Proceedings of the International Conference on Model-Based Reasoning: Scientific Discovery, Technological Innovation, Values, held May 17–19, 2001, in Pavia, Italy

ISBN: 0-306-47244-9

© 2002 Kluwer Academic/Plenum Publishers, New York
233 Spring Street, New York, New York 10013

http://www.wkap.ml/

10 9 8 7 6 5 4 3 2 1

A C.I.P. record for this book is available from the Library of Congress.

All rights reserved

No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording, or otherwise, without written permission from the Publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work.

Printed in the United States of America

Model-Based Reasoning

Science, Technology, Values

Edited by

Lorenzo Magnani
University of Pavia
Pavia, Italy

and

Nancy J. Nersessian
Georgia Institute of Technology
Atlanta, Georgia

Kluwer Academic / Plenum Publishers
New York, Boston, Dordrecht, London, Moscow

Perceptual Simulation in Analogical Problem Solving

David L. Craig*, Nancy J. Nersessian#, and Richard Catrambone†
*College of Architecture, david.craig@arch.gatech.edu; #College of Computing, nancyn@cc.gatech.edu; †School of Psychology, GA, USA, rc7@prism.gatech.edu, Georgia Institute of Technology, Atlanta

Abstract: We argue that in some situations solutions to target problems can be generated based on perceptual simulations of analogous but superficially dissimilar source problems. Unlike processes associated with perceptually neutral representations, those associated with perceptual simulations should be highly sensitive to context, whether conveyed through textual, visual or physical mediums. Consistent with this claim, we found that when the context of a source problem was manipulated to accentuate kinesthetic structure, people were more likely to generate solutions to an analogous problem. We also found that subjects who drew diagrams of solutions to source problems were most likely to transform them into solutions to an analogous problem when the contextual properties of their diagrams afforded such transformations.

1. INTRODUCTION

Imagine an engineer working on a cable system with an unusual set of constraints. The task requires that the engineer add a couple of tensioning wires to the system at right angles to each other. Although attaching the cables to the system is trivial, the problem is that the parts they attach to move back and forth, and when they do the two cables run into each other, causing the whole system to shut down. Ultimately, if the cables are positioned where they need to be, they will somehow have to be able to *pass through* one an-

Model-Based Reasoning: Science, Technology, Values, edited by
L. Magnani and N.J. Nersessian, Kluwer Academic/Plenum Publishers, New York, 2002

other without either of them disconnecting. Initially the problem seems impossible. The design engineer scratches her head and thinks if there might be a way to reposition or eliminate the tensioning cables altogether. She then thinks more carefully about the business of one thing passing through another without sacrificing its physical integrity. Suddenly she remembers the revolving door to her building which she passes through everyday day on the way to her lab. The revolving door is designed to let people pass through effortlessly while nonetheless keeping the inside of the building sealed off from the outside at all times. One part of the door, that is, always closes just as another opens and vice versa. With this insight she quickly sketches out a gate to add to one of the cables (Figure 1), which, like the revolving door, will allow the other cable to pass through it without either disconnecting, using the principle of redundancy (French, 1994).

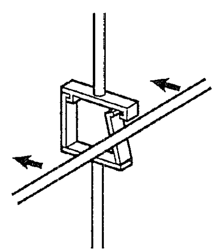

Figure 1. A gate mechanism, employing redundant latches, that allows one cable to pass through another at right angles.

In the scenario described above, the revolving-door exists in a context far removed from the cable problem. Yet the way the cable problem is analyzed somehow triggers the retrieval of the revolving door and the general principle behind it. The example illustrates the nontrivial process of solving a problem by constructing an analogy. Although the engineer may not have possessed explicit knowledge about how to design intersecting cable systems, the problem is recognized as being similar to one encountered elsewhere and one for which there is a known solution.

Because of their remarkable ability to transcend the limits of domain-specific knowledge, analogies play a central role in design, science, and other creative endeavours, as well as in everyday learning and problem solving. In this chapter we look at how mental representations are formed and transformed, and how differences in potential source experiences - the kinds of experiences that can lead to analogies - may result in differences in the construction process. As we look at the analogical problem solving process our primary focus is on perceptual factors. Going back to the example described above, one might ask if an engineer did in fact use her knowledge of a revolving door to solve a cable problem, would it have mattered if she imagined a particular revolving door (e.g., the one she uses every day), and if so, would the particular features of that door have mattered? Would it have mattered, for example, whether the door was big and heavy, whether it operated automatically or manually, or whether it was surrounded by a solid wall or large glass panels? Would it have mattered if the engineer had imagined walking through the door or if she instead visualized a schematic drawing of the sort produced by an architect? In most existing research on analogical problem solving these sorts of features are, strictly speaking, taken to be superficial, in that they bear no obvious relationship to the underlying principles used to solve the problem. As such, they are commonly disregarded as features that get stripped away before the real thinking gets underway.

As the chapter unfolds, we articulate a theory that explains how perceptual factors such as those described above could come into play in analogical problem solving. In particular, we present a theory that explains how analogies might be constructed via perceptual simulations. A perceptual simulation is a type of representation that likely involves the reactivation of patterns of neural activity in the perceptual and motor parts of the brain initially activated in experiencing something first hand. One possible component of perceptual simulations that happens to be fairly well understood is visual imagery. Perceptual simulations, however, can also involve other sensory modalities, including those associated with experiences of force and space. They can be richly or sparsely detailed, highly structured or flat. They can also be used to represent abstract concepts. Despite their generality, perceptual simulations contrast with the type of representation standardly used to explain high-level thought, namely predicate-based representations. Given the differences between the two, one might expect an explanation of analogical problem solving based on perceptual simulations to result in predictions somewhat different from those made by researchers employing the standard view. In particular, because perceptual simulations demand a heightened degree of physical realism, the physical setting one imagines might affect how one represents principles underlying a physical system and ultimately how one makes use of such principles in constructing analogies, even in cases where the setting is superficially distinct from that of the target problem.

Our motivation for trying to explain how analogies might be constructed via perceptual simulations stems from an interest in better understanding how external representations - including stories, pictures and even bodily reenactments - influence problem solving. In most research on analogy, external representations have been treated primarily on the assumption that they *contain* information. Reasoning with external representations, moreover, often seems limited to extracting information, in that once the information is extracted it is converted into a perceptually neutral format and at that point transformed independently of perceptual processes. To use an external representation like a diagram, for example, is to see pertinent information and discard the perceptual substrate. Alternatively, external representations might be seen as setting up perceptual simulations. The perceptual context implied in an external representation might, when integrated into a simulation, have a significant effect on how such a simulation can be transformed, particularly if the transformational processes are also perceptual.

