

---

## Revising Domain Knowledge with Cross-Domain Analogy

---

**Scott E. Friedman**

SFRIEDMAN@SIFT.NET

SIFT, LLC, 211 North First Street, Suite 300, Minneapolis, MN 55401 USA

**David M. Barbella**

BARBELLA@U.NORTHWESTERN.EDU

**Kenneth D. Forbus**

FORBUS@NORTHWESTERN.EDU

Qualitative Reasoning Group, Northwestern University, 2133 Sheridan Road, Evanston, IL 60208 USA

### Abstract

AI systems and human novices share a difficult problem: repairing incorrect models to improve expertise. For people, the use of analogies during instruction can augment the repair of science knowledge. Enabling AI systems to do the same involves several challenges: representing knowledge in commonsense science domains, constructing analogies to transfer knowledge, and flexibly revising domain knowledge. We address these issues by using qualitative models for representing knowledge, the Structure-Mapping Engine for analogical mapping, and a computational model of conceptual change for revising knowledge. In our simulation trials, we initialize the system with one of several student misconceptions of the day/night cycle from the cognitive science literature. The system automatically repairs these misconceptions using an analogy, expressed using natural language, by: (1) validating analogical correspondences via user feedback; (2) transferring knowledge from the base domain, and (3) constructing new explanations to repair misconceptions.

### 1. Introduction

Repairing incorrect beliefs about the world is a common task during formal education. Research suggests that people more effectively repair incorrect beliefs in commonsense science domains when they are presented with an analogy (Vosniadou et al., 2007; Brown & Clement, 1989). Unfortunately, our AI systems lack the flexibility of human students when revising commonsense science beliefs with analogy, due to four general challenges:

1. Representing imprecise and incomplete knowledge in commonsense science domains.
2. Creating a mapping between a base and target domains, when the latter has misconceptions.
3. Transferring knowledge to the target domain.
4. Revising beliefs in the target domain to accommodate transferred knowledge.

This paper presents a model that addresses these challenges. Our claims are as follows:

1. Qualitative process theory (Forbus, 1984) can represent the commonsense science knowledge of novice students.
2. The Structure-Mapping Engine (Falkenhainer et al., 1989) can compute analogical mappings and transfer knowledge across domains, even in the presence of misconceptions.
3. A computational model of explanation-based conceptual change (Friedman, 2012) can evaluate and selectively incorporate knowledge transferred via analogy.

We implemented our model using the Companions cognitive architecture (Forbus et al., 2009). We simulate four student misconceptions about the day/night cycle from the cognitive science literature. For each misconception, we use the analogy from Vosniadou et al. (2007): that the day/night cycle is like the cooking of gyros (a meat roasting vertical rotisserie, commonly found in Greek restaurants). Using formal models of the base (gyros) and target (day/night) domains, and hints about correspondences provided via natural language, the system automatically (1) computes a mapping between the base and target, (2) interactively checks its analogical correspondences with a human teacher, (3) transfers beliefs, and (4) revises its knowledge accordingly.

We begin with an overview of novice mental models of the day/night cycle and how analogies have facilitated student learning in this domain. Then we summarize relevant background and describe how our model works. The simulation experiment and its results are described next, ending with related and future work.

Children are not blank slates when they first encounter explanations of the day/night cycle in formal education. Many students use mental models that are flawed, or that are internally consistent but scientifically incorrect (Sadler, 1987; Baxter, 1989; Vosniadou & Brewer, 1994). Some of these flawed mental models, as well as the correct model, are illustrated in Figure 1. These mental models used by students describe processes in the world, including the process of the moon or clouds moving to block out the sun, the sun turning off or on, the sun moving and




Figure 1. Mental models of day and night, including simplified correct model (E). Superscripts indicate studies in which mental models were used by students: (1) Sadler (1987); (2) Baxter (1989); (3) Vosniadou and Brewer (1994).

Table 1. Analogy given to students. All students saw the non-italicized text; only students in the analogy condition saw the italicized text.

---

The Earth is round. Day changes to night because the Earth turns around its axis. *The earth is moving just like gyros turns around on the vertical spit while roasting.* The sunlight reaches only one side of the Earth and on this side it's day. *In the same way, the fire cooks only one side of gyros, the one that is turned towards it.* On the other side of the Earth, which is not reached by the sunlight, it is night. As the Earth turns around its axis, the side where there was night turns slowly towards the Sun and it becomes day. *In the same way, as the gyros turns around the spit, the side that was not cooked before gradually turns towards the fire and starts roasting.* On the contrary, the side, where there was day, slowly turns away from the Sun and it becomes night. The Moon is not responsible for the day and night cycle. Thus, it changes from day to night because the Earth turns around its axis, and the sunlight shines on a different side of Earth.

