

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

MAC/FAC: A Model of Similarity-based Retrieval

#### **Permalink**

<https://escholarship.org/uc/item/1gb7n02g>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 13(0)

#### **Authors**

Gentner, Dedre

Forbus, Kenneth D.

#### **Publication Date**

1991

Peer reviewed

# MAC/FAC: A Model of Similarity-based Retrieval

Dedre Gentner and Kenneth D. Forbus

The Institute for the Learning Sciences, Northwestern University  
1890 Maple Avenue, Evanston, IL, 60201

## Abstract

We present a model of similarity-based retrieval which attempts to capture three psychological phenomena: (1) people are extremely good at judging similarity and analogy when given items to compare. (2) Superficial reminders are much more frequent than structural reminders. (3) People sometimes experience and use purely structural analogical reminders. Our model, called MAC/FAC (for “many are called but few are chosen”) consists of two stages. The first stage (MAC) uses a computationally cheap, non-structural matcher to filter candidates from a pool of memory items. That is, we redundantly encode structured representations as *content vectors*, whose dot product yields an estimate of how well the corresponding structural representations will match. The second stage (FAC) uses SME to compute a true structural match between the probe and output from the first stage. MAC/FAC has been fully implemented, and we show that it is capable of modeling patterns of access found in psychological data.

## Introduction

Similarity-based reminders range from the sublime to the stupid. On one extreme is being reminded by octaves in music or the periodic table in chemistry. On the other extreme are times when a bicycle reminds you of a pair of eyeglasses. Most often, reminders are somewhere in between, such as when a bicycle reminds you of another bicycle. Our theoretical attention is inevitably drawn to spontaneous analogy, i.e., structural similarity unsupported by surface similarity, partly because it offers perhaps our best entree to studying the creative process. However, a good model must also capture the frequency of different outcomes, and research on the psychology of memory retrieval points inescapably to a preponderance of the latter two types of similarity – (mundane) literal similarity, based on both structural and superficial commonalities – and (dumb) superficial similarity, based on surface commonalities. Rare events are hard to model. A major challenge for research on similarity-based reminding is to devise a model that will produce chiefly literal-similarity and superficial reminders, but still produce occasional

analogical reminders.

This paper presents MAC/FAC, a model of similarity-based reminding which attempts to capture these phenomena. We first review psychological evidence on retrieval and mapping of similarity comparisons and describe the design of MAC/FAC. We then describe computational experiments which simulate the patterns of access found in a psychological experiment, and close by describing further avenues to explore.

## Framework

Similarity-based transfer can be decomposed into subprocesses. Given that a person has some current *target* situation in working memory, transfer from prior knowledge requires at least (1) *accessing* a similar (*base*) situation in long-term memory, (2) *creating a mapping* from the base to the target, and (3) *evaluating* the mapping. In the structure-mapping framework (Gentner, 1983, 1988), mapping is the process by which two representations present in working memory are aligned and further inferences imported. The process of computing a mapping from one situation to another is governed by the constraints of *structural consistency* and *one-to-one mapping*.

This account differs from most psychological treatments by defining similarity in terms of correspondences between structured representations. Matches can be distinguished according to the kinds of commonalities present. An *analogy* is a match based on a common system of relations, especially involving higher-order relations.<sup>1</sup> A *literal similarity* match includes both common relational structure and common object descriptions. *Surface matches* are based primarily on common object descriptions along with some shared first-order relations.

