

# Analogical Learning: Mapping and Integrating Partial Mental Models

Mark Burstein  
BBN Laboratories

Beth Adelson  
Tufts University

## Abstract<sup>1</sup>

Descriptions of scientific and technical systems take a number of different forms. Depending upon the *purpose* of a description, it may focus on a system's behavior, causality, physical or functional topology, or structural composition. An *analogical explanation* used to teach someone about such a system is also typically geared to one or another of these purposes. In this paper we describe some research leading to the development of a theory of the role of *explanatory model types* in the generation of analogical mappings. The work is motivated by the larger question of how explanations presented as analogies are applied by students learning about new domains. Our long term goals are (1) the development of a theory of *purpose-guided analogical learning*, based on a coherent taxonomy of mental model types, and (2) the development of a theory of the integration of partial mental models during learning, using principles for relating different explanatory model types.

## 1. Analogical Mapping of Different Explanation Types

In recent years, researchers in artificial intelligence and cognitive psychology have begun to focus more attention on the study of analogical reasoning and its role in learning and problem solving, particularly in scientific and technical domains. A number of these researchers have independently converged on a class of models of analogical learning which stress the mapping process (Gentner, 1983, Burstein, 1986, Thagard and Holyoak, 1985, Carbonell, 1986). By this class of process model, an underlying conceptual model of a familiar source or *base* domain is abstracted and mapped to an unfamiliar *target* domain. The mapped model is then used to build a new model in the target domain. Analogical learning theories need to specify the mapping process in some detail, since mapping dictates what can be postulated from any given analogy. For example, *what* is mapped must be constrained by the *purpose* of the analogy (Kedar-Cabelli, 1985, Burstein, 1985, Thagard and Holyoak, 1985, Winston et al., 1983).

Along with specifying the mapping process, theories of analogically-based learning must also account for the use of multiple analogical models. Typically, no single analogical model can be found that completely and accurately describes a non-trivial target system. That is, when subjects generate full and correct explanations of a system, they often need to use *multiple partial models*, of varying types, and at varying levels of abstraction (Burstein, 1983, Collins and Gentner, 1983, Collins, 1985, Adelson, 1984, Sternberg and Adelson, 1978, Coulson et al., 1986). For example, Collins and Gentner found that untutored subjects used as many as three analogical models to answer novel questions about evaporation (Collins and Gentner, 1982, 1983, 1987). Different kinds of models were used to answer different kinds of questions and frequently several models were used together.

---

<sup>1</sup>This work was supported in part by grants from ARI and NSF.

In another case study, Burstein (Burstein, 1985, 1986) found that three analogies were commonly used to teach students about variables and assignment for the programming language BASIC. One analogy described assignment as being like "putting things in boxes". Another analogy related assignment statements to algebraic equalities. A third analogy related the encoding and retrieval processes of human memory to analogous processes in the computer. Again, the analogies suggested partial models of different types. The box and memory analogies each describe an action that causes a result and so, as we explain in detail below, we regard these as *mechanistic causal* analogies. The analogy to algebraic equalities contributes a *behavioral* model. It explains how to infer values for variables without providing a mechanism accounting for *how* those values are derived and assigned by the computer.

## 2. Developing a Theory of Purpose-guided Analogical Learning

Our current research is aimed at developing a theory covering the two related issues of *mapping* and *integrating* partial mental models. Specifically, we are addressing the questions of: (1) how analogically-based models of different types are mapped to a new domain, and (2) how these different kinds of partial models are integrated in a target domain. This paper focuses on the mapping process, although we will touch on the integration issue as well.

Both parts of our theory are being developed around a taxonomy of explanatory partial model types. The research presented here is part of a series of protocol experiments being conducted to produce a detailed account of the process that maps partial models of various types. The protocols are being used to identify the kinds of relations that are (and are not) included in target models developed by mapping partial models of a given type. We have found that models of different types are distinguished by the different kinds relations that they contain, and that models of a given type map to form new models of the same type (i.e., models containing the same relations). By making the distinction between model types explicit in our theory of analogical learning, we hope to provide an account of analogical structure mapping that has clear pragmatic constraints on the amount and type of information mapped at one time from a base domain.