---

entering the earth, or the earth rotating. They make different assumptions about the world, they describe different types of processes, and they even make different predictions of how changes in the system (e.g., the absence of a moon) would affect day and night.

How, then, can students repair these flawed mental models, and what cognitive processes are responsible? Vosniadou et al. (2007) hypothesized that a scientific explanation that includes an analogy would help students repair their mental models more than an analogy-free explanation. Their experiment on 111 Greek third and fifth graders provides evidence for this hypothesis. They used an analogy between gyros and the Earth, shown in Table 1. Students given the analogy were more likely to change their explanations to a scientific (or partially-scientific) model than students who were not given the analogy.

## 2. Background

Our model uses qualitative process theory for knowledge representation and reasoning, model formulation to assemble this knowledge into explanations, and analogical reasoning to transfer this knowledge across domains for reuse. We discuss each in turn.

### 2.1 Compositional Modeling and Qualitative Process Theory

Simulating human reasoning about dynamic systems makes several demands on knowledge representation. First, it must be capable of representing incomplete and incorrect domain knowledge. Second, it must represent processes (e.g., rotation on an axis, translation along a path) and qualitative proportionalities (e.g., the closer a light source, the brighter the light). Our system meets these demands by using qualitative process theory (Forbus, 1984) for representing domain knowledge and reasoning about physical processes.

We use *model fragments* (Falkenhainer & Forbus, 1991) to formally represent conceptual entities, processes, and views (Forbus, 1984) that describe multiple entities. A model fragment called `RemoteLighting` is shown in Table 2. It has two participant slots: (1) a source of type `LightSource` which fills the `providerOf` role<sup>1</sup> and (2) a lit-object of type `SpatialThing`

---

<sup>1</sup> We use the OpenCyc ontology ([www.opencyc.org](http://www.opencyc.org)) as a source of everyday concepts and relationships.

Table 2. A model fragment that describes a relationship between a light source and a lit object.

---

```

ModelFragment RemoteLighting
Participants:
  ?source LightSource providerOf
  ?lit-object SpatialThing recipientOf
Constraints:
  (spatiallyDisjoint ?source ?lit-object)
Conditions:
  (greaterThan (Light ?source) 0)
Consequences:
  (qprop
    (Light ?lit-object)
    (Light ?source))
  (positivelyDependsOn
    (Light ?lit-object)
    (visibleFrom ?lit-object ?source))

```

---

which fills the `recipientOf` role. These participants are constrained to be `spatiallyDisjoint` in order for a `RemoteLighting` instance to exist over any given source and `lit-object`. Thus, the constraints describe the instantiation criteria for a model fragment. According to the condition, if the light of source is greater than zero, then the `RemoteLighting` consequences are inferred. Thus, the conditions describe the behavioral criteria of the model fragment.

The consequence statements of `RemoteLighting` include two influences: (1) the light of `lit-object` is qualitatively proportional to that of the source, and (2) the light of `lit-object` is positively affected if `lit-object` is `visibleFrom` the source. The qualitative proportionality is also known as an *indirect influence* in qualitative process theory. Indirect influences describe monotonically increasing or decreasing relationships between quantities. The second influence is between a proposition and a quantity, and it asserts that the quantity will be greater in a state where the proposition holds than in a state where the proposition does not, all else equal. Once these consequences are asserted, they describe how the light of an object might change, e.g., its light source brightens or dims, or something occludes it.

## 2.2 Model Formulation

Model fragments and QP theory provide formalisms for representing and organizing a domain theory as reusable pieces. *Model formulation* is the process of assembling these reusable pieces into a qualitative model that has predictive and explanatory power. In this work, we use the abductive model formulation algorithm of (Friedman, 2012). The algorithm is given (1) a set of model fragments called a *domain theory*, (2) a set of propositional statements (e.g., about astronomy) called a *scenario*, and (3) a proposition to justify, such as:

```

(greaterThan
  (MeasurementAtFn (Light Chicago) Day)
  (MeasurementAtFn (Light Chicago) Night))

```




Figure 2. Subset of justification structure resulting from model formulation. Circles are propositional beliefs and triangles are justifications that associate antecedent beliefs at left with consequences at right. Dashed lines indicate additional structure not shown here.

