

Analogical Asides on Case-Based Reasoning

Mark T. Keane

Department of Computer Science,
Trinity College Dublin,
Dublin 2, IRELAND

Abstract. This paper explores some of the similarities and differences between cognitive models of analogy and case-based reasoning systems. I first point out a paradox in the treatment of adaptation in analogy and in case-based reasoning; a paradox which can be only resolved by expanding the role of adaptation in cognitive models of analogy. Some psychological research on the process of adaptation in human subjects is reported and then the implications of this research are propagated into analogy and then on into CBR. The argument is that some of the existing stages in CBR should be integrated into a more stream-lined architecture that would be more efficient than current schemes.

1. Introduction

The present paper is part of a strong tradition of inter-communication between case-based reasoning (CBR) and psychologically-oriented work in cognitive science. My research roots are in cognitive psychology and the backdrop to my current case-based reasoning research is a broad, but unfinished, canvas portraying aspects of cognitive models of analogy (see Keane, 1985, 1987, 1988, 1990). Up until now, I have kept these two areas quite separate in my head, perhaps to reduce some cerebral cross-talk. This separation has probably been unwise because it has hidden a number of interesting connections between the two areas; cognitive models of analogy have a lot to learn from CBR. In this paper, I hope to convince you that the converse is also the case; that CBR can also learn from cognitive models.

If you put analogy models and CBR systems side by side, a striking paradox appears: a strange and curious state of affairs arises that reflects badly on analogy models. I call this paradox the "adaptation paradox". In the following sections, I elaborate this paradox and explain how I think it can be resolved. This resolution hinges on fixing some deficits in cognitive models of analogy. Having resolved this paradox, it occurred to me that the solution had implications for CBR; implications which have been followed up in joint research with Barry Smyth at

the Hitachi Dublin Laboratory in Trinity College (see Smyth & Keane, in this volume). But, enough of the preamble, what of the paradox ?

2. The "Adaptation Paradox" in Analogy Research

The "adaptation paradox" arises out of the differential treatment that adaptation receives in case-based reasoning and in cognitive models of analogy. The paradox can be stated in terms of three propositions, that are supported by current theory and/or empirical research but they propagate disruptive waves through analogy theory. The first proposition is that:

- (1) In CBR, adaptation is a necessary and significant step.

The re-use of past experience is trivial if it is just the direct application of concrete knowledge or the instantiation of established, abstract knowledge (e.g., schemas). One of the key insights in CBR, and analogy, is that prior knowledge that is not directly relevant to the current situation can be made relevant by perturbing or adapting it to fit the current situation. So, most CBR systems have finely-crafted adaptation rules to carry out this perturbation. Indeed, I was struck by Ian Smith's proposal (workshop communication) that adaptation is *the* unique and distinctive process in case-based reasoning, even though much CBR research concentrates on retrieval.

The second proposition is based on an accepted interpretation of people's problem solving behaviour using analogies. Cognitive theory of analogy acknowledges that there is an adaptation stage, but few investigators acknowledge that it occurs in the classical experiments in the field. People's use of analogy is typically seen as just being the product of analogical mapping. Thus, the second proposition is that:

- (2) People solve problems by analogy through the use of mapping, with little or no adaptation

The single study known to me in the literature that examines adaptation, is one in which people solve verbal maths problems by analogy (see Novick & Holyoak, 1991). The analogues used for the problems, in this study, could be mapped easily but did not fit the problems particularly well, so Novick & Holyoak observed people struggling to adapt the analogue to the problem. They characterised this adaptation as a slow, conscious process which could involve the making of further inferences and the use of other background knowledge (see also Keane, 1990, on validation). Thus, given that people typically solve insight

problems by analogy in a swift fashion, it is generally assumed that they simply map the analogy and do not use any conscious, adaptation process.