Future research will focus more heavily on the issue of model integration. Integrating multiple models is an important part of the learning process, since there are many situations that can only be explained by a combination of inferences from several different partial models. As we explain in section 4, we foresee that two kinds of integration processes may be used for combining partial analogical models. First, there is a reasoning process that functions to relate newly acquired partial models within a domain<sup>2</sup>. This process is based on principles about how, in general, partial models are inter-related. A second kind of integration process may be used when *adapting* partial models from an already active analogical *source* domain. In this case, information about how partial models are related in the source domain, can be used to avoid reasoning from first principles in the target domain. Since an understanding of the integration process is dependent on an understanding the mapping process, we have chosen to focus first on the mapping process.

---

<sup>2</sup>This process is general to learning in that it functions whether the models have been acquired by analogy or directly by observation or instruction in the target domain.

### 3. A Taxonomy of Explanatory Model Types

Our investigations of both the mapping and integration processes depend heavily on a well-defined taxonomy of model types. For our initial experiments, we have developed a working taxonomy which we expect will capture a broad set of models used in analogical learning (Figure 3-1). In developing our taxonomy of models for *analogical learning*, we have revised a taxonomy developed by colleagues at BBN initially to *categorize* textbook explanations of complex physical systems (Stevens and Collins, 1980, Stevens and Steinberg, 1981, Collins, 1985). The taxonomy was formulated during the development of the STEAMER ICAI system (Williams, Hollan and Stevens, 1981). It provides relevant background support for our work in that it has been used to show that people often use *several* different kinds of explanatory models of a single system in trying to produce a full explanation (Collins, 1985, Weld, 1983).

- 
- **Structural models** are used to describe systems in a time-invariant manner. Structural models include:
    1. **Componential models** simply list components.
    2. **Topological models** specify configurations where the logical or functional connections between components are preserved.
    3. **Geometric models** preserve the quantitative, spatial relations between components.
  - **Dynamic models** describe changes that occur in a system over time. Dynamic models include:
    1. **Functional/Behavioral models** describe a system as a "black box", in terms of inputs and outputs.
    2. **Internal Structure models** break the system down into interactions between various components. These models include:
      - **Mechanistic Causal models** describe unique behaviors for each component and break events into causal chains. These include *Action Flow* models, where some substance or energy flows through the system and more abstract *Information Flow* models, where information is described as passing between components. These models typically dictate how the outputs of individual components cause state changes in other, topologically connected components, leading to an account of the behavior of the system as a whole.
      - **Aggregate models** describe systems where the components behave in a uniform manner, subject to global constraints. In these models, components are represented prototypically, in terms of general behavioral characteristics of the group. Individual features of components are represented by distributions of values.
      - **Synchronous models** describe causal systems where events or forces occur synchronously.

---

Figure 3-1: A Taxonomy of Explanatory Model Types

The data presented here is being used to develop a model of the analogical mapping process. Our theory suggests that *explanatory purpose*, the type of model required by the student for his ongoing problem solving or question answering activity, strongly affects the selection of base domain features and relationships for mapping. Different explanatory purposes are characterized by different types of models, which in turn can be shown to be based on different structural relations (causal relations, function/goal relations, spatial or topological relations, etc.) For example, within mechanistic causal models,

temporal/causal relationships are used to relate the behaviors of a connected set of components, in order to explain a system's overall behavior. However, in aggregate models, local constraints are replaced by global constraints (e.g., conservation laws), and individual components give way to representative or prototypical entities with properties characterized as distributions.