The algorithm chains backwards by finding rules and model fragments that can infer the statement, binding participant slots and antecedent variables accordingly, and repeating on the antecedent statements recursively, as in Rickel & Porter (1997). If conditions of model fragments are not already known, they are assumed.

The output is a network of justifications and assertions, as shown in Figure 2. The justified belief is shown at the right, and the supporting justification structure extends to the left. In this case, the system inferred that the light of Chicago positively depends on whether the sun is visible from Chicago, and then inferred that this must be the case during the day and not the night. The change in visibility is explained by the sun entering and leaving a container, as in Figure 1(D). The situations of entering and leaving a container are other model fragments that are instantiated and activated just like `RemoteLighting` above. The assertions (circular belief nodes) in Figure 2 describe model fragment instances that might be active in various scenario states. Some of the assertions are assumptions that have no justification, but must hold for other assertions to be true.

### 2.3 The Structure-Mapping Engine

The Structure-Mapping Engine (SME) (Falkenhainer et al., 1989) is a computational model of Gentner's (1983) psychological theory of analogy and similarity. SME takes two cases as inputs: a *base*, such as a description of gyros cooking on a rotisserie, and a *target*, such as a (potentially) flawed description of the day/night cycle. SME automatically produces one or more *mappings* between the base and target, each of which consists of: (1) *correspondences* between base and target items; (2) a numerical *similarity score* that describes the quality of the mapping; and (3) *candidate inferences*, which are statements that might hold in the target case based on its correspondences to the base case. Candidate inferences may not be deductively sound; they are hypotheses that are supported by the comparison. Candidate inferences are how knowledge is transferred across domains by analogy.

### 3. Learning by Analogy and Explanation

Just as people learn from analogies across domains (Vosniadou et al., 2007), our system repairs its knowledge base by (1) transferring knowledge via analogy and then (2) using transferred knowledge to construct new qualitative models. Here we describe our approach, beginning with how knowledge is organized.

#### 3.1 Organizing Knowledge by Explanations

Our system records the beliefs and justifications of the best explanations for its beliefs. The explanation for a belief  $b$ , e.g.,  $b = \text{“Chicago is brighter in day than at night,“}$  is computed in three steps:

1. Justifying  $b$  via model formulation (e.g., Figure 2).
2. Reifying all *minimal explanations* of  $b$  (sequences of justifications that fully justify  $b$  without redundancy). There may be many, depending on the justification branching factor.
3. Computing the numerical cost of each explanation and retaining the minimal-cost explanation as the *preferred explanation* for  $b$ .

The preferred explanation  $x$  for a belief  $b$  acts as an index into memory. All of the beliefs and model fragment instances in  $x$  are now believed by the system, whereas beliefs and model fragments in non-preferred explanations of  $b$  are not necessarily believed by the system. This is important, because if the system switches its preferred explanation for  $b$  to a different explanation  $x'$  this could cause a belief revision. All beliefs in  $x'$  would then be believed and beliefs exclusive to  $x$  would no longer necessarily be believed.

Since beliefs are indexed using preferred explanations, the system revises its beliefs when either (1) new information yields a new, lower-cost explanation or (2) new information changes the cost ordering of existing explanations.

Certain artifacts within an explanation incur a cost. These artifacts and their associated costs include:

- **Logical contradictions** (cost = 100).
- **Assumptions** (cost =  $-\log(P(b))$ ) are penalized based on the probability  $P(b)$  of the belief  $b$ , where  $b$  is a belief that is not supported by instruction, observation, or antecedent beliefs.
- **Model fragment instances** (cost = 2) are penalized to minimize the number of causal factors.
- **Model fragments** (cost = 4) are penalized to minimize the types of causal factors.
- **Assumed quantity/state changes** (cost = 40) are changes within a system that cannot be explained.

Importantly, an artifact’s cost is only incurred once. For example, if the model fragment `RemoteLighting` is present in a previous preferred explanation, then a new explanation that uses it will incur zero cost for that artifact. This means that the system is biased to reuse existing models, model types, and assumptions. This promotes simplicity within and across explanations.

### 3.2 Finding and Handling Contradictions

Some artifacts, such as contradictions, involve more than one belief. For a contradiction cost to be incurred, each of the beliefs comprising that contradiction must be in a preferred minimal explanation or within the *adopted domain knowledge* context. Adopted domain knowledge consists of beliefs learned from instruction and observation. They might not be used in preferred explanations, but they constrain the space of preferred explanations because they influence the cost computation.

