

Incremental Structure-Mapping

Kenneth D. Forbus

Ronald W. Ferguson

The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL, 60201, USA
forbus@ils.nwu.edu
ferguson@ils.nwu.edu

Dedre Gentner

Psychology Department
Northwestern University
2029 Sheridan Road
Evanston, IL, 60208, USA
gentner@nwu.edu

Abstract

Many cognitive tasks involving analogy, such as understanding metaphors, problem-solving, and learning, require the ability to extend mappings as new information is found. This paper describes a new version of SME, called I-SME, that operates incrementally. I-SME is inspired by Keane's IAM model and the use of incremental mapping in Falkenhainer's PHINEAS learning system. We describe the I-SME algorithm and discuss tradeoffs introduced by incremental mapping, including parallel versus serial processing and pragmatic influences. The utility of I-SME is illustrated by two examples. First, we show that I-SME can account for the psychological results found by Keane on a serial version of the Holyoak & Thagard attribute mapping task. Second, we describe how I-SME is used in the Minimal Analogical Reasoning System (MARS), which uses analogy to solve engineering thermodynamics problems.

1. Introduction

Many cognitive tasks involving analogy, including metaphor understanding, problem solving, and learning require the ability to process information incrementally. In metaphor understanding, readers often build up correspondences across several sentences (Boronat & Gentner, in preparation). In problem solving, students using a worked example to solve a related novel problem may go back and forth between them, seeking additional ways to interpret the new problem in light of the old. In conceptual change, new data can lead to analogies being modified or abandoned.

Burstein (1986) was the first to computationally model incremental processing in analogical learning.

Falkenhainer's (1987, 1990) PHINEAS demonstrated that SME could be used in a *map/analyze* cycle to model the incremental use of analogy in discovering physical theories. The first general-purpose incremental analogical matcher was Keane's IAM (Keane & Bradshaw, 1988).

This paper presents I-SME, an incremental version of SME (Falkenhainer, Forbus, & Gentner, 1986; 1989) that simulates the computations proposed by Structure-Mapping theory (Gentner, 1983). I-SME has all the capabilities of SME, including the ability to generate novel candidate inferences. Unlike SME, I-SME can extend its existing interpretations when given new information about base or

target. Section 2 discusses psychological and computational issues raised by incremental mapping. Section 3 describes the I-SME algorithm. Section 4 illustrates its operation on two kinds of tasks: (1) an attribute-mapping task that shows I-SME can explain order effects found by Keane (in press), and (2) modeling student behavior in solving thermodynamics problems by referring to worked examples. Section 5 discusses other issues and future work.

Given propositional representations for a base and target,

1. **Local Match Construction:** For each pair of expressions in base and target, if their predicates are identical, create a match hypothesis (MH) expressing the possibility that they match. Corresponding arguments of items participating in an MH are also matched if either (1) their predicates are identical, (2) the predicates are functions, or (3) they are entities. This network of match hypotheses provides the grist for subsequent processing.
2. **Kernel Construction:** Find kernel mappings by starting at MH's which are not themselves the argument of any other. For each such MH_i, if it is structurally consistent, then it and every MH below it comprise a kernel. Otherwise, recurse on the MH's that are its arguments.
3. **Structural Evaluation:** Propagate scores through the network of MH's, using argument connections to pass activation downward. This "trickle-down" of scores implements the systematicity preference. The score of a mapping is the sum of the scores of its MH's.
4. **Merge:** Construct up to *n* global interpretations of a match by using a greedy algorithm to find structurally consistent combinations of kernels that maximize structural evaluation scores.

Figure 1: The SME Algorithm

2. Issues in incremental mapping

Serial vs Parallel: What parts of the matching process should be serial versus parallel? In SME (Figure 1), processing is essentially parallel within the first three