The final proposition is a theoretically plausible assumption that leaves the jaws of the paradox clamped firmly on our vitals:

(3) In theory, more adaptation should be required in analogy than in CBR

In CBR, the base cases are semantically-close to the target problem, they tend to be literally similar. In analogy, the two domains are semantically distant; they may be as far apart as the atom and the solar system or the legs of a chair and a theorem prover. So, it is plausible to assume that the adaptation required in analogies should be much more extensive and complicated than it is in CBR. Yet, as we have seen, propositions 1 and 2 suggest that all of the action on adaptation is going in CBR and not in analogy.

This paradox has nasty implications for either CBR or analogy research. It suggests that either CBR researchers are misguided in spending so much effort on adaptation or that analogy researchers are living some form of existentialist "bad faith". I think that the former is false but that the latter is partly true. So, let us examine the theoretical assumptions made by cognitive theorists in more detail.

3. Cognitive Models of Analogy

Cognitive theory of analogy posits five main processing stages: representation, analogue retrieval, analogical mapping, adaptation and induction. Clearly, the way an analogue and a problem are *represented* has an important impact on subsequent problem solving attempts by analogy (see e.g., Novick, 1988). The *retrieval* of an analogous case to solve a problem is harder if the domains are semantically-distant to one another (e.g., Gentner, Ratterman & Forbus, 1993; Keane, 1987). *Analogical mapping* is the core process during which the analogy is drawn between the two domains. After an analogy is mapped it may have to be *adapted* to make the solution fit the target problem (see Novick & Holyoak, 1991; Keane, 1994). With respect to *induction*, Gick & Holyoak (1983) have shown that subjects can induce a generalisation from the correspondences between two analogous stories that facilitates subsequent problem solving.

3.1 Theory of Analogical Mapping

One of the other notable differences between analogy research and CBR is that retrieval and mapping are kept quite separate. In CBR, retrieval and

mapping/matching are usually merged. In analogy, the mapping stage is quite complex and has received the bulk of theoretical attention. Recently, considerable theoretical consensus has emerged in the treatment of analogical mapping. Keane, Ledgeway & Duff (1991, in press) have used a three-levelled, meta-theoretical framework to characterise this theory (see Palmer, 1989; Marr, 1982): the levels of description being concerned with informational constraints, behavioural constraints and hardware constraints.

At the highest level, *informational constraints* characterise the informational aspects of the task situation. This level is concerned with describing what an analogy is; that is, what needs to be computed to produce appropriate outputs given certain inputs (akin to Marr's computational level). Below this level are *behavioural constraints* that have to capture the empirical facts of people's observable analogical behaviour (Marr's algorithmic level). Hence, this level should include constraints that predict when one analogy is harder than another, the relative differences in processing times for different analogies and the sorts of errors that people produce. Finally, there is the level of *hardware constraints* that aim to capture the neurological primitives of analogical thought, about which we know little (Marr's hardware level).

There is considerable agreement on the informational constraints that are important in analogical mapping: they are structural, similarity and pragmatic constraints (see e.g., Gentner, 1983; Holyoak & Thagard, 1989; Keane, 1990). The most important set of constraints are *structural constraints*. These constraints are used to enforce a one-to-one mapping between the two domains (Falkenhainer, Forbus & Gentner, 1986, 1989; Holyoak & Thagard, 1989). Structural constraints rely on several techniques:

- *make matches only between entities of the same type*; only attributes are matched with attributes, objects with objects and two-place predicates with two-place predicates. For example, in matching REVOLVES(A B) and REVOLVES(C D), the REVOLVES predicate would never be matched with the object C. This reduces the total number of matches that need to be considered (see Gentner, 1983; Holyoak & Thagard, 1989).
- *exploit structural consistency*; that is, if the propositions REVOLVES(A B) and REVOLVES(C D) match, then the arguments of both should also be matched appropriately, A with C and B with D. This is especially useful in eliminating many-to-one and one-to-many matches (see Falkenhainer, Forbus & Gentner, 1986, 1989).

- *favour systematic sets of matches* (Gentner's systematicity principle); that is, if one has two alternative sets of matches then the mapping with the most higher-order connectivity should be chosen. This constraint aids the choice of an optimal mapping from among many alternative mappings.