To illustrate, suppose that the system explains that the sun enters the earth to cause the night, and exits the earth to cause the day. This is illustrated in Figure 1(D) and its justifications structure is partially plotted in Figure 2. Suppose that this explanation  $x$  is the system’s currently-preferred explanation for this phenomenon, and that it contains the following statement:

```
(soleCauseOf-Type
  (not (visibleFrom Chicago TheSun))
  EnteringContainer)
```

This statement asserts that `EnteringContainer` is the only type of event that disables the Sun being visible from Chicago, with a similar statement connecting `ExitingContainer` with enabling Chicago’s sun-visibility.

Next, suppose that we instruct the system that Chicago’s sun-visibility is actually enabled and disabled by rotation processes (which we will accomplish using analogy, as described below). The following belief is added to adopted domain knowledge:

```
(soleCauseOf-Type
  (not (visibleFrom Chicago TheSun))
  Occluding-Rotation)
```

This contradicts the current preferred explanation. Contradictions are detected using domain-general rules, so the following contradiction will be incurred with a cost of 100:

```
C = {(soleCauseOf-Type
      (not (visibleFrom Chicago TheSun))
      EnteringContainer),
      (soleCauseOf-Type
      (not (visibleFrom Chicago TheSun))
      Occluding-Rotation)}
```

The system might be able to reduce cost by revising its explanation preference or construct a new explanation altogether to replace this one.

### 3.3 Transferring Domain Knowledge via Analogy

Suppose that the agent holds the beliefs that the sun is contained in the earth at night, and then it exits to cause the day. If the system does not have information to the contrary, this will be the lowest-cost, preferred explanation for the day/night cycle. An analogy is one way to provide new information to cause the system to change its model.

The analogy given to students in the Vosniadou et al. (2007) study is reproduced in Table 1. While our natural language system is capable of processing some complete analogies (Barbella & Forbus, 2011), the language system and the compositional models currently use different vocabularies. Since we needed to construct predicate calculus representations representing the different misconceptions in any case, to factor out natural language understanding errors, we also created a predicate calculus model of the gyros situation, but left out statements such as “The moon is not responsible...”, which are implications of the analogy. We did use natural language to express the key correspondences of the analogy, which were used to constrain the comparison via SME, using natural language, specifically, “The fire is to gyros, as the sun is to the earth.”

Statements such as these are interpreted as *analogical dialogue acts* (Barbella & Forbus, 2011) that establish correspondences between the entities of the domain. Specifically, the system inferred two required correspondence constraints that were used by SME:

```
(requiredCorrespondence Gyros PlanetEarth)
(requiredCorrespondence Fire TheSun)
```

The cases to be compared were constructed via *dynamic case constructors* (Mostek et al., 2000) that search preferred explanations and adopted domain knowledge for statements that mention Gyros or Fire (for the base case) and statements that mention PlanetEarth or TheSun (for the target case).

Given that the student’s understanding of the target domain can be quite far from correct, it is entirely possible that a student might not understand the analogy. An analogy can involve many correspondences and candidate inferences, some of which might be quite hard for students to articulate, and in any case would be far more than is reasonable to ask about in conversation. Consequently, the Companion was programmed to only ask about entity correspondences. For each entity correspondence in the best mapping produced by SME, it produces a multiple choice question for its human teacher. Here are two examples:

**Companion:** I believe *Cooking-State* corresponds to *Day*.

- a. That is correct.
- b. Those don't match.
- c. No, *Cooking-State* does not belong in the mapping.
- d. No, *Day-State* does not belong in the mapping.
- e. No, neither belongs in the mapping.

**User:** a

**Companion:** I believe *Gyros rotating to occlude GyroSide-Cooking from Fire* corresponds to *TheSun entering PlanetEarth so it is occluded from Chicago*.

- a. That is correct.
- b. Those don't match.
- c. No, *Gyros rotating to occlude GyroSide-Cooking from Fire* does not belong in the mapping.
- d. No, *TheSun entering PlanetEarth so it is occluded from Chicago* does not belong in the mapping.
- e. No, neither belongs in the mapping.

**User:** d



The underlying representation is Davidsonian (i.e., events are reified), which is why occlusion events can be asked about. After all of the entity correspondences have been queried (for this analogy, there are ten), the Companion computes an updated list of SME constraints from the user’s answers. Based on the second answer (d) shown above, the Companion will not allow anything to correspond to the model fragment instance of the sun entering the earth. If any new constraints were identified, a new SME mapping is computed.