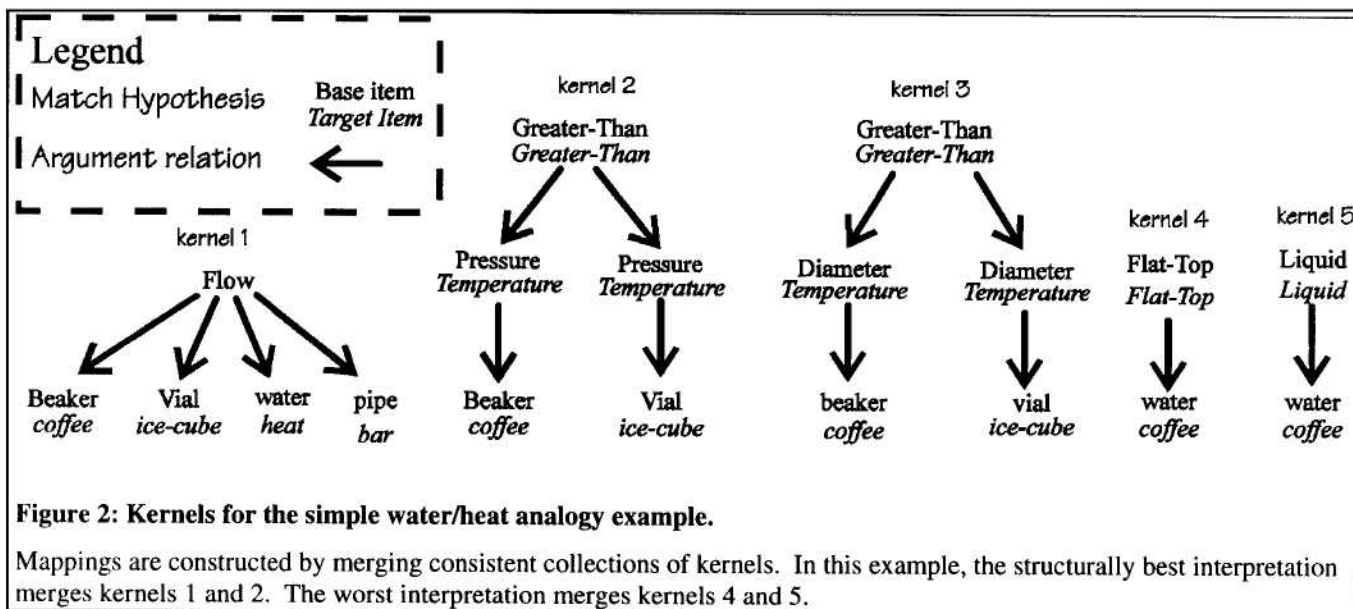


Figure 2: Kernels for the simple water/heat analogy example.

Mappings are constructed by merging consistent collections of kernels. In this example, the structurally best interpretation merges kernels 1 and 2. The worst interpretation merges kernels 4 and 5.

stages: only the final step of constructing global mappings is serial¹. By contrast, IAM is serial throughout: Even decisions about local matches are made sequentially, so that exploring alternate interpretations requires backtracking. SME avoids backtracking by creating, in parallel, networks representing all alternative local matches between items plus intermediate clusters (i.e., *kernel mappings*). We believe that a combination of initial parallel and later serial operation will best model human analogical processing. Some serial processing is essential: One cannot combine all information in parallel when some of it is not yet available. However, we believe the fully serial approach of IAM is unlikely to scale up to cognitively interesting representations. The natural place for serial processing is in the Merge step, because the kernels, which represent coherent, structurally-consistent collections of local matches, form a more appropriate unit of analysis for limited-resource serial processing than individual local matches themselves. (Figure 2 shows the kernel mappings for a simple example.) When base and target share large systematic structures, the number of kernels is small. Serial, capacity-limited merging of kernels could thus provide a plausible explanation for the "More is Less" phenomena (Gentner & Ratterman, 1991) where additional knowledge can improve both the rapidity and accuracy of mapping.

Pragmatics: Our original position (Gentner & Clement, 1988, Gentner, 1989) was that pragmatic concerns affected

only the pre- and post-processing of the mapping subsystem, with mapping being entirely based on structural consistency and relational identity matches. Others, such as Holyoak (1985) claimed "The distinction [between surface and structural similarities] thus crucially depends on the goal of the problem solver." Both positions have become less extreme. Holyoak & Thagard (1989) have incorporated structural consistency into ACME, and Hummel & Holyoak (1992) have pointed out problems in using pragmatics to influence the online mapping process. On our side, we have adopted for some purposes a technique called *pragmatic marking* (Forbus & Oblinger, 1990), to filter kernels that cannot yield novel candidate inferences relevant to a goal. Incremental matching is subject to misleading matches when only partial information is available, and early poor decisions can mislead subsequent processing. The use of this limited form of pragmatic information during mapping can help overcome these problems. However, we continue to believe that in order to capture the generativity of human analogical reasoning, the mapping process needs to have enough leeway to detect unanticipated structural correspondences. The incremental algorithm in this paper cycles between a largely autonomous mapping process and intervention by other goal-oriented processes.