Similarity constraints can disambiguate between alternative matches. When these constraints is applied only identical concepts are matched between the two domains (Gentner, 1983) or, more loosely, semantically-similar concepts are matched (Gick & Holyoak, 1980; Holyoak & Thagard, 1989). Semantic similarity can be used to disambiguate matches; if one match in a set of one-to-many matches is more similar than the others, it can be preferred. Finally, there are *pragmatic constraints* (e.g., Holyoak, 1985; Holyoak & Thagard, 1989; Keane, 1985). Again, these constraints may disambiguate a set of matches. For example, in a certain analogical mapping situation, one match may be pragmatically more important (or goal-relevant) than other alternatives and so it will be preferred over these alternatives.

Informational constraints constitute a high-level specification of what makes a particular comparison between two domains an *analogical* comparison. They constitute a *competence* theory of analogical mapping (see also Gentner, 1989). Computationally, they can be and have been implemented in a variety of different models (as we will see later). As such, they capture the significant informational aspects of analogical comparisons. However, this level of description on its own is not sufficient to constitute a cognitive model (cf. Newell, 1990, for more general arguments on this point). For an adequate cognitive model of analogical mapping we need to elaborate the *behavioural constraints* on analogising. Keane et al. (in press) elaborate two such constraints: working memory limitations and the effects of background knowledge (see Keane, 1988, 1990).

3.2 Models of Analogical Mapping

The three main cognitive models of analogical mapping in the literature all instantiate the above informational constraints to varying degrees and some also include behavioural constraints.

Falkenhainer et al.'s (1986, 1989) *Structure Mapping Engine* (SME) implements both structural and similarity constraints in a serial way. SME finds all the legal local matches between two domains and then combines these into alternative interpretations of the comparison. SME is explicitly designed to construct all possible maximal interpretations for a given comparison between two domains.

When SME has an appropriate set of match rules -- the analogy match rules -- it instantiates Gentner's (1993) structure-mapping theory. However, it can also be used as a tool, when different match-rule sets are used. More recently, versions of the model have been developed to include pragmatic constraints (Forbus & Oblinger, 1990) and to take account of behavioural constraints (see Forbus, Ferguson & Gentner, 1994). SME uses serial processing techniques, although portions of the model have been parallelized recently (see Forbus, et al., 1994).

Holyoak & Thagard's (1989) *Analogical Constraint Mapping Engine (ACME)* uses parallel constraint satisfaction in an interactive network to find the optimal mapping between two domains. It implements the structural, similarity and most pragmatic constraints. ACME establishes a network of units or nodes. Each node represents a legal match between two predicates. The excitatory and inhibitory connections between these nodes implement the various constraints. So, ACME uses a localist representation of analogical matches and parallel constraint satisfaction to settle on the optimal mapping between two domains.

Keane's (Keane & Brayshaw, 1988; Keane, 1990; Keane et al., in press) *Incremental Analogy Machine (IAM)* implements all the informational and behavioural constraints mentioned above using serial constraint satisfaction. It generates a single, optimal interpretation based on a small subset of the possible mappings between the two domains. IAM builds up this mapping incrementally by selecting a small portion of a base domain for mapping, mapping it and then moving on to map another portion. Typically, it will construct a single mapping that will tend to be the optimal interpretation. However, if it has to, IAM can consider several alternative interpretations. Again, it deals with these alternatives incrementally, one after the other. So, if the first mapping that is built is less than optimal, IAM will undo the matches found and try an alternative mapping.

ACME and IAM try to produce a single, optimal mapping because this seems to be more psychologically plausible. People are less likely to generate all possible mappings. More recently, modifications have been made to the SME algorithm to reduce the number of mappings it generates and to enable it to map analogies incrementally (see Forbus & Oblinger, 1990; Forbus, et al, 1994).

4. Assumptions that Banish Adaptation

The predictions about people's normal analogical behaviour are largely derived from these mapping models. Hence, it is these models that give rise to the

adaptation paradox elaborated earlier. Consider the normal account of what happens when people solve a problem by analogy.