The purpose of the chapter is not to prove that analogies are always constructed via perceptual simulations, nor to rule out all alternative explanations. We do, however, intend to show how findings produced in our lab might be explained more efficiently in terms of perceptual simulations than in terms of perceptually neutral representations. We start by explaining in more detail what perceptual simulations are. We then present a brief overview of analogical reasoning processes and describe how they might be accomplished via perceptual simulations. Finally, we describe a series of experiments designed to highlight the role of perceptual factors in analogical problem solving using perceptual simulations, both static and dynamic. In each experiment, a different type of external representation is used. In the first experiment perceptual factors are manipulated via narratives; in the second they are manipulated via physical activity; and finally, in the third they are manipulated via diagrams. We conclude the chapter by briefly discussing how our findings might impact more wide ranging studies of creative problem solving.

2. PERCEPTUAL SIMULATION

Representations have two significant aspects: format and content. In most research on analogy and other types of conceptual reasoning, the two have been considered independently. That is, researchers have typically employed representational forms that bear no relationship to the content they are used to represent. Representations encoded in such formats are amodal representations. Natural language is a good example of an amodal representational system in that the sound or appearance of a word typically bears no relationship to its referent - the thing it describes. Amodal representations can be used to represent nearly anything, precisely because they are analogous to natural language. They could, for instance, take the form of predicates, where a predicate is simply a statement about some thing or about a relationship between things of the type ROUND(ball) or PASS_THROUGH (ball, hoop).

A notable property of amodal representations is that they are perceptually neutral - not in the sense that they fail to represent perceptual information (ROUND(ball) is obviously perceptual), but rather in the sense that they do not preserve information about the modalities through which their referents

are clearly experienced. Predicates like DIAMETER(ball, 30cm) and DIAMETER(hoop, 31cm) are perfectly useful for reasoning about whether, for instance, a ball will fit through a hoop, but they do not suggest what it is like to hold something 30cm in diameter or even how the experience of size differs from, say, the experience of weight. Some have argued that this property makes amodal representations attractive. Amodal representations, in other words, allow one to reason with information drawn from different modalities all at the same time (Palmer, 1975). The real appeal, though, of amodal representations is that they can efficiently represent things that are not perceptual. Predicates, for example, are capable of representing significant differences between things that seem otherwise perceptually similar (Smith and Medin, 1981). Oranges and orange plastic balls may look the same (and feel the same and so forth), but predicate-based representations of the two can get right to the heart of their differences, representing, for example, that one is nutritious while the other is inert, or that one is organic while the other is made in a factory.

Modal representations contrast with amodal representations in that modal representations do preserve information about the modalities through which their referents might be experienced (Barsalou, 1999; Glenberg, 1997). A modal representation is a perceptual state or a long-term copy of one. It is an analog representation, meaning that it preserves constraints inherent in the patterns of neural activation found in the perceptual and motor cortices of the brain. Reactivating a modal representation in working memory means perceptually simulating something, or, literally, reactivating a particular perceptual state. Like predicate-based representations, modal representations need not be complete. One would not, in other words, be limited to remembering things exactly as they were experienced. Modal representations could be schematic in the sense that bits and pieces of perceptual experiences could be stored separately and reactivated in novel combinations. A long-term modal representation of an orange, for example, may preserve certain properties, like shape, size, and taste, independently, allowing one to later imagine things like orange-flavored candy or orange-colored cars. That modal representations could function in such a way is supported by the fact that the perceptual system itself tends to break up stimuli (e.g., patterns on the retina) into distinct perceptual features and relations.

There are a number of reasons to believe that people do store modal representations and do make use of them, via perceptual simulations, in various reasoning tasks. One widely researched type of modal representation is mental imagery. As a number of studies have demonstrated (see Kosslyn, 1994, for a review), people who are compelled to "picture" something in their head exhibit behavior that overwhelmingly suggests they are processing short-term representations in the primary visual cortex. When, for example, people are asked to answer questions about the relative locations of points shown on a memorized spatial map, their responses tend to lag depending on the distance between the points, as if they were actually visualizing and scanning between them. If, by contrast, such a spatial map were represented using predicates, all distances would be encoded in the same way, independent of actual length, and one would expect processing speed to be independent of actual distance. There is also evidence that people use modal representations in basic conceptual tasks. One finding, for example, which Barsalou (1993) calls "linguistic vagary," is that when people are asked to list features associated with everyday objects, they provide significantly different answers at different times, even if the difference in time is just a matter of days. As Barsalou argues, such results are hard to explain using amodal representations. Predicates, for example, should produce fairly consistent output because predicates are analogous to natural language. Listing features based on modal representations, on the other hand, would be open to the vagaries of perception (physical perspective, situational context, attention, etc.), resulting in the type of behavior that is, in fact, found. None of these findings, of course, refute amodal representations outright. One could always jerry-rig an amodal representation to function like a modal representation and vice versa. Nonetheless, some types of behavior do seem more consistent with modal representations. More of this type of evidence will be presented as the chapter progresses.

The most significant limitation of modal representations is in encoding abstract conceptual information. Modal representations encode only sensory experiences and, hence, may seem inherently superficial. How, then, would one explain known differences between an orange and an orange plastic ball with modal representations? Although some perceptual differences might tip one off (oranges are tangy, for example), the most substantial differences are not superficial, such as the fact that oranges are organic and plastic balls are synthetic. Anticipating this problem, Barsalou (1999) argues that abstract concepts like organic and synthetic could be modally represented in terms of complex sequences of events in which particular objects and actions are highlighted. Organic, for example, might require a modal representation of a sequence of events in which an orange grows on a tree, while synthetic might require a representation of something being manufactured in a factory. Problematically, though, concepts like organic and synthetic are not specific to plants and factories. Rather, they seem generic and capable of entering into reasoning independent of such situations.