#### 4. Protocols of Mapping By Model Type

Our theory suggests that the type of model selected and mapped during learning is constrained by the aspect of the target situation made salient by the learning task. This prediction generates the following hypothesis: When explaining a given aspect of a target domain, subjects should map only a subset of the base domain and that subset should be coherent<sup>3</sup> and reflective of the purpose of the learning task.

The following situation was used to provide an initial test of this hypothesis. Subjects were provided with an analogical model of a computer programming constructs<sup>4</sup>. The subjects were all naive to the target construct and familiar with its base domain analog. Subjects were then asked to answer a set of questions about various aspects of the initially unfamiliar target domain concept<sup>5</sup>. This procedure was then repeated until subjects had been taught about the three constructs used in this study: queues, stacks and sorting. As an example, Figure 4-1 shows the texts used in teaching the concept of queues and the questions that subjects received following the texts<sup>6</sup>.

Our hypothesis that purpose constrains selection and mapping in a way that results in a coherent and appropriate partial mental model will be supported if subjects who have a complete base domain model map only the part of the base domain model that is relevant to the question they are answering.

In order to see whether subjects were selecting and mapping partial models we constructed preliminary behavioral and causal representations for each of the concepts that the subjects had been taught (queues, stacks, and sorting). As an example, a sketch of our behavioral and causal representations for queues, in both the base and target domains, is given in figure 4-2. The representations can be thought of as vertical behavioral or causal chains of which can be read from top to bottom.

The answers to each behavioral and causal question appearing in the recorded protocols for each subject were analyzed to see how clearly they corresponded to our behavioral and causal representations of the concept. The number of times that elements from the causal or behavioral representations occurred was counted for both the behavioral and causal questions.

In answering *behavioral* questions about queues, stacks and sorting, subjects made, on the average, 2.0 references to behavioral elements and no references to portions of the causal models that were not

---

<sup>3</sup>In the sense suggested by Gentner's *systematicity principle* (Gentner, 1983).

<sup>4</sup>As described below, the subjects varied in their levels of programming experience.

<sup>5</sup>In this study we selected behavioral and causal models as a subset of the taxonomy presented in Figure 3-1.

<sup>6</sup>The full design of the study is not presented here. Text type was crossed with question type and order of presentation was counterbalanced over the full set of protocols; subjects received only one written description of the base domain situation which stressed either behavioral, causal or behavioral and causal aspects of the already familiar base domain situation. This was followed by a causal question, a behavioral question and a question about the relationship between behavior and causality. Differences in the written descriptions of the already familiar base domain did not have a discernible effect on the results described below.

---

*Behavioral Analogy:*

Frequently, at Mary Chung's (restaurant) there are people waiting to be seated. Mary keeps track of who to seat in such a way that the first person to be seated next is always the person who, among those currently waiting, came in first.

In sending files to the printer the same situation often occurs; several files need to be printed but only one can be printed at a time. In this case the computer resolves the problem in the same way that Mary does.

*Causal Analogy:*

Frequently at Mary Chung's, Mary has a list of people who are waiting to be seated and served. Whenever a person enters Mary puts their name at the bottom of the list. Whenever a table becomes vacant Mary calls the name at the top of the list, gives that person a table and then crosses that name off the list.

In sending files to the printer the same situation often occurs; several files need to be printed but only one can be printed at a time. In this case the computer resolves the problem in the same way that Mary does.

*Behavioral Question:* If Karen and then Janet and then Amy and then Karen all typed print commands one right after the other, in that order, what would the computer do?

*Causal Question:* Describe what you would do if you were the computer keeping track of print requests.

*Behavioral/Causal Integration Question:* Explain why/how the computer's method of keeping track of print requests produces the correct result? That is, what is the relationship between what gets done and how it gets accomplished?

---

Figure 4-1: Presented Versions of an Analogy to Queues

also part of the behavioral model. That is, *all* references to components of the causal models were references to elements that appeared in our representations of *both* the behavioral and the causal representation (Figure 4-2), because they referred to the goal and starting states in those representations. No references were made to portions of the causal model that were purely descriptions of mechanism.