Gick & Holyoak (1980) demonstrated analogical problem solving by giving subjects story analogues to target, insight problems. For example, in many experiments they used Duncker's (1945) radiation problem:

A doctor is attempting to destroy a malignant tumour using rays. The doctor needs to use high-intensity rays to destroy the tumour but these high-intensity rays will destroy the healthy tissue surrounding the tumour. If the doctor uses low-intensity rays then the healthy tissue will be saved but the tumour will remain unaffected too.

This dilemma can be solved by a "convergence solution": the doctor can send low-intensity rays from a number of different directions so that they converge on the tumour, summing to a high intensity to destroy it. However, only about 10% of subjects produce this solution spontaneously (Gick & Holyoak, 1980).

Now consider the following story:

A general is launching an assault on a fortress. He is prevented from using his whole army to take the fortress because the roads leading to it are mined to explode if large groups of men pass over them. So, he divides his army up into small groups of men and sends them along different roads to the fortress so that they converge on it.

When subjects were given a longer version of this story to memorise and then asked if they could use it to solve the radiation problem, the rates of convergence solutions to the problem rise to about 80% (see Gick & Holyoak, 1980). In short, people can use the analogous story to solve the problem.

The general, however, knew just what to do. He divided his army into small groups of men and dispatched each group to the head of a different road. When all was ready he gave the signal and each group marched down a different road. Each group continued to the fortress so that the entire army finally arrived together at the fortress, at the same time. The fortress fell and the king was forced into exile.




Figure 1: The Text of the Solution to the General Story, the Suggested Propositional Representations of its Contents and the Analogous Solution Produced

Gick & Holyoak (1980) consider the base and target to be represented propositionally. Figure 1 gives some indication of the sorts of representations they use; it shows the original text of the solution to the General story and the propositional representations for it and the solution to the radiation problem. People solve these sorts of analogies relatively quickly; the analogous solution generally enters consciousness fully-formed and their reports on the mappings tend to be ad-hoc rationalisations (see Gick & Holyoak, 1980, Expt. 1). So, according to current analogy theory, whatever people are doing here, it is not adaptation. But, how can this be ?

If you examine the representational assumptions made by Gick & Holyoak, it is clear that they are a little 'convenient'. The original version of the story has many irrelevant details in it that do not find their way into the assumed representation (see Figure 1). Gick & Holyoak recognised this representational move (although the point has been ignored since); they maintained that both domains were represented at many levels of abstraction and that there was an optimal representational level at which the analogy was drawn (i.e., the level represented in Figure 1). However, they had no coherent account of how these levels were formed and have never investigated this facet of analogising. Indeed, on reflection, the idea that every domain is encoded at multiple levels of abstraction is a shade implausible, as it further complicates an already complicated mapping process. By the way, no coherent solution to the problem of finding *the* optimal

representation for two domains has yet been found; some form of re-representation or re-interpretation is required, but this must be done in a fashion that is sensitive to the requirements of a mapping, that has yet to be constructed (but see O'Hara & Indurkha, this volume, for one possible direction).

So, if we cannot assume such convenient representations, what *is* going on when people draw analogies ?

5. Undoing Some Assumptions Makes Adaptation Appear

There are two simple counter-assumptions we can make that resolve the adaptation paradox and make analogy models look a little better:

- The representations are messy
- Adaptation is happening, but in a speedy and automatic fashion

First, we can assume that the representations are not the convenient ones shown in Figure 1, but are rather much messier representations, with parts that do not correspond well between the two domains. Second, given these ill-fitting representations we can assume that the base analogues have to be adapted to fit the target problem. So, the solution to the General story should be adapted to fit the requirements of the radiation problem. For example, "signalled to march" might have to be modified to be "sent". However, we should also note that this adaptation is being carried out in a speedy, automatic fashion, quite unlike the slow conscious form of adaptation postulated by Novick & Holyoak.

Taking this theoretical stance, I have recently carried out some experiments that attempt to show that this form of covert adaptation occurs in the solution of insight problems by analogy (see Keane, 1994, for details). These experiments use analogues which have two alternative mappings, each of which could be used to solve the problem; adaptability of these mappings is then manipulated systematically.