To represent generic, abstract concepts, we propose that people might recruit modally represented, perceptual structures to *define* them. The kinds of perceptual structures employed would be spatial and kinesthetic - in other words, structures abstracted from first-hand, bodily experiences of space and force. Consider again the concepts organic and synthetic. Organic might be represented using spatial and kinesthetic structures associated with expansion from within, something most people should be familiar with from a physical

point of view. Synthetic, on the other hand, might be represented using spatial and kinesthetic structures associated with accretion. Obviously, biologists entertain more complicated representations of organic and synthetic, but these might just consist of a more complicated network of spatial and kinesthetic structures. Following Lakoff and Johnson (Johnson, 1987; Lakoff, 1987; Lakoff and Johnson, 1998) and Mandler (1992), we refer to spatial and kinesthetic structures as image schemas. Image schemas serve primarily to structure mental images, each possessing structure that can be analyzed but that nonetheless is elemental in form. An image schema for convergent force, for example, serves to structure that which is compressed in relation to that which is doing the compressing. By incorporating a convergent-force image schema in a perceptual simulation (e.g., by overlaying it on an image), one could represent things ranging from a person being compressed by a crowd to a walnut being compressed by a nutcracker. Importantly, though, image schemas need not always coexist with other modalities, visual or otherwise. An image schema for compression, for example, could be activated by itself simply as an independent representation of space and force. Although experiences of space and force typically coexist with other sense modalities, they too can be abstracted via perceptual processes and reactivated separately, in the same way that one might imagine a color or a shape as independent of an object.

To review so far, both modal and amodal representations, as theoretical entities, would seem capable of representing knowledge in schematic, composable units, ranging from minute perceptual details to abstract concepts. These units of information, however they are formatted, are, in turn, likely bound in rich associative networks in long-term memory known as schemas, scripts, frames, or mental models (Barsalou, 1992; Gentner and Stevens, 1983; Minsky, 1975; Palmer, 1975; Rumelhart and Ortony, 1977; Schank and Abelson, 1977). Despite the similarities in the types of information that modal and amodal representations might encode and the way that information might be associated in long-term memory, one significant difference between modal and amodal representations that could impact higher-level reasoning, particularly reasoning associated with analogies, is in how structural relations are encoded, specifically in how information is *encapsulated* in them. As mentioned before, in modal representations, structural relations must be represented in image schemas. A major difference between an image schema and a predicate-based representation of the same thing is that the former is considerably more condensed. The kind of information implicit in an image schema, in other words, would take a lot of predicates to represent explicitly. That image schemas automatically encapsulate a wealth of information may prove important in things like constructing analogies. An increase in representational power, that is, might increase the potential overlap between two domains and thereby increase the chance that similarities between them will be detected.

A second important difference between modal and amodal representations that we will consider has to do with how information is processed in working memory. In general, to complete any reasoning task requires that some subset of information stored in long-term memory be activated. The simplest reasoning tasks might be accomplished deductively, by tracing links in an activated part of the associative network. To tell what color ravens are, for example, one might activate raven and follow the link to color. More complex tasks, however, might require that one create new representations in working memory by restructuring units of information from long-term memory. To reason about a raven in a raven's nest, for example, might involve using the structural relation "in" to bind existing representations of raven and nest, regardless of whether such a thing has ever been witnessed or imagined in the past. The difference, in this case, between modal and amodal representations is that in the former all information included in working memory must be perceptually coherent. In addition, if one transforms a modal representation in working memory, those transformations must also be perceptually coherent. Because of the demands of perceptual coherence, context might ultimately affect how a representation is constructed and transformed over time.

3. ANALOGICAL PROBLEM SOLVING

Analogical problem solving is heavily dependent on mental representations. To start, a representation of a problem in working memory - a target representation - must somehow activate another representation - a source representation - stored in long-term memory. This process, known as retrieval, could be triggered by the fact that both representations share common features, parts, events, scenes or abstract concepts. However, because any representation active in working memory is likely to overlap with a large number of representations in long-term memory, additional matching criteria are typically required. One possibility, suggested by Gentner and her colleagues (e.g., Clement and Gentner, 1991; Gentner 1983), is that the retrieval process selects only those representations that yield highly systematic matches. Systematicity, in this case, refers to the degree of consistency with which high-level relations match between representations, high-level relations being those that structure other relations as opposed to features, parts or objects. That one object revolves around another, for example, is a low-level relation, because it structures two objects. That gravity *causes* one object to revolve around another, however, is a high-level relation, since the causal relation connects two other abstract relations, namely gravitational pull and revolve.

Once a source representation is retrieved (on the promise that it can help solve the problem at hand), that representation must be aligned with the target

representation. In some cases, the retrieval process itself may produce partial alignment, as is required, for example, in checks on systematicity. In any case, the alignment process brings into correspondence parts of the source and target representations that match exactly, which in turn allows one to bring things that do not match into correspondence. Cigarettes and time bombs, for example, can be brought into correspondence despite their many differences, because both are known to cause eventual destruction (Gentner, 1983). Ideally, alignment is followed by a mapping process in which structural relations that exist in the source representation but not the target, and which are systematically connected to already aligned relations, are mapped onto the target. The mapping process, of course, is what in ideal situations produces a solution to the target problem.

In some cases, analogies can be constructed using knowledge from closely related domains, even the same domain as that of the target problem. Referring back to the example at the start of the chapter, if the design engineer was familiar with another cable system that satisfied constraints similar to the problem on which she was working, she may have been able to use it as a source for constructing a close-domain analogy. The closeness of two domains is ultimately a matter of the number of superficial similarities between them, superficial similarities being those that involve features that are causally unrelated to the goals and constraints of the problem. If the design engineer, for example, was familiar with an old cable problem involving 10-gauge cables and was working on a new problem that also involved 10-gauge cables, cable size would constitute a superficial similarity. In this case, size is not important in how the problem is solved, but it could help in retrieving the old experience (Gentner, Ratterman and Forbus, 1993; Holyoak and Koh, 1987). Analogies can also be drawn from unrelated domains, such as in the example where a revolving door is aligned with the cable problem. In the experiments discussed in this chapter, we focus on analogies at this end of the spectrum - those, in other words, for which few, if any, superficial similarities exist between the source domain and the problem at hand. We refer to such instances as abstract analogies.

Given the potential complexity of mental representations, the processes associated with retrieval, alignment, and mapping can themselves be complex. At a fundamental level, though, it comes down to matching units of information between representations. In the case of amodal representations, matching by itself is straightforward. Given two predicate-based representations, for example, one could just match identical predicates. In the case of modal representations, matching might also be straightforward, in that one could look for identical perceptual features, parts, objects, scenes, events or image schemas. Modal representations might in fact have an advantage over amodal representations in making superficial matches since modal representations tend to make perceptual features, such as visual patterns, easier to detect. Cars and people, for example, might not seem analogous until one visually simulates the front of a car (headlights, grill, bumper, etc.) and compares that with an image of a face (Davies, Nersessian and Goel, 2001).