3. The Incremental SME Algorithm

The Incremental SME algorithm (I-SME) accepts information about base and target descriptions to compare. Information can be presented to it all at once, one fact at a time, or in any combination desired. (When I-SME is given every item in base and target at the same time, its results are identical to those of SME.) I-SME maintains a small set of global mappings that represent its best interpretations of how the base and target compare, according to the information it has received. Each

¹ The fact that most SME implementations happen to run on serial machines is irrelevant here. ACME (Holyoak & Thagard, 1989), is also parallel in conception but is commonly implemented on serial machines, and the extraction of a mapping from its network is a serial process outside the simulation. Hummel & Holyoak (1992) explored incremental mapping, but their scheme does not handle higher-order relations, making it unsuitable for problem-solving.

EXTEND:

Given new base items $B1..Bi$ and new target items $T1..Tj$, and previous global mappings $M1..Mk$ (if any),

1. Extend set of local match hypotheses $\{Mhi\}$ by testing $B1..Bi$ against new and old target items, and by testing $T1..Tj$ against old and new base items. New elements of the match hypothesis network that violate the parallel connectivity constraint are marked as inconsistent, and *nogood* links are introduced to enforce the 1:1 constraint in merge steps.
2. Update the set of kernels. Starting with new root MH's (those which are not arguments of other MH's), search downward in parallel until the first structurally consistent MH is found. It and all of its descendants comprise a new kernel. New match hypotheses can introduce new root mappings or subsume old ones. Let K' be the set of new kernels.
3. Carry out structural evaluation on the new match hypotheses and K' . Filter K' via pragmatic constraints, if any.
4. Extend $M1..Mk$ by merging in kernels from K' whenever possible. (If no previous mappings exist, use the greedy algorithm to generate an initial set from the full collection of kernels K . This algorithm produces at most n mappings, all of which are within $p\%$ of the best in terms of their structural evaluation.) Keep at most n mappings.

REMAP:

1. Let K be the entire set of kernel mappings. Let K_F be the result of filtering K via pragmatic constraints, if any.
2. Clear the previous global mappings.
3. Use greedy algorithm to find new set of mappings starting from K_F , as in step 4 of EXTEND

Figure 3: Summary of the I-SME Algorithm

mapping (as in SME) has a structural evaluation score indicating the overall quality of the match and a set of candidate inferences representing surmises about the target domain justified by importing information from the base. Figure 3 summarizes the I-SME algorithm.

The I-SME algorithm follows the same basic sequence of the SME algorithm. The major features are (1) the datastructures and processing are carefully organized so as to preserve previously computed results that are still valid, rather than recomputing them, (2) the default operation is to *extend* a set of existing mappings, rather than to start from scratch, (3) the *remap* operation provides for backtracking at the level of kernels if the accumulated mappings become suboptimal, and (4) mappings are now first-class entities that can be referred to by systems using I-SME (e.g., a problem solver can have as a goal extending or verifying a particular mapping).

The intuition here is that normally people first try to incorporate new information into an ongoing mapping, but that they can reinterpret the analogy if necessary. I-SME does not automatically remap, since we believe remapping criteria tend to be task-specific. To help external systems make such decisions I-SME does provide an estimate of what fraction of the total possible structural evaluation the current mappings represent.

I-SME currently supports two filters on kernels.²:

- (REQUIRED $B_i T_i$): Exclude any root mapping that places another target item T_j in correspondence with B_i , or places another base item B_j in correspondence with T_i .
- (IDENTICAL-FUNCTIONS): Exclude any root mapping that places non-identical functions in correspondence.

The REQUIRED filter was used in PHINEAS; it enforces externally imposed correspondences (e.g., the speaker says that "heat is like water"). The IDENTICAL-FUNCTIONS filter represents a conservative strategy often used by subjects in an unfamiliar domain, when they may reject all but the most certain mappings. The utility of this constraint shown in Section 4.2.