There is considerable evidence that people are good at mapping. People can readily align two situations, preserving structurally important commonalities, making the appropriate lower-order substitutions, and mapping additional predicates into the target as *candidate inferences*. For example, Clement & Gentner (in press) showed people analogies and asked which of two lower-order relations,

<sup>1</sup>We define the order of an item in a representation as follows: Objects and constants are order 0. The order of a statement is one plus the maximum of the order of its arguments.

both shared by base and target, was most important to the match. Subjects chose relations that were governed by shared higher-order relations. In a second study, subjects showed the same sensitivity to connectivity and systematicity in choosing which predicates to map as candidate inferences from base to target. Further, people rate metaphors as more apt when they are based on relational commonalities than when they are based on common object-descriptions (Gentner & Clement, 1988) and they rate pairs of stories as more sound when they share higher-order relational structure than when they share object-descriptions (Gentner & Landers, 1985; Rattermann & Gentner, 1987). We also find effects of relational structure on judgments of similarity (Goldstone, Medin & Gentner, in press; Rattermann & Gentner, 1987) and on the way in which people align perceptually similar pictures (Markman & Gentner, 1990).

An adequate model of human similarity and analogy must capture this sensitivity to structural commonality, by involving structural representations and processes that align them. This would seem to require abandoning some highly influential models of similarity: e.g., modeling similarity as the intersection of independent feature sets or as the dot product of feature vectors. However, we show below that a variant of these nonstructural models can be useful in describing some aspects of access.

*Similarity-based Access from Long-term Memory:* There is psychological evidence that access to long-term memory relies more on surface commonalities and less on structural commonalities than does mapping. For example, people often fail to access potentially useful analogs (Gick and Holyoak, 1980). Ross (1984, 1987) further showed that, although people in a problem-solving task are often reminded of prior problems, these reminders are often based on surface similarity rather than on structural similarities between the solution principles.

In our research we used the "Karla the hawk" stories to investigate the determinants of similarity-based access. We put people in the position of trying to access analogy and similarity matches from long-term memory and asked which kinds of comparisons were easiest to retrieve (Gentner & Landers, 1985; Rattermann & Gentner, 1987). Subjects first read a large set of stories. Two weeks later, they were given new stories which matched the original ones in various ways. Some were true analogs of the first stories; others were surface matches, sharing lower-order events and object descriptors but not higher-order relational structure. Subjects were asked to write out any prior stories recalled while reading the new stories. Afterwards, they rated all the pairs for soundness: i.e., how well inferences could be carried from one story to the other.

The results showed that, although subjects rated the analogies as much more sound than the surface matches, they were more likely to retrieve surface matches. Surface similarity was the best predictor of memory access, while similarity in relational structure was the best predictor of subjective soundness and also of subjective similarity. This dissociation held not only between subjects, but also within subjects. That is, subjects given the soundness

task immediately after the cued retrieval task judged that the very matches that had come to their minds most easily (the mere-appearance matches) were highly unsound (i.e., unlikely to be useful in inference). This suggests that analogical access may be based on qualitatively distinct processes from analogical inferencing<sup>2</sup>.

*Comparison to Current Approaches.* Some models of similarity assume smart processes operating over richly articulated representations. Most case-based reasoning models have this character (Schank, 1982; Kolodner, 1988). These models are rich enough to capture processes like case alignment and adaptation. But their models of memory access involve intelligent indexing of structured representations, which can predict superhuman access behavior; that is, that people should typically access the best structural match, even if it lacks surface similarity with the current situation. Further, models that assume that elaborate structural mapping processes are used to compare the current situation with stored situations have the disadvantage of being hard to scale up to large data bases. The reverse set of advantages and disadvantages holds for approaches that model similarity as the result of a dot product (or some other operation) over feature vectors, as is commonly done in mathematical models of human memory (e.g., Medin & Schaffer, 1978) and in connectionist models of learning (Smolensky, 1988). These models, with their nonstructured representations and relatively simple processes, do not allow for the structural precision of people's similarity judgments and inferences. However, they provide an appealing model of access since: (1) these computations are simple enough to make it feasible to compute many such matches and choose the best (the scaling criterion); and (2) being simple, these models will not always produce the best match (the fallibility criterion). While this might be a disadvantage in a normative model, it could be an advantage in modeling human similarity-based access, provided that the best match is sometimes produced. Next we propose an approach that we think may offer the best of both kinds of models.