The most significant difference between modal and amodal representations for generating analogies potentially lies with matching abstract structure between representations. As mentioned above, abstract structure can be aligned by pairing identical predicates in the case of amodal representations and, similarly, by pairing identical image schemas in the case of modal representations. Modal representations, however, present another possibility. An image schema in a source representation might also be matched with a pattern of constraints in a target problem. A pattern of constraints, that is, might describe properties implicit in an image schema and hence be perceived as approximating it, resulting in alignment. In this case, the image schema serves as a sort of generic, encapsulated solution schema. If a source representation includes such a schema, one might expect it to be retrieved and aligned more easily with a target problem than one that lacks it, simply because any given image schema, like any perceptual state, can be described, and hence accessed, in so many ways.

In the three sections that follow, we describe experiments in which the perceptual context of source situations is manipulated using different kinds of external representations: narratives, physical activity, and diagrams. In the first two experiments context is manipulated for the purpose of either inhibiting or fostering the activation of particular image schemas (Craig, Catrambone and Nersessian, under review). The goal is ultimately to determine whether image-schematic structure will affect the ease with which analogies are made, as perceptual simulations predict. In the third experiment we study how the perceptual context of self-generated diagrams inhibits or fosters the sort of perceptual transformations required to convert a representation of a system in one domain into a solution to a problem in another.

4. PERCEPTUAL SIMULATION AND NARRATIVE SOURCES

When people read narratives their behaviour suggests that they begin linking together parts of long-term representations in working memory, producing short-term representations generally referred to as situation models, although sometimes also as mental models (Johnson-Laird, 1983; van Dijk and Kintsch, 1983). If one reads about a cat taking a nap, for example, one might retrieve a schema for cats and one for napping and link them together. Readers may embellish stories by inferring prototypical information (e.g., that cats sometimes purr when sleeping) as well as information that serves to add coherent connections between things that are specified in the story (e.g., that

cats nap on comfortable surfaces, like sofa cushions). The process of elaborating mental models likely serves to make the information more coherent, easier to remember, and easier to use in disambiguating subsequent references made in a text. It is so natural that readers often have a hard time remembering that which was explicitly mentioned in a text separately from that which they inferred (Bransford, Barclay and Franks, 1972).

We argue that readers may construct perceptual simulations that function like situation models with additional perceptual constraints. Evidence of additional perceptual constraints includes the fact that access to remembered details in a text often depends on spatial proximity (rather than, say, conceptual proximity) to the current focus of attention (e.g., Glenberg, Meyer and Lindem, 1987; Bower and Morrow, 1990) and on perceptual orientation (e.g., Franklin and Tversky, 1990). It also includes the fact that judgments about a text (regardless of whether the described situations are highly spatial or visual) are hampered when people are asked to hold unrelated visual images in memory, but not when they are asked to hold numerical sequences, even though the latter is considered more difficult (Fincher-Kiefer, 2001). Whether readers always generate perceptual simulations to track information presented in a text is, of course, open to speculation. Some texts - mystery novels, for example - might be too complex to track with one coherent perceptual simulation. At this point we are just considering what the consequences of constructing perceptual simulations might be for analogical problem solving, especially when contrasted with predicate-based situation models.

As it turns out, narratives have been studied extensively in analogical problem solving. In studies done by Gick and Holyoak (1980; 1983), for example, subjects were given a short story to read about a general leading an attack on a dictator's fortress. In the story, the general, knowing that he must use his entire army but also knowing that the army is too big to pass safely over the roads, splits up his men and has them converge on the fortress from different directions. Subjects are asked to study this story under the guise of a recall task. They are then given a problem to solve in which a doctor is trying to destroy a tumor using radiation, with the constraint that a beam of radiation strong enough to kill the tumor will also destroy vital tissue it passes through on the way to the tumor. Gick and Holyoak found that subjects who studied the fortress story were more likely to suggest converging weak beams of radiation on the tumor from different directions than were subjects who read an unrelated story. They argued that subjects may have induced a solution schema from the fortress story that was abstract enough to be aligned with the radiation problem. In focusing on abstract information contained the fortress story, Gick and Holyoak did not address the perceptual context itself. Interestingly, though, they found that subjects were less likely to come up with analogous solutions to the radiation problem if they read an abstract summary of the story than if they read the story with all the contextual details.

Either the abstract summary inadvertently missed a crucial piece of information, or the context of the story did, in fact, play a significant role in constructing analogies.

We hypothesize that if readers generate perceptual simulations to represent events described in a story, the image-schematic structure embedded in those simulations might depend to a large extent on how the story is written - from what perspective, for example, and in what context. Image-schematic structure could even depend on the use of metaphors in a text. Sentences like, "He blew him out of the water with a scathing rebuke," for example, might activate a simple, linear force schema, which in turn could be used to structure a simulation of two people arguing, despite the fact that the characters do not use physical force. In the case of Gick and Holyoak's (1980) fortress story, we argue that the presence of a particular image schema, namely a convergent-force image schema, will significantly increase the chance that the story is used to construct an analogical solution to the radiation problem. A convergent-force image schema, which might be learned through everyday activities like pinching and squeezing (or being pinched and squeezed for that matter), might be useful because it encapsulates information that can be aligned with constraints of the radiation problem. Specifically, the radiation problem requires that radiation be weak in all areas surrounding the tumor but strong at the site of the tumor, just like intensity grows at the point of focus in the convergent-force image schema. Incidentally, though, the fortress story used by Gick and Holyoak probably will not elicit such a schema, if only because the army in the story does not literally apply pressure to the fortress. Instead, it seems more likely that readers will represent a sequence of events in which a resource (the army) is split up and reunited at the target (the fortress) to get around an obstacle (landmines along the road). If that is the case, the fact that intensity grows at the point of focus will be left to be inferred from a sequence of maneuvers and hence not be readily available for aligning with the pattern of constraints inherent in the radiation problem.