In answering *causal* questions about queues, stacks and sorting, subjects made, on the average, 5.8 references to causal elements and on the average, .9 references to elements that were in both the causal and behavioral representations. Again, all of the behavioral elements that appeared in subjects' answers to causal questions were elements referring to the start and goal in those representations. This implies that a causal account of *how* something works is, in some sense, not coherent unless some goal or purpose oriented statement is included regarding *why* the mechanism is needed (Adelson, 1984).

It seems that, as our theory predicts, subjects are mapping *coherent*, purpose-oriented partial models from the target to the base domain. The models are coherent in that they provide an adequate basis for responding to questions of a given type. (This can be seen in the excerpts from the protocols that are presented below.) The models are partial in that they do not provide a complete account of the system being examined.

If our result suggests that purpose, as characterized by the *type* of model, does constrain selection and mapping, it also confirms that there are interesting and intuitively plausible interdependencies between the kinds of information mapped with each type of model. Earlier we mentioned that any learning theory needed to include an account of how partial models were integrated. Our subjects responses to

behavioral and causal questions suggest that the initial and goal states that occur in both types of models can be used to provide a bridge between the two. This is an example of the kind of knowledge that is required for our postulated *model integration process*. That is, a behavioral and a causal model of a given construct can be related by their shared start and goal states. Other types of models will need similar kinds of criteria to be related to or integrated with each other. Since these criteria are not domain-specific, the integration process should be applicable whether the models have been acquired through analogy or learned directly.

In the remainder of this section, we present excerpts from the protocols that motivate some central issues for theories of analogical learning.

Subject J<sup>7</sup> received the behavioral version of the analogy. She was then asked the following behavioral question.<sup>8</sup>

- 
- E: (*Behavioral Question:*) If Karen and then Janet and then Amy and then Karen all typed print commands, one right after the other, in that order, what would the computer do?
- J: Give me the order again.
- E: Karen, Janet, Amy, Karen.
- J: All at the same moment?
- E: One right after the other. So not exactly a tie.
- J: Ok, well it would say that Karen is first... and, uh, it's just like, for me it's just like the thing you get when you call AAA and they tell you "Don't hang up, your calls are being taken in order."
- 

J's last response, ("It's just like the thing you get when you call AAA,") is interesting because it brings up the issue of our theory's *justification process* in which newly mapped models are tested and debugged. After mapping the behavioral model, J is reminded of a recent situation where a computer-like machine was performing the kind of behavior she had just been told that computers use for print queues. That is, by our theory, when J mapped the behavioral model of first-come-first-serve from the Mary Chung (base) domain to the computer (target) domain, she appropriately placed a computer in the "agent" role that Mary had played. This new behavioral description in the target domain then triggered a *reminding* of a similar behavioral model, where the agent, a machine similar to a computer, achieved the same goal (Schank, 1982). Recalling the behavior of the AAA phone system can be seen as an attempt by J to justify or test the adequacy of her newly formed behavioral model of a queue by comparing it to a similar, already-known behavioral model.

The fact that J could have mapped a causal model when answering the behavioral question, but did

---

<sup>7</sup>J is a researcher in music cognition and a self-taught (LOGO) programmer.

<sup>8</sup>E is the experimenter.