I-SME's computational complexity is quite low. Most of its processing can be done in parallel, with bounds as reported for SME in (Falkenhainer, Forbus, & Gentner, 1989). The greedy merge algorithm is linear in the number of kernels, which in turn is between log (best case) and linear (worst case) in the size of the match hypothesis network. The match hypothesis network, in turn, is between linear (best case) and quadratic (worst case) in the size of the base and target descriptions. In other words, adding incremental operation does not result in any increase in computational complexity for I-SME with respect to SME. While we have not yet optimized the code, already I-SME is fast even on smallish computers.³

4. Examples

I-SME has been successfully tested with a corpus of examples from our previous work.⁴ Here we describe two computational experiments that exploit its incremental capabilities.

² Currently these constraints are implemented as hard filters, but we have experimented with "soft" implementations that allow their violation when doing so would yield an otherwise excellent match.

³ On a PowerBook 165c, small examples like Figure 2 take about 1/2 second, a typical qualitative physics example takes about two seconds, and matching stories involving 67 and 52 expressions respectively requires 34 seconds.

⁴ We have also used I-SME in a new version of our MAC/FAC retrieval model and in a simulation of symmetry and regularity detection (Ferguson, submitted).

(smart steve)	(hungry fido)
(tall bill)	(friendly rover)
(smart bill)	(hungry rover)
(tall tom)	(friendly blackie)
(timid tom)	(frisky blackie)

Singleton first condition (easy): Subjects mean time to solve = 178.0 seconds (Keane, in press)
Average # of I-SME remaps per example: 1.5

(timid tom)	(hungry fido)
(tall tom)	(frisky blackie)
(tall bill)	(friendly blackie)
(smart bill)	(hungry rover)
(smart steve)	(friendly rover)

Singleton last condition(hard): Subjects' mean time to solve = 363.1 seconds (Keane, in press)
Average # of I-SME remaps per example: 3.0

Figure 4: Examples of attribute mapping task

4.1 The attribute mapping task

Keane (in press) demonstrates that providing information incrementally can lead to order effects in analogical processing. Consider the pairs of statements at the top of Figure 4 (inspired by Holyoak & Thagard's (1989) attribute mapping task). If subjects are given each line, one at a time, they start by matching SMART with HUNGRY and STEVE with FIDO. In this example there is only one statement about FIDO and one statement about STEVE, and the hypotheses that suggest themselves as each new item is added lead to a solution where every entity and every attribute has a correspondence. By contrast, the pairs of statements on the bottom are more difficult: The natural initial guess, that TIMID goes with HUNGRY and TOM goes with FIDO, cannot lead to an interpretation in which everything has a unique match. Keane showed that subjects take, on the average, twice as long to do such problems when the singleton items are last as when they are first.

We have some concerns about this task as a test of analogical mapping, since it involves no higher-order structure, which is the hallmark of analogical processing. Nevertheless, the order effects Keane demonstrated seem likely to appear even with more natural stimuli, so we used I-SME with the condition that it remapped whenever an entity or attribute could not be added to its current mapping.⁵ As Figure 4 shows, when I-SME is executed on

the representations Keane used with IAM⁶, it produces the same kinds of order effects found in human subjects: The singleton-last condition, which on the average took subjects twice as long, resulted in I-SME performing on the average twice as many remaps.

Problem: Suppose helium enters a heater at 40° C and at a rate of 2 kg/sec. If the outlet temperature is 100° C, what is the heat transfer rate to the helium?

Worked Example: Suppose air enters a heater at 50° C and at a rate of 1 kg/sec. If the outlet temperature is 70° C, what is the heat transfer rate to the air?

Solution: Let h1 be the heater.

1. $h(\text{inlet}(h1)) = 324,253 \text{ kJ/kg}$ because (since air is an ideal gas), $h(\text{inlet}(h1)) = C_p(\text{air}) * T(\text{inlet}(h1))$ and

$$C_p(\text{air}) = 1003.414 \text{ J/kgC and}$$

$$T(\text{inlet}(h1)) = 323.15 \text{ K}$$

2. $h(\text{outlet}(h1)) = 344,321 \text{ kJ/kg}$ by the same argument, since

$$T(\text{outlet}(h1)) = 343.15 \text{ K}$$

3. $q(h1) = 20068 \text{ kJ/kg}$, because
 $h(\text{outlet}(h1)) = h(\text{inlet}(h1)) + q(h1)$

4. $Q(h1) = 20068 \text{ kW}$, because
 $q(h1) = Q(h1) / \text{mass-flow}(h1)$
 $\text{mass-flow}(h1) = \text{mass-flow}(\text{inlet}(h1))$
 $\text{mass-flow}(\text{inlet}(h1)) = 1.0 \text{ kg/sec}$

Figure 5: A typical problem and worked example given to MARS (summarized in English).