### The MAC/FAC model

The complexity of the phenomena in similarity-based access suggests a two-stage model. Consider the computational constraints on access. The large number of cases in memory and the speed of human access suggests a computationally cheap process. But the requirement of judging soundness, essential to establishing whether a match can yield useful results, suggests an expensive match process. A common solution is to use a two-stage process, where a computationally cheap filter is used to pick out a subset of likely candidates for more expensive processing (c.f. Bareiss & King, 1989). MAC/FAC uses this strategy. The

<sup>2</sup>The finding is not that higher-order relations do not contribute to retrieval. Adding higher-order relations led to non-significantly more retrieval in two studies and to a small but significant benefit in the third. The point is simply that higher-order commonalities have a much bigger effect on mapping once the two analogs are present than they do on similarity-based retrieval.

puzzling phenomena noted previously, we claim, can be understood in terms of the interactions of its two stages.

Figure 1 illustrates the components of the MAC/FAC model. The inputs are a pool of memory items and a *probe*, i.e., a description for which a match is to be found. The output is a memory description and a comparison of this description with the probe.

There is little consensus about the global structure of long-term memory. Consequently, we assume only that at some stage in access there is a pool of descriptions from which we must select one (or a few) which is most similar to a probe. We are uncommitted as to the size of this pool. It could be the whole of long-term memory, or a subset of it if one postulates mechanisms for restricting the scope of search, such as spreading activation or indexing<sup>3</sup>.

Both stages consist of a *matcher*, which is applied to every input description, and a *selector*, which uses the evaluation of the matcher to select which comparisons are produced as the output of that stage. Conceptually, matchers are applied in parallel within each stage. Since the role of the MAC stage is to produce plausible candidates for the FAC stage, we discuss FAC first.

### The FAC stage

The FAC matcher is simply the literal similarity computation defined by structure-mapping. Its output is a set of correspondences between the structural descriptions, a numerical *structural evaluation* of the overall quality of the match, and a set of candidate inferences representing the surmises about the probe sanctioned by the comparison. In subsequent processing, the structural evaluation provides one source of information about how seriously to take the match, and the candidate inferences provide potential new knowledge about the probe which must be tested and evaluated by other means. We implement this computation using SME, the Structure-Mapping Engine (Falkenhainer, Forbus & Gentner, 1989).

We use literal similarity rather than analogy in order to get the high observed frequency of surface reminders, which would mostly be rejected if FAC were strictly an analogy matcher. We believe this choice is ecologically sound because mundane matches are often the best guides to action. Riding a new bicycle, for instance, is often just like riding other bicycles (Gentner, 1989; Medin & Ortony, 1989). Associating actions with particular complex descriptions makes good computational sense because such associations can often be made before one can delineate exactly which aspects of a situation are relevant.

Currently FAC selects as output the best match, based on its structural evaluation, and any others within 10% of it. In pilot studies we have experimented with various criteria, such as broadening the percentage, selecting a fixed

number, and so forth. We settled on the 10% criteria because it generally returns a single result, only producing multiple results when there are two extremely close candidates. Depending on the assumptions one makes about subsequent processing, a modification which places a strict upper bound on the number produced (say, two) may also be appropriate.

Sometimes a probe reminds us of nothing. There are several ways this can arise in the MAC/FAC model. First, the FAC stage may not receive any candidates from the MAC stage (see below). Second, FAC might reject all candidates provided. This shows up by no match hypotheses being created; this has occurred, albeit rarely. Third, there could be a threshold on structural evaluations, so that matches below a certain quality simply were not considered. We view this as psychologically plausible, but do not include such thresholds currently because we have not yet found good constraints on them.