Figure 4-2: Base and Target Domain Models of Queues

<b>Behavioral Model</b>		
<b>Target</b>		<b>Base</b>
<i>Scene 1:</i>		<i>Scene 1:</i>
Print file command issued.	[B1]	Patron enters restaurant.
	<i>then</i>	
	[B2]	Mary takes person's name.
	<i>and</i>	
File 'queued' for printing	[B3]	Patron waits till called.
	<i>then</i>	
(when print request is first among those remaining)	[B4](when Patron is first among those waiting to arrive)	
	<i>then</i>	
File is printed.	[B5] Patron is seated.	
	= <i>Goal State</i>	
<b>Causal Model</b>		
<b>Target</b>		<b>Base</b>
<i>Scene 1:</i>		<i>Scene 1:</i>
Print file command issued.	[C1]	Patron enters restaurant.
	<i>and</i>	
Printer not free.	[C2]	All tables full.
	<i>initiates</i>	
	[C3]	Mary takes person's name.
	<i>enables</i>	
Request put on end of queue.	[C4]	Mary puts name on bottom of list.
	<i>results</i>	
File in queue	[C5]	Name on list
<i>Scene 2:</i>		<i>Scene 2:</i>
Printer free.	[C6]	Table empty.
	<i>and</i>	
File is on top of Queue	[C7]	Patron name is on top of list
	<i>enables</i>	
Queue pointer incremented	[C8]	Mary crosses out Patron's name.
	<i>and</i>	
File is Removed from Queue	[C9]	Mary calls Patron's name.
	<i>enables</i>	
	[C10]	Patron hears name
	<i>initiates</i>	
	[C11]	Patron goes to Mary
	<i>enables</i>	
File passed to Printer	[C12]	Patron follows Mary to table
	<i>enables</i>	
File is printed.	[C13]	Patron sits at table.
	= <i>Goal State</i>	

not do so seems clear when her response to the behavioral question is compared to her response to the causal question that followed:

- 
- E: (*Causal Question:*) Describe what you would do if you were the computer keeping track of print requests.
- J: Label them, I suppose, or put them if you want to think of it spatially.
- E: If you want to draw it ...
- J: OK, you've got a potential for something like this...  
(draws a row of boxes or 'slots')  
and you can *shove* things into these slots ...  
and Karen calls first  
so that becomes Karen, (writes Karen in 1st slot)  
and Janet..  
and then Amy and then Karen (writes each name in a successive slot)  
and I assume it's the same Karen. But, but, this (slot)  
is probably assigned some label like '1' and so she (Karen)  
will be '4' here and '1' there.  
So Karen is 4, or whatever,  
and when that is finished (points to the first slot)  
this becomes '1'. (points to the second slot)
- When that is finished (points to the second slot)  
this becomes '1'. (points to the third slot)
- 

J clearly has *pieces* of the causal mechanism, rather than a coherent causal model that had been mapped during the previous behavioral question, since she is able to answer the causal question, but in a bit-by-bit rather than an all-of-a-piece manner.

Another subject, K, received the causal version of the analogy and then was asked a causal question<sup>9</sup>. K's answer is interesting because, for a non-programmer, she develops a surprisingly good causal model. She does this by using the *analogically related causal description* she had been given. As suggested above, in order to produce a well-motivated and coherent description, she includes the related behavioral goal in her mapping from the base domain:

- 
- E: (*Causal Question:*) If Mark and Annie and Beth had typed requests to the computer, one after another, how would it (the computer) keep track of that?
- K: Can the computer work like a calculator that has a memory?
- E: Yes, it can. That's right, you can have told the computer in advance.
- K: It can take all of these names, put it in the memory, and then pull the first one out.
- 

<sup>9</sup>K is an antique book dealer. She has used several PC-based software packages, but she does not program.

E: Exactly right.

K: So that would be virtually the same thing. So as it's printed, it won't be in the memory anymore?

E: That's right. Let's say at this point Mark's stuff has been printed, and scratched off and Annie's stuff and my stuff are waiting to be printed. Then Glenn down the hall, has something to be printed, where would the computer put that request, in order to keep track in just the same way Mary Chung kept track?

K: Well, it would have to have a program that would put that at the end of the list.

E: (*Behavioral Question:*) So what...describes how Mark got to be first?

K: First come first serve.

