with improved performance on the transfer task, explicit knowledge of the actual correspondences between the domains was unrelated to task performance. These findings are consistent with suggestions that low-level perceptual and motor processes, as the foundation of the cognitive system, are able to serve as a basis for much more abstract conceptual understanding.

This research also raises some interesting issues about the most preferable conditions for transfer. As discussed, there is considerable research suggesting that concreteness may sometimes present an impediment to analogical reasoning. This suggests that transfer should operate best from a fairly abstract, amodal mental representation (even if this representation was initially acquired through generalization across multiple concrete instances, as through comparison). On the other hand, proponents of a more situated and embodied approach to cognition would predict that representations that are acquired through direct interaction with systems that follow natural physical laws—arguably the most "concrete" representations we possess—often provide the best means of understanding more abstract concepts. Although further research will clearly be needed to directly compare these alternatives, the current findings are at least consistent with this latter approach.

The fact that performance on the transfer task was unrelated to explicit knowledge about the correspondences between the two systems was particularly interesting. Analogical reasoning is generally thought of as involving fairly slow, deliberate, and intentional processes. In finding that structural knowledge may be acquired and applied without an explicit understanding of that application, we are demonstrating a different kind of transfer than is generally considered—and perhaps a kind that is even more pervasive in daily life (also see Day & Gentner, 2007a&b).

The goal of understanding and facilitating knowledge transfer is clearly an important one, both for psychologists and educators. We believe that the current studies provide the beginnings of an interesting way of considering the role of concrete, perceptual, motoric knowledge in the understanding of more abstract domains.

References

Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22,* 577–609.

Boroditsky, L. (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition, 75,* 1–28.

Clark, A. (1998). Embodied, situated, and distributed cognition. In W. Bechtel, & G. Graham (Eds.), *A companion to cognitive sciences* (pp. 506–517). Malden, MA: Blackwell Publishers Inc.

Clement, C. A. , Mawby, R. & Giles, D. (1994). The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language, 33,* 396–420.

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

Day, S. B. & Gentner, D. (2007b). Hidden structure: Indirect measurement of relational representation. *Proceedings of the Twenty-ninth Annual Meeting of the Cognitive Science Society.*

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19,* 141–205.

Gentner, D., Loewenstein, J., Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology, 95,* 393–408.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology, 25,* 524–575.

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

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15,* 1–38.

Goldstone, R. L., & Barsalou, L. (1998). Reuniting perception and conception. *Cognition, 65,* 231–262.

Goldstone, R. L., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology, 46,* 414–466.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427–466.

Lakoff, G., & Núñez, R. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being.* New York: Basic Books.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology, 16,* 371–416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 13,* 629–639.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory and Cognition, 15,* 456–468.