To test our hypothesis we constructed two different stories, both derived from Gick and Holyoak's (1980) fortress story and both sharing the same predicate structure. The differences between the stories lie with the objects described and, hopefully, in the sense of force those objects elicit. One story (the "beehive story") is specifically meant to elicit a convergent-force image schema. It describes a man trying to kill a queen bee in a cell in the center of a beehive. The man first considers driving a big stick into the hive but decides not to, knowing that the big stick might set off the drones. Instead he pushes in two smaller sticks with less force from different sides so that they meet in the middle. Because the story describes the application of physical pressure in a convergent manner, we argue that it naturally elicits a convergent-force image schema. The other story (the "modified fortress story") tells of a man trying to kill a dictator in a bunker in the center of a compound. He first considers marching a big army into the compound but decides not to, knowing that the big army might set off landmines. Instead he marches in two smaller groups of men with less force from different sides so that they meet in the middle. The men in the fortress story do not literally apply pressure on the fortress. Instead their power adds up like a generic resource. As a result, a reader might construct a perceptual simulation of the events in the fortress story, but one that probably does not include a convergent-force image schema.

We started by giving 102 Georgia Tech undergraduates the radiation problem (Phase 1). After working on that problem for 4 minutes, we then gave them two stories to study, first an unrelated distracter story and second either the beehive story or the modified fortress story, depending on whether they were in the beehive or fortress condition respectively. After studying the stories, we gave them the radiation problem again (Phase 2), telling them that we were interested in seeing whether a period of incubation would help them solve the problem. Finally, we gave them the radiation problem a third time (Phase 3), along with a hint that one of the two stories they read earlier might be useful in solving it. Because the beehive story explicitly suggests the use of convergent force, we predicted that subjects in the beehive condition would perform better on the radiation problem than subjects in the fortress condition. After eliminating 9 subjects who produced a convergence solution to the radiation problem in phase 1, 20 out of 49 subjects in the beehive condition produced a convergence solution in Phase 2 as compared to 9 out of 44 subjects in the fortress condition. Of the remaining subjects, 22 out of 29 subjects in the beehive condition produced a convergence solution in Phase 3 as compared to 15 out of 35 subjects in the fortress condition. In both cases, the difference between the beehive condition and the fortress condition is reliable

The fact that significantly more subjects in the beehive condition solved the radiation problem before the hint suggests that, in general, the radiation problem was more similar to beehive story than to the fortress story. However, the fact that subjects also performed significantly better in the beehive condition *after* the hint suggests that a significant part of the advantage is due to structural similarities. Once, that is, subjects were given a hint to consider the stories, retrieval should not have been a problem, and hence superficial differences should not have given one story an advantage over the other (Holyoak and Koh, 1987). There could be other reasons why subjects who read the beehive story performed better on the radiation problem besides the presence of a convergent-force image schema. One is that sticks could have been perceived as looking like beams of radiation. Another is that people might have metaphorically represented the radiation problem itself using force. It is not uncommon, for example, for people to describe radiation as "hitting" a target. Third, and finally, there might be something about radiation and tumors that is semantically more similar to sticks and bees than to armies and dictators. The results, however, are consistent with the explanation that image-schematic structure accounts for some of the differences. Our subsequent experiments explore the image-schema hypothesis further.

5. PERCEPTUAL SIMULATION AND PHYSICAL ACTIVITY

Physical activity and conceptual thought have come together primarily in studies of gesture and language. Some of these studies are worth considering here to the extent that they hint about the significance of gesture *for* speech, not just as something that accompanies verbal communication. One finding in the gesture literature, for example, is that when gesturing is inhibited, speech production becomes choppy (Krauss and Hadar, 1999). Another is that blind children tend to gesture like sighted children, even when talking to other children they know are blind (Iverson and Goldin-Meadow, 1997). Gestures, in this case, appear to do more than embellish speech. They seem, rather, to be functionally tied to speech production and perhaps even to the conceptual thinking that underlies speech.

Complementing studies on gesture and language are those that directly address the impact of physical activity on conceptual thought. One relevant finding is that physical activity can prime sensibility judgements. When people, for example, are asked to briefly form a gesture with their hand and are subsequently asked to determine whether a given sentence makes sense, they tend to answer more quickly if the gesture is consistent with what the sentence describes (pinching one's fingers together, for example, would be considered consistent with a sentence like "He aimed a dart" but not with a sentence like "He rolled a ball") (Klatzky, Pellegrino, McCloskey, and Doherty, 1989). The fact that gestures can prime sensibility judgements, even when subjects are blind to the purpose of the gesture, is consistent with the idea that physical activity can actually help generate perceptual simulations. It seems unlikely, by contrast, that an unanalyzed gesture would activate specific predicates, given the lack of specificity in what a simple hand gesture might mean.

We hypothesize that physical activity can be used to embellish perceptual simulations by adding structure in the form of image schemas, since physical activity itself is always spatial and kinesthetic. Moreover, we argue that for physical activity to be useful it need not explicitly mimic events or situations under consideration. A lieutenant, for example, might wave his hands to simulate maneuvers on the battlefield, as an engineer might use her arms to simulate the physics behind cantilevered bridges. If physical activity can add image-schematic structure to perceptual simulations, then it should be possible to add a convergent-force image schema to the fortress story used in the last section, perhaps by having people physically enact the events in the story using their hands. A convergent-force image schema might, for example, be elicited by any sort of compressive activity, provided it is consciously attended to in the context of the story (just because it is acted out does not mean it will be used in a representation of the story). If a convergent-force image schema is activated, it could in turn help promote alignment and mapping with the radiation problem for the same reason we argued that the beehive story in the previous section would help.

To test our hypothesis we considered two possibilities, first that the abstraction involved in drawing a diagram of the story might invoke a convergent-force schema (such as in physics diagrams) and, second, that the activity of physically colliding two objects might also invoke it (by activating a kinesthetic sense of force). We gave 104 Georgia Tech undergraduates in three conditions the fortress story to study, telling them that we were interested in how people recalled stories in different modalities. After they studied the story, we had them verbally recall it in their own words. Subjects in the first condition (the "sketch" condition) were additionally asked to draw a simple map-like diagram depicting the action in the story during recall. Subjects in the second condition (the "gesture" condition) were additionally asked to slide wood blocks on the desktop to show what happened in the story during recall. There were three blocks total, one for the dictator's bunker measuring 1.5"x1.5"x1.5" and two for the army measuring 1"x1.5"x3.5". Subjects in both conditions were asked to carefully focus on what they were doing as they recalled the story. In the third condition (the "verbal" condition), subjects were not given any instructions in addition to verbally recalling the story. After the recall phase, all subjects were given the radiation problem under the guise of a separate problem-solving study. After working on the problem for 8 minutes, they were given it again with the hint that the fortress story might be useful.