Table 1 Summary of the Stories Used by Keane (1994)

Setting

People are stuck in a towering inferno and are trying to evacuate the building.

Plan 1:

A helicopter is positioned and a winch is thrown to the people, they grab hold of it and tie onto the winch, then swing from one building to the next.

Plan 2:

A helicopter is positioned and a winch is lowered to people, they reach for the winch, with a stick, tie on and are carried to the ground.

Plan 3:

Helicopter is positioned and a winch is lowered to people, they lasso the winch with a rope, tie on and are carried to the ground.

6. Evidence for Covert Adaptation in Analogy

Keane (1994) has developed a variant of the standard experiment used to test analogical problem solving. Like the Gick & Holyoak experiments, it uses story analogues and an insight problem. The insight problem is Maier's (1931) two-string problem which involves a situation in which two strings, hanging from the ceiling of a room, have to be tied together. However, the two strings are too far apart for a person to reach one string while holding the other. Traditionally, subjects propose a variety of solutions to this problem including a *stick* solution (use a stick to bring the string that is furthest away, within reach) and a *swing* solution (swing one of the strings so that it comes within reach while holding the other string).

A number of different story analogues were used with this problem. All the stories had a standard structure: setting, failed plan, successful plan (see Table 1). The plans used in the stories were designed to suggest either swing or stick-reaching or lasso-reaching solutions to the problem. These solutions should be differentially adaptable, as solutions to the two-string problem. The stick plan should be less adaptable than the swing plan. The stick plan proposes that a stick be used to reach one of the strings, a stick that is not present in the problem statement. Hence, the analogy suggests various actions of reaching and grabbing

but also the creation of a new object in the two-string problem. The swing plan simply suggests the new actions of swinging and grabbing, without the need to introduce a new object. As such, the analogous swing solution is easier to adapt. It may also be easier to adapt because, unlike the stick plan, it does not violate implicit constraints in the problem situation (i.e., that the problem is to be solved with the objects given). The *lasso plan* is similar to the stick plan, with the exception that the rope is reached by lassoing it with another piece of rope. Intuitively, although this plan is very similar to the stick plan, it is a less adaptable solution to the problem because a non-rigid object (e.g., the lasso) is less effective for reaching things than a rigid object (e.g., the stick).

We paired the stick plan with one or other of the other plans in two different versions of the story: swing-stick or lasso-stick. The prediction was that if people were taking into account the relative adaptability of the plans in assessing which one to map then in the stick-swing case, the stick plan should be avoided but in the stick-lasso case, the stick plan should be preferred. Current models would predict a 50-50 choice of one or other of the plans, as they do not take adaptation factors into account. This is exactly what we found in the experiment, as roughly three times more people use the stick plan as their solution when it is the more adaptable plan (see Figure 2).

So, this experiment shows, albeit in a very constrained situation, that people are sensitive to the adaptability of different aspects of an analogue when forming an analogical mapping. In short, all of the time in previous experiments when people seemed to be just mapping an analogue, it is likely that they were adapting it to the problem; but in a swift, automatic fashion. People are clearly powerful adaptation specialists. Exactly how they adapt things is a story for another day.