### The MAC stage

Even though the FAC stage is reasonably efficient<sup>4</sup>, it is too expensive to consider running it exhaustively on realistic-sized memories as the "inner loop" in an analogical processing system. The MAC stage uses an extremely cheap matcher to estimate how well FAC would rate comparisons, to filter candidates down to a manageable number.

One estimate is the number of match hypotheses that FAC would generate in comparing a probe to a memory item, the *numerosity* of the comparison. If very few local matches are hypothesized, then clearly the best global interpretation cannot be large. On the other hand, numerosity is not a perfect estimator, since having a large number of local matches does not guarantee a large global interpretation. This is true because (1) match hypotheses can end up being ungrounded because some of their arguments cannot be placed into correspondence (and are thus ignored), and (2) the mutual incompatibilities introduced by the 1:1 constraint may prevent a single large interpretation from forming, yielding instead several small ones.

The most straightforward way to compute numerosity is to actually generate and count the match hypotheses. This is what our original version of MAC/FAC did (Gentner, 1989). It also partly what ARCS (Thagard et al 1990) does. ARCS builds much of the network which ACME would build between target and base but between the probe and every item in memory. We view these solutions as psychologically and computationally implausible. Even with parallel and/or neural hardware, it is hard to see how the expense of generating match hypothesis networks between a probe and everything in a large pool of memory can provide realistic response times. Instead, we turn to a novel means of estimating numerosity.


Let  $P$  be the set of functors (i.e., predicates, functions, and connectives) used in the descriptions that constitute

<sup>3</sup>In current AI systems indexing often yields a unique description; we view this property as unlikely to scale. For example, there could be dozens or even hundreds of experiences which are similar enough to be put in the same index entry, yet different enough to make it worthwhile to save them as distinct memories.

<sup>4</sup> $O(n^2)$  for match hypothesis generation, where  $n$  is the number of items in base or target, and roughly  $O(\log(n^2))$  to generate a global interpretation, using the greedy merge algorithm of Forbus & Oblinger (1990).



Figure 1: The MAC/FAC model



memory items and probes. We define the *content vector* of a structured description as follows. A content vector is an n-tuple of numbers, each component corresponding to a particular element of P. Given a description D, the value of each component of its content vector indicates how many times the corresponding element of P occurs in D. Components corresponding to elements of P which do not appear in statements of D have the value zero. One simple algorithm for computing content vectors is to simply to count the number of occurrences of each functor in the description. Thus if there were four occurrences of IMPLIES in a story, the value for the IMPLIES component of its content vector would be four<sup>5</sup>. Thus content vectors are easy to compute from a structured representation and can be stored economically.

The MAC matcher works as follows: Each memory item has a content vector stored with it. When a probe enters, its content vector is computed. A score is computed for each item in the memory pool by taking the dot product of its content vector with the probe's content vector. These scores are fed to the MAC selector, which produces as output the best match and everything within 10% of it, as in the FAC stage. (We plan to add a threshold so that if every match is too low MAC returns nothing.)

Clearly, measuring similarity using content vectors has critical limitations, since the actual relational structure is not taken into account. But the dot product can be used to estimate relative similarity, since it is a good approximation to numerosity. (Essentially, the product of each corresponding component is an overestimate of the number of match hypotheses that would be created between functors of that type.) Content vectors are insufficient be-

cause they do not provide the correspondences and candidate inferences which provide the power of analogy. But by feeding MAC's results to the structural matcher of the FAC stage, we obtain the required inferential power.

This MAC matcher has the properties we desire. It is cheap, and could be implemented using a variety of massively parallel computation schemes, including connectionist. Next, we demonstrate that MAC/FAC provides a good approximation of psychological data.