We predicted that subjects in both the sketch and the gesture conditions would perform better on the radiation problem than subjects in the verbal condition simply because the additional activities might enhance and reinforce their memory of the story. More importantly, though, we predicted that subjects in the gesture condition would perform best overall because they would be more likely to simulate the story using a convergent-force image schema and would in turn have an easier time retrieving and aligning the story with based on the constraints of the radiation problem.

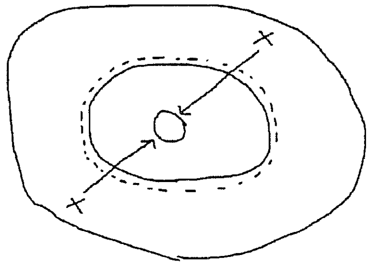

Figure 2. An abstract sketch illustrating the events in the fortress story drawn by a subject prior to working on the radiation problem.

Consistent with our predictions we found that significantly more subjects solved the radiation problem in the gesture condition than in the sketch condition. After eliminating subjects who already knew about treating tumors with converging beams of radiation, we found that 10 out of 33 subjects in the sketch condition solved the radiation problem before the hint to consider the story, versus 17 out of 33 in the gesture condition and 4 out of 21 in the verbal condition. The difference between the gesture condition and the sketch condition (52% vs. 30%) is reliable. The difference between the solution rates in the sketch condition and the verbal condition (30% vs. 19%), however, is not. There were also no reliable differences in solution rates among any of the conditions after the hint was given to consider the story. The lack of significant differences in solution rates after the hint could indicate that the task was not sensitive enough to reflect structural differences. Unlike in the previous experiment, no distracter story was used. Subjects thus might have been able to use a "solution-focusing" strategy based on a surface representation of the story and on pragmatic associations (e.g., between agents and goals) to construct a successful solution to the radiation problem (Gick and Holyoak, 1980; Keane, 1985).

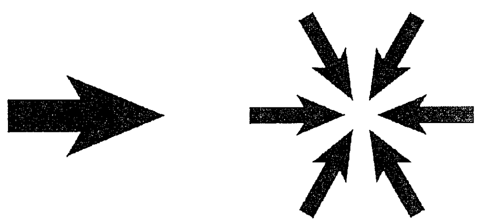

Figure 3. Convergence arrows similar to those used by Gick and Holyoak (1983) to help illustrate the events in the fortress story.

Before concluding this section, an additional comment on diagrams is due. Self-generated diagrams, as we found out, do not help as much as physically colliding blocks in facilitating analogous solutions to the radiation problem. We argued that although the activity of sketching is physical and although the resulting images might resemble elements of the target problem (Figure 2), sketching is not likely to activate a convergent-force image schema. One might still ask, however, why *viewing* such diagram (whether self-generated or not) does not elicit a convergent-force image schema either. The answer is that visual displays might elicit the schema, but that to do so requires visual features atypical of the sketches our subjects produced. First, viewing simple displays, like those showing converging arrows (Figure 3), is not likely to help much. Gick and Holyoak (1983), for example, found that subjects given converging-arrow diagrams in addition to the fortress story performed no better on the radiation problem than subjects who read the fortress story without the diagrams. Simple diagrams of this sort only convey a spatial schema, a pattern, in other words, similar to the convergent-force image schema but absent the force component. Without the force component, they fail to convey the fact that intensity builds at the point of focus, which we argued before might make alignment with the radiation problem significantly easier. Two experiments using modified visual displays have addressed this issue. In one study, by Beveridge and Parkins (1987), subjects were shown transparent, colored plastic strips that, when fanned out radially (in a star-like pattern), became increasingly light at the periphery where the strips did not overlap and dark at the center where they did overlap. The plastic strips show intensity growing at the center (albeit without convergent forces), and subjects who viewed this display did, in fact, perform significantly better on the radiation problem than subjects who viewed a static arrow diagram. In a separate study by Pedone, Hummel and Holyoak (2001), subjects were shown an animated version of the static converging-arrow diagram. Such a display appears ideal for eliciting a convergent-force image schema since the animated arrows literally "collide" at the center. Consistent with this analysis, subjects who viewed the animated diagram performed significantly better on the radiation problem than subjects who viewed a static diagram.

6. DYNAMIC PERCEPTUAL SIMULATION AND DIAGRAMMATIC SOURCES

Up to this point we have referred to perceptual simulations as if they were static. Because the perceptual system is capable of processing motion and other dynamic patterns independently of other perceptual features (such as size and shape), it is also plausible that people can store and later invoke dynamic image schemas for the purpose of structuring dynamic simulations. A simple image schema representing linear motion, for example, could be stored and then reactivated to imagine things in motion (Mandler, 1992). As with static simulations, coherence might determine when and how dynamic schemas are used. Physical properties of objects, that is, need to mesh with actions in a physically consistent way (Glenberg, 1997). An interesting finding that illustrates this possibility comes in a study by Schwartz and Black (1996). In one their experiments, Schwartz and Black asked subjects to view a diagram showing two gears meshed together, one larger than the other, and then decide whether marks on the circumferences of each gear would line up with one another if the gears were rotated. By comparing the angular disparity of the marks with response times, Schwartz and Black identified different strategies used by subjects to complete the task, one of which was mentally rotating the gears in a synchronous fashion. More interestingly, they found that when the gears were drawn to have rough surfaces rather than smooth surfaces, subjects were more likely to use the mental rotation strategy over all others, as if roughness made it easier to imagine one gear driving the other. In this case, the physical affordances implicit in the drawings might have affected the ease with which simulations based on those drawings were transformed.

In this section we consider the possibility that diagrams might help people construct perceptual simulations of source situations that can in turn be dynamically transformed into solutions to target problems. The process is similar to one of abstracting deep structure from a source situation and embedding it into a target situation. In other work we have referred to this as "generic abstraction" (see, e.g., Nersessian, 1999). In this section we focus on how perceptual simulations are transformed and specifically at how the perceptual context implicit in a diagram might influence the process. Unlike the experiments previously discussed, the experiments presented here all assume that source and target representations are available, thereby eliminating problems associated with retrieval.

We hypothesize that if subjects have access to a source solution but are not sure how to map it onto a target problem, a diagram of the former could help by setting up a perceptual simulation - static or dynamic - that can be transformed to fit the target situation. Crucially, though, we also hypothesize that the kinds of physical affordances implicit in such diagrams can have an impact, either enhancing or impeding the process of dynamical simulation.