The adopted domain knowledge is extended using the candidate inferences of the best mapping. For the mapping in this example, thirty inferences are recorded, including:

```
(isa Skolem-rotation Rotation-Periodic)
(objectRotating Skolem-rotation PlanetEarth)
(soleCauseOf-Type
  (not (visibleFrom Chicago TheSun))
  Occluding-Rotation)
(holdsIn Night (between Chicago TheSun PlanetEarth))
(ingredients PlanetEarth Skolem-Beef)
(ingredients PlanetEarth Skolem-Lamb)
```

Some of these beliefs include terms with prefix *Skolem-*. These represent new entities and processes that did not previously exist in the target domain. The inferences in adopted domain knowledge affect learning in two ways: (1) some beliefs, such as those describing the earth rotating, permit the construction of new explanations; and (2) some beliefs contradict beliefs in preferred explanations, which raises their cost and biases the system to change its explanation preference and revise its beliefs. We discuss this next.

### 3.4 Cost-Based Belief Revision

The system previously believed that the mechanism preventing visibility was the sun entering the earth, but has just made the analogical inference that a rotation process is responsible. These beliefs are contradictory. When a contradiction is detected, the system attempts to reduce cost by changing its explanation preferences. We use a greedy revision algorithm (Friedman, 2012) that takes a contradiction as input and then (1) finds all explanations that contain a belief from the contradiction and (2) for each explanation, it greedily revises the preference to a new explanation if it will reduce cost. This is not guaranteed to find the globally minimum cost, but it is efficient and guaranteed to terminate.

## 4. Simulation

We conducted six simulation trials. On each trial, the system is initialized with (1) domain knowledge of how gyros rotate on a spit and cook one side at a time and (2) one of six possible starting models of the day/night cycle: a misconception in Figure 1 (A-D); the correct model in Figure 1(E); or no model, where the system is a blank slate. We do not build the system’s explanation by hand; it automatically constructs this from its domain theory.

For each starting model, we record the number of incorrect analogical correspondences corrected by the user, whether the correct model is reached without user assistance (i.e., without answering correspondence queries), and whether the correct model is reached with user assistance. For the unassisted trials, we remove all constraints, even those from the original

Table 3. Simulation results for six starting models. SME errors are incorrect correspondences between base entities and target entities.

Starting model	# SME errors	Correct unassisted	Correct assisted
Moon blocks sun at night	4	No	Yes
Clouds block sun at night	4	No	Yes
Sun shuts off at night	2	Yes	Yes
Sun enters earth at night	3	Yes	Yes
Correct (rotation) model	0	Yes	Yes
No model	0	Yes	Yes

analogical dialogue act, to determine whether the misconception interferes with unconstrained analogy.

The results are shown in Table 3. When there were no misconceptions (e.g., in the “correct model” or “no model” conditions), there were no incorrect correspondences. Conversely, when misconceptions exist in the target domain, incorrect correspondences are computed, and the number of incorrect correspondences varies with the misconception. Without instructor intervention, incorrect correspondences confuse the transfer of knowledge and hinder belief revision. For example, in the “Moon blocks sun” example, one of the incorrect correspondences includes `Gyros` corresponding to `MoonOfEarth`. This causes the inference that Chicago is on the moon, and ultimately leads the system to revise its beliefs, concluding that the moon rotates to cause Chicago’s day/night cycle.

## 5. Related Work

Two earlier systems bear close relation to the approach we have reported here. One of these, CARL (Burstein, 1986) was a computer model of analogical learning that learned from multiple analogies presented by a teacher. It modeled students who were tutored on the programming language BASIC. CARL coordinated multiple analogies such as (1) variable assignment is like “putting things in boxes” and (2) variable assignment is like algebraic equality. It constructed analogical hypotheses and then revised them when they did not hold in the target domain. Unlike CARL, our system uses qualitative models, it evaluates new hypotheses and explanations using a cost function, and it interacts with the user to debug analogical correspondences (whereas CARL interacted to debug the hypotheses themselves). Since our simulation makes an average of 40 analogical inferences (i.e., hypotheses) in a single analogy, debugging all of them with the user is impractical.

ECHO (Thagard, 2000) is a connectionist model that uses constraint-satisfaction to judge hypotheses by their explanatory coherence. This system creates excitatory and inhibitory links between consistent and inconsistent propositions, respectively. Its “winner take all” network means that it cannot distinguish between there being no evidence for competing propositions versus balanced conflicting evidence for them. ECHO requires a full explanatory structure as its input. By contrast, our system generates its justification structure automatically and evaluates it using a cost function.