### Computational Experiments

We have successfully tested MAC/FAC on a variety of descriptions, including simple metaphors and physics scenarios. Here we compare the performance of MAC/FAC with that of human subjects, using the "Karla the Hawk" stories. For these studies, we wrote sets of stories consisting of base stories plus four variants, created by systematically varying the kind of commonalities. All stories share first-order relations, but vary as follows:

	Common h.o. relations	Common object attributes
LS:	Yes	Yes
SF:	No	Yes
AN:	Yes	No
FOR:	No	No

As discussed above, subjects rated analogy (AN) and literal similarity (LS) as more sound than surface (SF) and FOR matches (matches based only on common first-order relations, primarily events). Previously, we tested SME running in analogy mode on SF and AN matches and found that it correctly reflected these human soundness rankings (Forbus & Gentner, 1989; Skorstad et al, 1987). Here we seek to capture human retrieval patterns: Does MAC/FAC duplicate the human propensity for retrieving SF and LS matches rather than AN and FOR matches. The idea is to give MAC/FAC a memory set of stories, then probe with various new stories. To count as a retrieval, a story must make it through both MAC and FAC.

<sup>5</sup>We have also experimented with normalized content vectors, to minimize the effects of size discrepancies. So far we have seen no significant empirical difference between these algorithms, but we suspect that normalization will be necessary when adding retrieval thresholds.

Table 1: Proportion of correct retrievals given different kinds of probes

1. Memory contains 9 base stories and 9 FOR matches; probes were the 9 LS, 9 SF, and 9 AN stories.
2. The rows show proportion of times the correct base story was retrieved for different probe types.

Probes	MAC	FAC
LS	1.0	1.0
SF	0.89	0.89
AN	0.67	0.56

In the psychological experiment, the human subjects had a memory set consisting of 32 stories, of which 20 were base stories and 12 were distractors. They were later presented with 20 probe stories which matched the base stories as follows: 5 LS matches, 5 AN matches, 5 SF matches and 5 FOR matches and told to write down any prior stories of which they were reminded. The proportions of reminders for different match types were .56 for LS, .53 for SF, .12 for AN and .09 for FOR. Across three variations of this study, this retrievability order has been stable:  $LS \geq SF > AN \geq FOR$ .

For the computational experiments, we encoded predicate calculus representations for 9 of the 20 story sets (45 stories). These stories are used in all three experiments described below.

*Simulation Experiment 1.* In our first study, we put the 9 base stories in memory, along with the 9 FOR stories which served as distractors. We then used each of the variants – LS, SF, and AN – as probes. This roughly resembles the original task, but MAC/FAC’s job is easier because (1) it has only 18 stories in memory, while subjects had 32, in addition to their vast background knowledge; (2) subjects were tested after a week’s delay, which may have caused some memory deterioration.

Table 1 shows the proportion of times the base story made it through MAC and through FAC. MAC/FAC’s performance is much better than that of the human subjects, perhaps partly because of the differences noted above. However, its results show the same ordering as those of human subjects:  $LS > SF > AN$ .

*Simulation Experiment 2.* To give MAC/FAC a stronger challenge, we put the four variants of each base story into memory. This made a larger memory set (36 stories) and also one with many competing similar choices. Each base story in turn was used as a probe. This is almost the reverse of the task subjects faced, and is more difficult.

Table 2 shows the mean number of matches of different similarity types that succeed in getting through MAC and through FAC. There are several interesting points here. First, the retrieval results (i.e., the number that make it through both stages) ordinally match the results for human subjects:  $LS > SF > AN > FOR$ . This degree of fit is encouraging, given the difference in task. Second, as expected, MAC produces some matches that are rejected by FAC. This number depends partly on the criteria for the two stages. Here, with MAC and FAC both set at 10%,

Table 2: Mean numbers of different match types retrieved when base stories used as probes

1. Memory contains 36 stories (LS, SF, AN, and FOR for 9 story sets); the 9 base stories used as probes
2. Other = any retrieval from a story set different from the one to which the base belongs.