To explore our hypothesis, we conducted an experiment in which we gave 209 Georgia Tech students two problems to work on along with a hint that the two might be solved in similar ways. The first problem, which was intended to serve as a potential source, asked subjects to design a door system that would let people in and out of a room without letting the air inside be contaminated by the air outside. The second problem, which was intended to serve as a target problem, asked subjects to design a pole capable of passing through another pole at right angles without breaking apart. We anticipated that if subjects came up with a solution to the door problem that involved redundant doors (either a vestibule system or a revolving door), they might be able to use it to come up with an analogous solution to the pole problem, namely one that incorporated redundant connections. We also predicted that although diagrams depicting a redundant-door solution might help facilitate the construction of analogous solutions, it would depend on the perceptual affordances implicit in those diagrams. Because the experiment is exploratory, we were not able to specify in advance what sort of differences might exist.

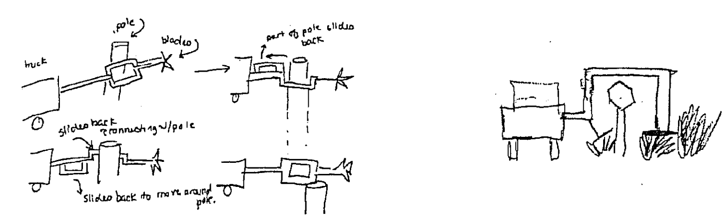

Figure 4. A good solution (left) to the "pole" problem that uses a redundant gate mechanism to allow one pole to pass through another, and a bad solution (right).

Of the 209 subjects who participated, 184 produced a "redundant-door" solution to the door problem. Of these subjects, 31 (17%) produced a solution to the pole problem that also involved some kind of redundant mechanism (Figure 4). In contrast, only 2 out of the 25 (8%) subjects who did not generate a redundant door solution to the door problem solved the pole problem. Drawing a diagram of a redundant door solution did not seem to help at all. Of the 131 subjects who drew diagrams, 22 (17%) solved the pole problem, compared with 9 of 53 (17%) who produced textual descriptions without diagrams.

The most interesting findings had to do with the types of diagrams people created. We found, for example, that the physical perspective of the diagram was correlated with whether subjects subsequently solved the pole problem (Figure 5). Of the 58 subjects who drew diagrams from a top view, 14 (24%) solved the pole problem, compared with only 8 of the 75 (9%) who used either used an oblique view, a side view, or an undetermined (either side or top) view. The difference in solution rates between those who drew top-view diagrams and those who drew other diagrams is reliable. We also found that the way the door system was drawn in relation to adjoining walls was correlated with whether subjects subsequently solved the pole problem. In particular, if the door system shared walls with a surrounding building – if, in other words, surrounding walls were continuous with the door system – subjects were much less likely to produce an analogous solution to the pole problem (Figure

6). Of the 104 subjects who drew physically continuous wall systems, only 10 (10%) solved the pole problem, as compared to 12 of 27 (44%) who drew discontinuous wall systems. These differences are also reliable.

The correlation between the frequency with which analogous solutions are generated and the perspective and physical continuity of source diagrams is consistent with the idea that subjects use diagrams to generate dynamic perceptual simulations and that they used these simulations, in turn, to try to adapt the source solution to the target context. In the analogy we are exploring, a significant difference in context exists. In the door problem, the object being designed (the door) is fixed relative to what passes through it (people). In the pole problem, however, the opposite is true. The object being designed (the pole) is moving relative to that which passes through it (another pole that happens to be stationary). If subjects dynamically simulate the door system and transform it into a pole solution, the motional context will have to be reversed. Subjects, in other words, will have to get the door system in motion. If this is the case, drawing the walls of the building to be continuous with the door system might cause the door to be perceived as being physically *encumbered*, and hence firmly rooted. The perspective of a diagram might produce an even more glaring perceptual discrepancy. If a door system is simulated from a side view it will be almost impossible to reconcile with the context of the pole problem, since, given the way the pole problem is stated, it is hard to see what is going on from the side. Based on our data, the physical affordances in the diagrams pertaining to the creation of a *dynamical* perceptual simulation appear to influence the transformation from source to target situation.

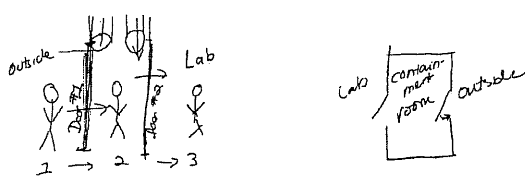

Figure 5. A solution to the door problem viewed from the side (left) and a solution viewed from above (right).

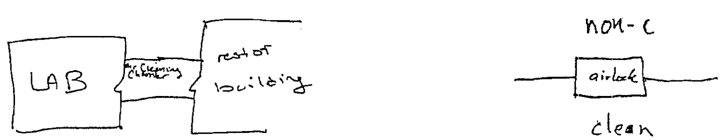

Figure 6. A solution to the door problem in which the door system shares a wall with the surrounding context (left) and one in which the door system does not share walls.

7. SUMMARY

The research discussed in this chapter has focused on the perceptual context of source situations - perceptual properties, in other words, that are causally disconnected from the goals and constraints of a problem as conveyed through predicate-based representations. First, we argued that different perceptual contexts could activate different image schemas, resulting in representations that differ in how they encapsulate information and thus in how easily they can be aligned with target problems. Our findings in the first two experiments are consistent with this analysis, although we acknowledge that they do not rule out other influences. We also argued that how different situations are portrayed through diagrams might affect how easily they are dynamically transformed into solutions to target problems. Our findings in the third experiment are consistent with this analysis in that contextual properties of source diagrams - notably physical perspective and physical continuity - were found to correlate with the production of analogous solutions to a target problem.

More work is needed to better understand the relationship between the perceptual context of source situations and analogies based on those situations, and to further develop the theory of perceptual simulation. Following up the research presented here, we are currently working on new studies that make a more concerted effort to control for superficial properties of source situations - both semantic and perceptual - when manipulating image-schematic structure. We are also taking a closer look at the role of diagrams in analogical problem solving, this time controlling more precisely for the presence of superficial features when manipulating things like physical perspective and physical continuity.