---

After confirming that the computer could store a list, K correctly described the mechanism of the queue, by mapping the steps that pull requests off the front of the list and add new ones at the end (statements 2 and 4 respectively). K's description of this mechanism is a clearcut example of a causal analogical mapping from an explicitly described base domain causal mechanism. It can be contrasted with subject L<sup>10</sup>, who received the behavioral version of the analogy and was asked to describe the queue's causal mechanism.

---

E: Frequently at Mary Chung's there are people waiting to be seated. Mary keeps track of who to see in such a way that the first person to be seated next is always the person who among those currently waiting came in first. Ok, so that's the analogy is that the next guy who is going to get a seat is the one who ..

L: ... came in the longest period of time ago.

E: Yes, ok, now, in sending files to the printer the same situation often occurs.. if you are not working on a personal computer. It is that several files need to be printed, but only one can be printed at a time. And in this case, the computer resolves the problem in the same way that Mary does.

E: (*Causal question:*) Describe what you would do if you were the computer keeping track of the print requests and you have a bunch of requests at a particular time, and you can only print one at once, and they came in at different times. Ok, so can you describe to me how you would make this decision using the Mary Chung analogy.

L: I would give them all a number. And then whoever has the lowest number gets seated next. Then I take that number and throw it out. And so I don't have to worry about ever having to start over again at number 1 and seating the last person who came

---

<sup>10</sup>L is a theatre technician who does some programming as part of her technical work. L has approximately 2 years course-work experience with programming in Pascal and LISP.

in out of order, because I am juggling my numbers as I go.  
... I also used to work in a restaurants a lot.

---

Although we see in statement 2 that L clearly understood the *function* of a queue and her description of a mechanism satisfied that function, L's response to the causal question is interesting because it draws heavily on her knowledge of a *base* domain mechanism rather than a *target* domain mechanism. L chose to map a causal model by analogy from the base domain, despite the fact that she was *not* explicitly given a causal model of that domain and, could have generated a correct solution directly, rather than by analogy.<sup>11</sup>

L's generation of a causal model by analogy, using her knowledge of the restaurant (base) domain, rather than relying directly on her knowledge of programming suggests that domain-specific solutions to problems are not always preferred. Furthermore, the causal model that L retrieved was from a base domain situation that was *not* the one provided by the experimenter, but was instead familiar as a result of having worked in restaurants. This suggests that the base domain retrieval process, as well as the mapping process is quite complex. We suggest that it may be guided by a combination of factors: the *type of explanation* required; the problem solving goal; and the *currently* salient attributes of the base domain model presented during learning. Holyoak (Holyoak, 1985) has also suggested a model like this, although his model would not predict that an analogical solution might *dominate* an existing target domain solution. While our model also cannot account precisely for L's preference of a base-domain solution to the problem, the protocol clearly suggests that a full theory of analogical reasoning will need to include a fairly sophisticated account of the use of base domain features in the retrieval stage of the analogical reasoning process.

## 5. Summary

We have presented some preliminary protocol evidence of the role of *explanatory model types* in the retrieval and mapping stages of the analogical reasoning process. Our results suggest that:

- Subjects can generate and use coherent partial models.
- Behavioral and causal models can be related by shared initial and goal states.
- Retrieval of base domain solutions can be influenced by the type of explanation required, the problem solving goal, and the salient attributes of the base domain model presented during learning.

The evidence we are collecting is being used to develop a theory of constrained, purpose-directed analogical retrieval and mapping, and a theory of the process of *integrating multiple partial models* in learning about a new domain.

---

<sup>11</sup>Evidence that L could describe a purely target domain solution to the problem was gathered in a follow-up session several weeks after this protocol was collected. In that session, L was asked to describe an implementation for a queue directly rather than by analogy. She first repeated her earlier answer, but when she was asked if the numbers were necessary, she responded that the ordering of the list was sufficient, and proceeded to describe the set of steps involved in using an ordered list to implement the queue; adding new entries to the end and removing the "next" item from the front.