Retrievals	MAC	FAC
LS	0.78	0.78
SF	0.67	0.44
AN	0.33	0.11
FOR	0.22	0.0
Other	1.33	0.22

Table 3: Mean numbers of different match types retrieved with base stories as probes

1. Memory contains 27 stories (9 SF, 9 AN, 9 FOR); 9 base stories used as probes.

Retrievals	MAC	FAC
SF	0.89	0.78
AN	0.56	0.45
FOR	0.22	0.11
Other	1.11	0.11

the mean number of memory items produced by MAC is 3.3, and the mean number accepted by FAC is 1.5. Third, as expected, FAC succeeds in acting as a structural filter on the MAC matches. It accepts all of the LS matches MAC proposes and some of the partial matches (i.e., SF and AN), and while rejecting most of the inappropriate matches (i.e., FOR and matches with stories from other sets).

*Simulation Experiment 3.* In the prior simulation, LS matches were the resounding winner. While this is reassuring, it is also interesting to know which matches are retrieved when there are no perfect overall matches. Therefore we removed the LS variants from memory and repeated the second simulation experiment, again probing with the base stories. As Table 3 shows, SF matches are now the clear winners in both the MAC and FAC stages. Again, the ordinal results match well with those of subjects:  $SF > AN > FOR$ .

*Summary of Simulation Experiments.* The results are encouraging. First, MAC/FAC’s ordinal results match those of human subjects. In contrast, the closest competing model, Thagard et al’s (1991) ARCS model of similarity-based retrieval, when given the Karla the hawk story in memory (along with 100 fables as distractors) and the four similarity variants as probes, produced two violations in its order of asymptotic activation. Its asymptotic activations were LS (.67), FOR (-.11), SF (-.17), AN (-.27). Thus MAC/FAC explains the data better than ARCS. This is especially interesting because Thagard et al argue that a complex localist connectionist network which integrates semantic, structural, and pragmatic constraints is required to model similarity-based reminding. While such models are intriguing, MAC/FAC shows that a simpler

model can provide a better account of the data.

Finally, and most importantly, MAC/FAC's overall pattern of behavior captures the motivating phenomena: (1) it produces a large number of LS matches, thus satisfying the *primacy of the mundane* criterion; (2) it produces a fairly large number of SF matches, thus satisfying the fallibility criterion; (3) it produces a small number of analogical matches, thus satisfying the *existence of rare events* criterion; and finally, (4) its algorithms are simple enough to apply over large-scale memories, thus satisfying the *scalability* criterion.

## Discussion

We have presented MAC/FAC, a two-stage similarity-based model of access. The MAC stage uses content vectors, a novel summary of structured representations, to provide an inexpensive "wide net" search of memory, whose results are pruned by the more expensive literal similarity matcher of the FAC stage to arrive at useful, structurally sound matches. We demonstrated that MAC/FAC can simulate the patterns of access exhibited by human subjects. We believe that the psychological issues MAC/FAC raises are worth further study. MAC/FAC is reasonably efficient, even on serial machines, so we believe it could be a useful component in performance-oriented AI systems also.

In addition to the psychological issues raised earlier, there are several computational studies in preparation using MAC/FAC. These include:

*Experiments with larger knowledge bases:* A crucial question for any access model is how well it scales to substantially larger memories. Two avenues we are exploring are: (1) using the CYC knowledge base as a source of descriptions and (2) using MAC/FAC as a tool on the ILS Story Archive Project to aid in spotting potentially relevant links between stories.

*Larger-scale process models:* Several psychological questions about access cannot be studied without embedding MAC/FAC in a more comprehensive model of analogical processing. For example, there is ample evidence that subjects can "tune" their similarity judgements when the items being compared are both already in working memory. While it seems clear that MAC is impenetrable, it is hard to tell whether or not FAC is tunable or whether a separate similarity engine is required. Order effects in analogical problem solving (Keane, in press) suggest the latter. How can the access system be used to incrementally construct abstractions and indexing information to help structure long-term memory (c.f. Skorstad, Gentner, and Medin 1988)?