4.2 The Minimal Analogical Reasoning System (MARS)

We are using I-SME to create cognitive simulations of human problem-solving. For example, students often refer to worked examples to figure out how to solve new problems (Anderson & Thompson 1989; VanLehn, Jones, & Chi, 1992). MARS is designed to help explore this phenomena, providing an architecture whose processes and parameters should be adjustable to model different classes of student behaviors.

MARS currently solves problems in engineering thermodynamics. We chose that domain both because it is difficult and because we already have developed representations and tools for an intelligent learning environment for engineering thermodynamics (Forbus & Whalley, 1994), thus we can easily and automatically generate representations of problems and worked

⁵ Although we disagree with similarity tables as a modeling technique, to keep conditions between the IAM and I-SME and

ACME as close as possible we implemented a simple similarity table facility for this experiment.

⁶ We thank Mark Keane for providing his original representations.

examples.⁷ MARS receives as input two propositional representations, one representing the problem to be solved and the other representing a worked example. (Figure 5 summarizes in English a typical pair of inputs.) MARS has some qualitative knowledge about the kinds of processes that occur in components of thermodynamic cycles, a handful of numerical constants, and a simple equation solver. However, MARS currently has no built-in knowledge of thermodynamics equations. It begins by creating an initial match involving the objects and its qualitative understanding of them. MARS then begins a search for equations and numerical values that will allow it to find a numerical value for the goal quantity. To find equations, it uses the initial set of correspondences to identify information in the worked example that is potentially relevant. The initial mapping(s) are then extended with this new information, providing candidate inferences that are examined for usable information. Figure 6 shows MARS' solution (which is correct) for the problem of Figure 5. MARS extends its initial mapping five times, getting the equations it needs from the new candidate inferences sanctioned by the extended mapping. MARS uses the identical-functions filter constraint in extending its mappings to avoid misleading correspondences. In this example, for instance, removing that constraint allows MARS to match h (specific enthalpy) with T (temperature), leading to a simpler but incorrect solution. We believe that the appropriate use of mapping filters may be acquired as part of either domain expertise or general metacognitive skills.

1. Extend mapping to include $Q(h2)$.
2. From mapping: $q(h2)=Q(h2)/\text{mass-flow}(h2)$.
3. Extend mapping to include $\text{mass-flow}(h2)$.
4. From mapping: $\text{mass-flow}(\text{inlet}(h2))=\text{mass-flow}(h2)$.
5. Derived $\text{mass-flow}(h2)=2.0 \text{ kg/sec}$ from step 4 & givens.
6. Extend mapping to include $q(h2)$.
7. From mapping: $h(\text{outlet}(h2))=h(\text{inlet}(h2))+q(h2)$.
8. Extend mapping to include $h(\text{inlet}(h2))$.
9. From mapping: $h(\text{inlet}(h2))=C_p(\text{helium}) * T(\text{inlet}(h2))$.
10. From knowledge base, $C_p(\text{helium})=5176.85 \text{ J/kgC}$.
11. Derived $h(\text{inlet}(h2))=1621.13 \text{ kJ/kg}$, via steps 9 & 10 & givens.
12. Extend mapping to include $h(\text{outlet}(h2))$.
13. From mapping: $h(\text{outlet}(h2))=C_p(\text{helium}) * T(\text{outlet}(h2))$.
14. Derived $h(\text{outlet}(h2))=1931.74 \text{ kJ/kg}$ from steps 13, 10, and givens.
15. Derived $q(h2)=310.61 \text{ kJ/kg}$ from steps 7, 11, and 14.
16. Derived $Q(h2)=621.22 \text{ kW}$ from steps 2, 5, and 15.