## 6. Discussion

We have described a system that (1) uses qualitative models to represent humanlike misconceptions in a commonsense science domain and then (2) interactively works through an analogy to repair its misconceptions. We tested our system using four misconceptions about the day/night cycle from the cognitive science literature, and showed that learning is improved by user interaction during analogical reasoning.

Our simulation results diverge from Vosniadou et al.'s (2007) study in at least one respect. Like their study, reading an analogy improved the chances of mental model repair, but our simulation is more ready to change its beliefs than the students were. This is especially interesting because the analogy given to the students has direct instructional information that was not given to our simulation (e.g. "The Moon is not responsible for the day and night cycle.") There could be several reasons for this. First, the students did not have the opportunity to actively query a teacher, the way the simulation did. Second, there are very likely individual differences between students in how hard they push an analogy, and in their strategies for coping with inconsistencies that they find. A student who sees that her current mapping implies that Chicago is on the moon, for example, might simply drop that inference instead of looking for a better mapping.

Analogies are pervasive in science texts and in conversation, so we plan to integrate this analogical learning mechanism into our Companions-based learning by reading work. The processes of analogical mapping, contradiction detection, and explanation construction are useful processes for offline processing and rumination (Forbus et al., 2007) to accommodate new information.

## Acknowledgements

This work was funded by the Office of Naval Research, under the Socio-Cognitive Architectures Program and the Intelligent & Autonomous Systems Program.

## References

- Barbella, D., & Forbus, K. (2011). Analogical dialogue acts: Supporting learning by reading analogies. *Proceedings of the Twenty-Sixth Annual AAAI Conference on Artificial Intelligence*. San Francisco, CA: AAAI Press.
- Baxter, J. (1989). Children's understanding of familiar astronomical events. *International Journal of Science Education*, *11*, 502-513.
- Brown, D., & Clement, J. (1989). Overcoming misconceptions via analogical reasoning: Factors influencing understanding in a teaching experiment. *Instructional Science*, *18*, 237-261.
- Burstein, M. H. (1986). Concept formation by incremental analogical reasoning and debugging. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2). San Francisco: Morgan Kaufmann Publishers, Inc.
- Falkenhainer, B., Forbus, K., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1-63.

- Falkenhainer, B., & Forbus, K. (1991). Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51, 95-143.
- Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.
- Forbus, K., Riesbeck, C., Birnbaum, L., Livingston, K., Sharma, A., & Ureel, L. (2007). Integrating natural language, knowledge representation and reasoning, and analogical processing to learn by reading. *Proceedings of the Twenty-Second Annual AAAI Conference on Artificial Intelligence*. Vancouver, BC: AAAI Press.
- Forbus, K., Klenk, M., & Hinrichs, T. (2009). Companion cognitive systems: Design goals and lessons learned so far. *IEEE Intelligent Systems*, 24, 36-46.
- Friedman, S. (2012). *Computational conceptual change: An explanation-based approach*. Doctoral dissertation, Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, IL.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Klenk, M., & Forbus, K. (2009). Domain transfer via cross-domain analogy. *Cognitive Systems Research*, 10, 240-250.
- Mostek, T., Forbus, K., & Meverden, C. (2000). Dynamic case creation and expansion for analogical reasoning. *Proceedings of the Fifteenth Annual AAAI Conference on Artificial Intelligence*. Austin, TX: AAAI Press.
- Rickel, J., & Porter, B. (1997). Automated modeling of complex systems to answer prediction questions. *Artificial Intelligence*, 93, 201-260.
- Sadler, P. M. (1987). Misconceptions in astronomy. *Proceedings of the Second International Seminar on Misconceptions and Educational Strategies in Science and Mathematics*, 3, 422-425. Ithaca, NY.
- Thagard, P. (2000). Probabilistic networks and explanatory coherence. *Cognitive Science Quarterly*, 1, 93-116.
- Tomai, E., & Forbus, K. (2009). EA NLU: Practical language understanding for cognitive modeling. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference*. Sanibel Island, FL.
- Vosniadou, S., & Brewer, W. F. (1994). Mental models of the day/night cycle. *Cognitive Science*, 18, 123-183.
- Vosniadou, S., Skopeliti, I., & Gerakaki, S. L. (2007). Understanding the role of analogies in restructuring processes. In A. Schwering, U. Krumnack, K.U. Kuhnberger, & H. Gust (Eds.), *Analogies: Integrating multiple cognitive abilities*. Osnabruck, Germany: Institute of Cognitive Science.