Insights about perceptual context and perceptual simulation might also be gained reflexively, by applying the theory to real-world problem solving settings. How does an architect, for example, use drawings, physical models, embodied metaphors and fictional narratives when formulating analogies between building systems and other domains? How does a physicist make use of external representations in formulating theories? It might be found that perceptual context ultimately does more than guide interpretation or determine the choice of problem-solving strategies. Perceptual context, in fact, might be seen as a problem-solving tool itself, in the sense that problem solvers might consciously endeavour to create rich contexts for the purpose of facilitating their own thinking.

REFERENCES

Barsalou, L.W., 1992, Frames, concepts, and conceptual fields, in: *Frames, Fields, and Contrasts: New Essays in Semantic and Lexical Organization*, E. Kittay and A. Lehrer, eds., Lawrence Erlbaum, Hillsdale, NJ, pp. 21-74.

Barsalou, L.W., 1993, Structure, flexibility, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols, in: *Theories of Memory*, A.C. Collins, S.E. Gathercole, and M.A. Conway, eds., Lawrence Erlbaum, London, pp. 29-101.

Barsalou, L.W., 1999, Perceptual symbol systems, *Behavioral and Brain Sciences* 22:577-609.

Beveridge, M. and Parkins, E., 1987, Visual representation in analogical problem solving, *Memory and Cognition* 15:230-237.

Bower, G.H. and Morrow, D.G., 1990, Mental models in narrative comprehension, *Science* 247:44-48.

Bransford, J.D., Barclay, J.R., and Franks, J.J., 1972, Sentence memory: A constructive versus interpretive approach, *Cognitive Psychology* 3:193-209.

Clement, C.A. and Gentner, D., 1991, Systematicity as a selection constraint in analogical mapping. *Cognitive Science* 1:89-132.

Craig, D.L., Catrambone, R., and Nersessian, N.J., under review, The role of implicit perceptual factors in analogical problem solving.

Davies, J., Nersessian, N.J., and Goel, A.K., 2001, Visual models in analogical problem solving, *Georgia Institute of Technology Cognitive Science Technical Report GIT-COGSCI-2001/03*.

Fincher-Kiefer, R., 2001, Perceptual components of situation models, *Memory and Cognition* 2:336-343.

Franklin, N. and Tversky, B., 1990, Searching imagined environments, *Journal of Experimental Psychology: General* 119:63-76.

French, M.J., 1994, *Invention and evolution: Design in nature and engineering*, Cambridge University Press, New York.

Gentner, D., 1983, Structure-mapping: A theoretical framework for analogy, *Cognitive Science* 7:155-170.

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

Gentner, D. and Stevens, A.L., eds., 1983, *Mental Models*, Lawrence Erlbaum, Hillsdale, NJ.

Gick, M.L. and Holyoak, K. J., 1980, Analogical problem solving, *Cognitive Psychology* 12:306-35.

Gick, M.L. and Holyoak, K.J., 1983, Schema induction and analogical transfer, *Cognitive Psychology* 15:1-38.

Glenberg, A.M., 1997, What memory is for, *Behavioral and Brain Sciences* 20:1-19.

Glenberg, A.M., Meyer, M., and Lindem, K., 1987, Mental models contribute to foregrounding during text comprehension, *Journal of Memory and Language* 26:69-83.

Holyoak, K.J., and Koh, K., 1987, Surface and structural similarity in analogical transfer, *Memory and Cognition* 15:332-340.

Iverson, J.M. and Goldin-Meadow, S., 1997, What's communication got to do with it? Gesture in children blind from birth. *Developmental Psychology* 33:453-467.

Johnson, M., 1987, *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*, University of Chicago Press, Chicago.

Johnson-Laird, P.N., 1983, *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*, Harvard University Press, Cambridge, MA.

Keane, M., 1985, On drawing analogies when solving problems: A theory and test of solution generation in analogical problem-solving tasks. *British Journal of Psychology* 76:449-458.

Klatzky, R.L., Pellegrino, J.W., McCloskey, B.P., and Doherty, S., 1989, Can you squeeze a tomato? The role of motor representations in semantic sensibility judgments, *Journal of Memory and Language* 28:56-77.

Kosslyn, S.M., 1994, *Image and Brain: The Resolution of the Imagery Debate*. MIT Press, Cambridge, MA.

Krauss, R.M. and Hadar, U., 1999, The role of speech-related arm/hand gestures in word retrieval, in: *Gesture, Speech, and Sign*, L.S. Messing and R. Campbell, eds., Oxford University Press, New York, pp. 93-116.

Lakoff, G., 1987, *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, University of Chicago Press, Chicago.

Lakoff, G. and Johnson, M., 1998, *Philosophy in the Flesh*, Basic Books, New York.

Mandler, J.M., 1992, How to build a baby: II. Conceptual primitives, *Psychological Review* 4:587-604.

Minsky, M., 1975, A framework for representing knowledge, in: *The Psychology of Computer Vision*, P.H. Winston, ed., McGraw Hill, New York, pp. 211-277.

Nersessian, N.J., 1999, Model-based reasoning in conceptual change, in: *Model-Based Reasoning in Scientific Discovery*, L. Magnani, N.J. Nersessian, and P. Thagard, eds., Kluwer Academic/Plenum Publishers, New York, pp. 5-22.

Palmer, S.E., 1975, Visual perception and world knowledge: Notes on a model of sensory-cognitive interaction, in: *Explorations in Cognition*, D.A. Norman and D.E. Rumelhart, eds., Freeman, San Francisco, pp. 279-307.

Pedone, R., Hummell, J.E., and Holyoak, K. J., 2001, The use of diagrams in analogical problem solving, *Memory and Cognition* 29:214-221.

Rumelhart, D.E. and Ortony, A., 1977, The representation of knowledge in memory, in: *Schooling and the Acquisition of Knowledge*, R.C. Anderson, R.J. Spiro, and W.E. Montague, eds., Lawrence Erlbaum, Hillsdale, NJ, pp. 99-136.

Schank, R. and Abelson, R., 1977, *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*, Lawrence Erlbaum, Hillsdale, NJ.

Schwartz, D.L. and Black, J.B., 1996, Analog imagery in mental model reasoning: Depictive models. *Cognitive Psychology* 2:154-219.

Smith, E.S. and Medin, D.L., 1981, *Categories and Concepts*, Harvard University Press, Cambridge, MA.

van Dijk, T.A. and Kintsch, W., 1983, *Strategies of Discourse Comprehension*, Academic Press, New York.