Figure 6: A summary of MARS' solution to the problem

Currently, MARS does not extract any operators or control information from its worked example.⁸ This is one design choice that we believe should be varied to account for differences in individual students. Other such choices include how closely candidate inferences are scrutinized and the size and composition of the MARS' knowledge base. We are extending MARS to include a set of programmable options corresponding to what we believe will be theoretically important choices, and will explore its behavior through the space of these options to see how well it can account for variations in student performance.

5. Discussion

I-SME's technique of parallel computation of kernels and serial merging appears to be computationally effective and we believe that it is psychologically plausible. I-SME can account for the order phenomena discovered by Keane, while maintaining the efficiency of parallel operations for local processing. Furthermore, MARS demonstrates that I-SME can be used as a module in building sophisticated problem solvers that operate via analogy.

Modeling analogy requires a delicate interplay between goal-sensitive, strategic processes and processes of structural alignment and mapping that can run on their own. We think that the mix of parallel and serial processing in I-SME may best allow for planfulness and creativity.

6. Acknowledgments

We thank Mark Keane and John Everett for helpful feedback. This research was supported by the Cognitive Science Division of the Office of Naval Research.

⁷ Using representations designed and optimized for building a learning environment for engineering undergraduates, rather than specially designed for I-SME's use, provides a good test of I-SME's robustness.

⁸ Unlike for instance Carbonell's (1986) derivational analogy account, or analogical search control in Van Lehn & Jones (in press).

7. Bibliography

- Anderson, J.R. & Thompson, R. (1989). Use of analogy in a production system architecture. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. (pp 267-297). Cambridge: Cambridge University Press.
- Burstein, M. (1986). Concept formation by incremental analogical reasoning and debugging. In R.S. Michalski, J.G. Carbonell & J.M. Mitchell (Eds.), *Machine Learning II: An Artificial Intelligence Approach*. Los Altos, CA: Morgan-Kaufmann.
- Carbonell, J.G. (1986). Derivational Analogy: A theory of reconstructive problem solving and expertise acquisition. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell (Eds.), *Machine Learning: An AI Approach: Vol. 2* (pp.371-392). Los Altos, CA: Kaufmann.
- Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. In *Proceedings of IJCAI-87*. Los Altos: Morgan-Kaufmann.
- Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In Shrager, J. & Langley, P. (Eds.), *Computational Models of Scientific Discovery and Theory Formation*. San Mateo, CA: Morgan Kaufmann.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1986). The Structure-Mapping Engine. *Proceedings of AAAI-86*. Los Altos, CA: Morgan-Kaufmann.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1-63.
- Forbus, K.D. and Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. Hillsdale: Erlbaum.
- Forbus, K.D. and Whalley, P.B. (1994) Using qualitative physics to build articulate software for thermodynamics education. *Proceedings of AAAI-94*.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, pp 155-170.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). London: Cambridge University Press. Reprinted in *Knowledge acquisition and learning*, 1993, 673-694.
- Gentner, D. & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G.H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 22, pp. 307-358). New York: Academic Press.
- Gentner, D. and Ratterman, M.J. (1991). Language and the career of similarity. In S.A. Gelman & J.P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development*. (pp 225-277), London: Cambridge University Press.
- Holyoak, K.J. (1985). The pragmatics of analogical transfer. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol 1, pp.59-87). New York: Academic Press.
- Holyoak, K.J., and Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3), 295-355.
- Hummel, J.E., & Holyoak, K.J. (1992). Indirect analogical mapping. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp 516-521). Bloomington, Indiana.
- Keane, M.T., & Brayshaw, M. (1988). The Incremental Analogy Machine: A computational model of analogy. In D. Sleeman (Ed.), *Third European Working Session on Machine Learning*. London: Pitman/San Mateo, CA: Morgan Kaufmann.
- Keane, M.T., Ledgeway, T. & Duff, S. (in press) Constraints on analogical mapping: A comparison of three models. To appear in *Cognitive Science*.
- Van Lehn, K., Jones, R.M., and Chi, M.T. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2(1), pp. 1-59, Hillsdale, NJ: Erlbaum.
- Van Lehn, K., & Jones, R. (in press). Integration of analogical search control and explanation-based learning of correctness. In S. Minton & P. Langley (Eds.), *Machine learning methods for planning and scheduling*. Los Altos, CA: Kaufmann.