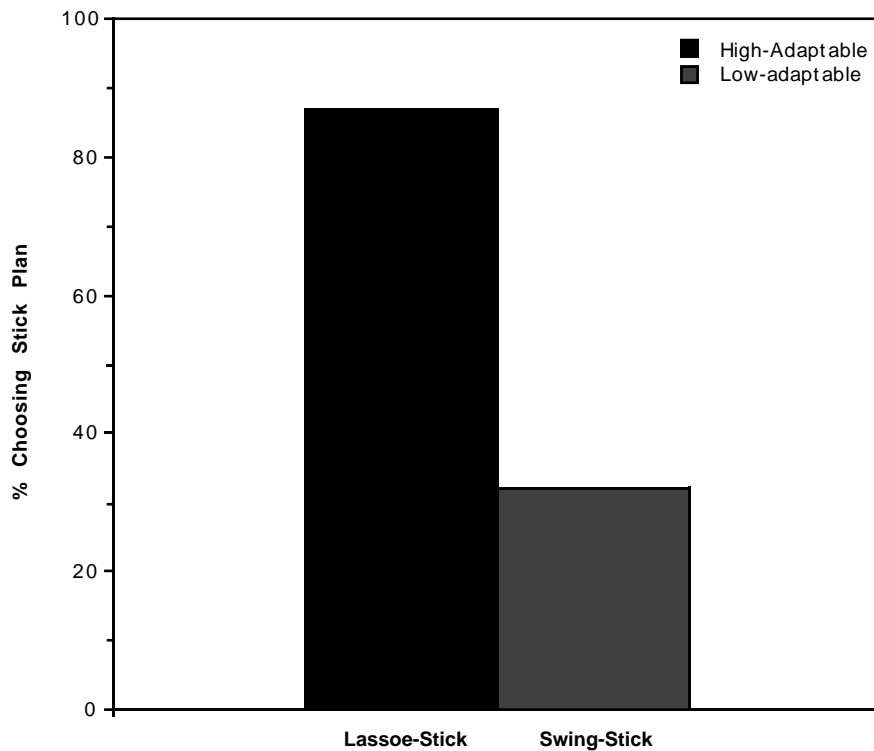


Figure 2 The Percentages of Subjects Selecting the Stick Plan in the Two Conditions of Experiment

7. Implications for Case-Based Reasoning

What then are the implications of this result for research in CBR? Well, some of the implications arise out of how cognitive models might be modified to incorporate aspects of adaptation. There are two different ways in which this form of covert adaptation might be integrated into analogy models. First, one could have the mapping mechanism produce a variety of alternative mappings and then have a post-processor that goes through these mappings evaluating them and scoring them to find the optimal one. This is the sort of architecture that Gentner (1989) has proposed for SME (see also Falkenhainer, 1987, 1990). The second option is to integrate adaptation directly into the mapping mechanism. This is the sort of solution that IAM and ACME are committed to because of the stance they take on mapping. So, as mapping is performed, so too is some form of adaptation. It is this solution that can be transposed into CBR systems.

CBR has been presented as a multi-stage process: retrieval/matching, adaptation and simulation. In cognitive theory, processing stages are often used as conceptual crutches rather than fundamental distinctions about a phenomenon. In time, when a phenomenon is understood better, these stages often fade away as

they are combined together. Maybe it is time to integrate the stages of CBR. Retrieval and mapping are already combined to some extent in many CBR systems. The implication I would like to take from these results is that retrieval and adaptation should also be integrated. So, the sort of integration demanded in cognitive models by the research outlined above may suggest a better, integrated architecture for CBR systems.

In a loose sense, you can think of these ideas as being implemented by rules of the form:

```
IF      X and Y are matched AND
        X can be adapted to Y
THEN favour this match over alternatives
```

Barry Smyth has arrived at a similar idea by a different route, working purely within CBR. His basic insight is that *CBR systems really need to retrieve the most adaptable case to the current problem, rather than the one that is the most semantically similar*. Semantic similarity is typically used in retrieval algorithms but it is just an approximation to a similarity metric based on adaptability. These ideas have been realised in a working system, called *Déjà Vu* (see Smyth & Cunningham, 1992; Smyth & Keane, this volume), that carries out automated programming in the domain of plant-control programs. We have yet to prove empirically that this type of integrated CBR is more efficient than standard approaches but at present it is clear that the weight of the theoretical arguments are in our favour.

8. Conclusions

In this paper, I have taken a circuitous route through cognitive theory in order to extrapolate some analogical parallels between cognitive models of analogy and case-based reasoning. In summary, it seems that given sufficient domain knowledge, people have efficient mapping/adaptation systems. Indeed, the behaviour they indulge in suggests that the mapping and adaptation stages should be integrated theoretically. If this idea is transferred into CBR it implies that retrieval and adaptation should be integrated, in as much as this is possible, without overburdening the retrieval stage. Such a retrieval mechanism would produce few false positives and should only retrieve highly-adaptable cases. As a further step, if I was feeling bold, I might propose that this integrated view of CBR should form the basis of a canonical architecture for CBR systems. However, at present, I would be happy for the latter proposal to be viewed simply as an attractive challenge.