## Acknowledgements

This research was supported by the Office of Naval Research (Contract No. N00014-89-J-1272) and by the National Science Foundation. We thank Ray Bareiss, Mark Burstein, Gregg Collins, Brian Falkenhainer, Rob Goldstone, Art Markman, Doug Medin, and Mary Jo Rattermann for discussions of these issues.

## References

- Bareiss, R. and King, J. 1989. Similarity Assessment in Case-Based Reasoning, Proceedings of the DARPA Case-Based Reasoning Workshop, Pensacola Beach, FL, pages 67-71.
- Clement, C. A., & Gentner, D. (in press). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*.
- Falkenhainer, B., Forbus, K., Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and examples, *Artificial Intelligence*, 41:1-63.
- Forbus, K. and Gentner, D. 1989. Structural evaluation of analogies: What Counts?, Proceedings of the Cognitive Science Society.
- Forbus, K. and Oblinger, D. 1990. Making SME greedy and pragmatic. Proceedings of the Cognitive Science Society.
- Gentner, D. 1983. Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7(2).
- Gentner, D. 1988. Analogical inference and analogical access. In A. Prieditis (Ed.), *Analogica* (pp. 63-88). Los Altos, CA: Morgan Kaufmann.
- Gentner, D. 1989. Finding the needle: Accessing and reasoning from prior cases. Proceedings: Case-based reasoning workshop, DARPA ISTO, Pensacola, Florida.
- 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* (Vol. 22, pp. 307-358). New York: Academic Press.
- Gentner, D., & Landers, R. 1985. Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Cybernetics and Society* (pp. 607-613), Tucson, AZ. New York: IEEE.
- Gick, M. L., & Holyoak, K. J. 1980. Analogical problem solving. *Cognitive Psychology*, 12:306-355.
- Goldstone, R. L., Medin, D. L., & Gentner, D. (in press). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology*.
- Keane, Mark T. (in press) Similarity and Ordering Constraints on Analogical Mapping.
- Kolodner, J.L. (Ed.) 1988. *Proceedings of the First Case-Based Reasoning Workshop*, Morgan Kaufmann, Los Altos, CA.
- Markman, A. B., & Gentner, D. 1990. Analogical mapping during similarity judgments. Proceedings of the Twelfth Annual Conference of the Cognitive Science Society, 38-44, Cambridge, MA. Hillsdale, NJ: Erlbaum.
- Medin, D. L., and Ortony, A. 1989 Psychological Essentialism. In S. Vosniadou and A. Ortony (Eds.), *Similarity and Analogical Reasoning*, 179-195. New York: Cambridge University Press.
- Medin, D. L., & Schaffer, M. M. 1978. Context theory of classification learning. *Psychological Review*, 85:207-238.
- Rattermann, M. J., & Gentner, D. 1987. Analogy and similarity: Determinants of accessibility and inferential soundness. In Proceedings of the Ninth Annual Meeting of the Cognitive Science Society, 23-34, Seattle, WA.
- Ross, B. H. 1984. Reminders and their effects in learning a cognitive skill. *Cognitive Psychology*, 16:371-416.
- Ross, B. H. 1987. This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13:629-639.
- Schank, R. 1982. *Dynamic Memory*. Cambridge University Press, Cambridge, England.
- Skorstad, J., Falkenhainer, B., and Gentner, D. 1987. Analogical processing: A simulation and empirical corroboration. Proceedings of AAAI-87.
- Skorstad, J. Gentner, D. and Medin, D. 1988. Abstraction processes during concept learning: A structural view. Proceedings of the Cognitive Science Society.
- Smolensky, P. 1988. On the proper treatment of connectionism. *Behavior and Brain Sciences*. 11:1-74.
- Thagard, P., Holyoak, K.J., Nelson, G., and Gochfeld, D. 1990. Analog retrieval by constraint satisfaction. *Artificial Intelligence*, 46:3.