## References

- Adelson, Beth. When novices surpass experts: How the difficulty of a task may increase with expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition*, July 1984.
- Burstein, Mark H. Concept Formation by Incremental Analogical Reasoning and Debugging. In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Ed.), *Proceedings of the International Machine Learning Workshop*. Champaign-Urbana, IL: University of Illinois, 1983. Also appears in *Machine Learning: Volume II*. pp. 351-370. Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1986.
- Burstein, Mark H. *Learning by Reasoning from Multiple Analogies*. Doctoral dissertation, Yale University, 1985.
- Burstein, Mark H. Concept Formation by Incremental Analogical Reasoning and Debugging. In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Ed.), *Machine Learning: Volume II*. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1986. Also appeared in the *Proceedings of the Second International Machine Learning Workshop*, Champaign-Urbana, IL., 1983.
- Carbonell, Jaime G. Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Ed.), *Machine Learning: Volume II*. Los Altos, CA: Morgan Kaufman Publishers, Inc., 1986.
- Collins, Allan. Component Models of Physical Systems. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*. Cognitive Science Society, 1985.
- Collins, Allan and Gentner, Dedre. Constructing Runnable Mental Models. In *Proceedings of the Fourth Annual Conference of the Cognitive Science Society*. Boulder, CO: Cognitive Science Society, 1982.
- Collins, Allan and Gentner, Dedre. Multiple Models of Evaporation Processes. In *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*. Rochester, NY: Cognitive Science Society, 1983.
- Collins, Allan and Gentner, Dedre. How People Construct Mental Models. In N. Quinn and D. Holland (Eds.), *Cultural Models in Thought and Language*. Cambridge, UK: Cambridge University Press, 1987. In press.
- Coulson, R., Feltovich, P., and Spiro, R. *Foundations of a Misunderstanding of the Ultrastructural Basis of Myocardial Failure: A Reciprocating Network of Oversimplifications* (Tech. Rep. 1). Southern Illinois University School of Medicine, Conceptual Knowledge Research Project, August 1986.
- Gentner, Dedre. Structure-Mapping: A theoretical framework for analogy. *Cognitive Science*, 1983, 7(2), 155-170.
- Holyoak, K. J. The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation*. New York, NY: Academic Press, 1985.
- Kedar-Cabelli, Smadar. *Analogy from a Unified Perspective* (Tech. Rep. ML-TR-3). Laboratory for Computer Science Research, Rutgers University., November 1985.
- Schank, R.C. *Dynamic Memory: A theory of learning in computers and people*. Cambridge University Press, 1982.
- Sternberg, R. J. and Adelson, B. Changes in Cognitive Structure via Metaphor. In *Annual Meeting of the Psychonomic Society*. San Antonio, Texas: , 1978.
- Stevens, A. and Collins, A. Multiple Conceptual Models of a Complex System. In Snow, R. E., Federico, P. and Montague, W. E. (Eds.), *Aptitude, Learning, and Instruction*. Hillsdale, N.J.: Erlbaum, 1980.
- Stevens, A. and Steinberg, C. *A Typology of Explanations and Its Application to Intelligent Computer Aided Instruction* (Tech. Rep. 4626). Bolt Beranek and Newman Inc., March 1981.

- Thagard, Paul and Holyoak, Keith. Discovering the Wave Theory of Sound: Inductive Inference in the Context of Problem Solving. In *Proceedings of the Ninth IJCAI*. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1985.
- Weld, Daniel S. *Explaining Complex Engineered Devices* (Tech. Rep. 5489). Bolt Beranek and Newman Inc., November 1983.
- Williams, M., Hollan, J., and Stevens, A. Human reasoning about a simple physical system. In Gentner, D. and Stevens, A. (Eds.), *Mental Models*. Erlbaum, 1981.
- Winston, P. H., Binford, T. O. Learning Physical Descriptions from Functional Definitions, Examples, and Precedents. In *Proceedings of AAAI-83*. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1983.