9. References

- Duncker, K. (1945). On problem solving. *Psychological Monographs*, 58 (Whole No. 270).
- Falkenhainer, B. (1987). An examination of the third stage in the analogy process: Verification-based analogical learning. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*. 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). Structure-mapping engine. *Proceedings of the Annual Conference of the American Association for Artificial Intelligence*.
- Falkenhainer, B., Forbus, K.D., & Gentner, D. (1989). Structure-mapping engine. *Artificial Intelligence*, 41, 1-63.
- Forbus, K.D., Ferguson, R.W., & Gentner, D. (1994). Incremental structure mapping. In A. Ram & K. Eiselt (Eds.), *Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Forbus, K.D. & Oblinger, D. (1990). Making SME greedy and pragmatic. *Twelfth Annual Conference of the Cognitive Science Society*. Hillsdale: Erlbaum.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. (pp. 267-297). Cambridge: Cambridge University Press.
- Gentner, D., Rattermann, M.J., & Forbus, K.D. (1992). The roles of similarity in transfer. *Cognitive Psychology*, 25, 431-467.
- Gick, M.L., & Holyoak, K.J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M.L., & Holyoak, K.J. (1983). Schema induction in analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Holyoak, K.J. (1985). The pragmatics of analogical transfer. *The Psychology of Learning and Motivation*, 19, 59-87.
- Keane, M. (1985). On drawing analogies when solving problems: A theory and test of solution generation in an analogical problem solving task. *British Journal of Psychology*, 76, 449-458.

- Keane, M.T. (1987). On retrieving analogues when solving problems. *Quarterly Journal of Experimental Psychology*, 39A, 29-41.
- Keane, M.T. (1988). *Analogical Problem Solving*. Chichester: Ellis Horwood (Simon & Schuster in N.America).
- Keane, M.T. (1990). Incremental analogising: Theory and model. In K.J. Gilhooly, M.T. Keane, R. Logie & G. Erdos (Eds.), *Lines of Thinking: Reflections on the Psychology of Thought*. Vol. 1. Chichester: John Wiley.
- Keane, M.T. (1994). Adaptation as a selection constraint on analogical mapping. In A. Ram & K. Eiselt (Eds.), *Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Keane, M.T., & Brayshaw, M. (1988). The Incremental Analogy Machine: A computational model of analogy. In D. Sleeman (Ed.), *Third European Working Session on Learning*. London: Pitman/San Mateo, Calif.: Morgan Kaufmann.
- Keane, M.T., Ledgeway, T. & Duff, S. (1991). Constraints on analogical mapping. In K.J. Hammond & D. Gentner (Eds.), *Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Keane, M.T., Ledgeway, T. & Duff, S. (in press). Constraints on analogical mapping: A comparison of three models. *Cognitive Science*.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard: Harvard University Press.
- Novick, L.R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 14, 510-520.
- Novick, L.R. & Holyoak, K.J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 17, 398-415.
- O'Hara, S., & Indurkha, B. (this volume). Incorporating (re)-interpretation into Case-Based Reasoning. *Proceedings of the First European Workshop on Case-Based Reasoning*. Amsterdam: Springer-Verlag.
- Palmer, S.E. (1989). Levels of description in information processing theories of analogy. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. (pp. 267-297). Cambridge: Cambridge University Press.
- Smyth, B., Cunningham, P. (1992). Déjà Vu: A Hierarchical Case-Based Reasoning System for Software Design. *Proceedings of the 10th European Conference on Artificial Intelligence*. (pp. 587 - 589). Vienna, Austria

Smyth, B., & Keane, M.T. (this volume). Retrieving adaptable cases.
Proceedings of the First European Workshop on Case-Based Reasoning.
Amsterdam: Springer-Verlag.