The Cognitive Processes Underlying Complex Analogies:
Theoretical and Empirical Advances

Mark T. Keane

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

Correspondence and proofs to be sent to:

Dr. Mark Keane,
Cognitive Science Group,
Department of Computer Science,
Trinity College,
Dublin 2,
IRELAND

E-mail: mkeane@cs.tcd.ie
Phone: +353-1-7021534
Fax: +353-1-6772204

Abstract

The psychological investigation of complex analogies (like the solar system / atom analogy) began in the late 1970s and in the intervening period has grown into a prototypical, area of Cognitive Science research. During this time the investigation of complex analogies have gone from rough general notions to detailed computational models. This paper traces the theoretical and empirical development in the area and plots the major advances which have occurred. Apart from reviewing the literature, the main aim of the paper is to present this area as a template for the development of other Cognitive Science research topics. Thus, it should have important implications for the methodology of modern research in Cognitive Psychology. Finally, some outstanding problems in the area and likely directions for future research are considered.

Introduction

While there is a long history of psychological research on proportional analogies (red:stop::green:go, see e.g., Sternberg, 1977), it is only relatively recently that we have come to appreciate how people comprehend and use complex analogies; the sort of analogies which are used in scientific creativity (e.g., the analogy between the solar system and the atom) and in everyday life (e.g., the analogy between the wars in Afghanistan and Vietnam). Furthermore, the problems faced by psychometric research on proportional analogies are very different to those faced by workers interested in complex analogies.

Complex analogies differ in many ways from proportional analogies. First, they are more complex in the sense that they involve a lot more information. Complex analogies, unlike their proportional counterparts involve intricate systems of inter-dependent relations, rather than just having a single relation (e.g., red *signals* stop as green *signals* go). For example, our knowledge of the solar system consists of many relations about planets *revolving around* the sun and *being attracted* by the sun, and causal dependencies between these relations (see e.g., Gentner, 1983). When we draw an analogy using this domain of knowledge, many of these relations are carried over to structure our understanding of the atom domain (assuming that we know little about atoms at the outset). Second, the format of complex analogies is not as constrained as that of proportional analogies (see Weitzenfeld & Klein, 1979). Proportional analogies always have the a:b::c:d format, whereas complex analogies appear in many different guises as non-literal aspects of language, as creative ideas, as poems, as arguments and as adverts. Third, in proportional analogies you are given both parts of the analogy, the base domain (i.e., red is to stop) and the target domain (green is to go), whereas in complex analogies you may only be given the target domain or a base domain or fragments of either domain. So, with complex analogies you may need some account of how the base domain is found or how the fragments of either domain are accessed in memory. In short, in complex analogies, there is a whole new set of issues which need to be dealt with, issues which do not arise in research on proportional analogies.

In this paper, I will concentrate on the research which has been carried out on complex analogies over the last 15 or so years. In it, I will trace the development of theories in this area, the empirical advances made with respect to these theories and the future course of research. So, at one level, the paper is a review of the literature, albeit one which stresses chronology. However, at another level it can be viewed as an exercise in the philosophy of science. For an ancillary aim of the paper is to examine the nature of progress in this area: research on complex analogies can be viewed as a practical example of how research on a cognitive phenomenon, with a strong cognitive science leaning, can advance. One of the hopes of the paper is that this account may act as a template or guide to work on other topics in cognitive psychology.

The Ebb and Flow of Complex Analogy Research

The development of theories of complex analogies can be viewed as consisting of a number of stages or epochs. The main stages I identify are the following:

- Early Theoretical Statements on Analogy
- Clarification of the Stages of Analogical Processing
- Controversies and Counter-Proposals on Analogical Processes
- The Emergence of Consensus in Computational Models

These stages are not strictly sequential. Some occur in parallel; for example, even though the clarification of the stages of analogical processing began in 1980, it is a task which still continues today . Furthermore, these divisions are not the only way to carve up the history of analogy research, but they do emerge as prominent milestones. In the following sections, we expand up on each of these stages and detail the main happenings in each.

Early Theoretical Statements

The starting point of complex analogy research must begin with Dedre Gentner's proposal of structure-mapping theory (see Gentner, 1979, 1982, 1983) . Even though Gentner did not provide much empirical support for the theory, for a number of years, many of her proposals still form the core of analogy theory today. The length of the stride that was taken with structure-mapping theory should not be underestimated. Psychologists always have a problem finding suitable theories when faced with a new area of research. The classic strategy is to use some normative account from another discipline; for example, the use of logic in research on deduction (see Johnson-Laird & Byrne, 1991) or probability theory in decision making (see Eysenck & Keane, 1990, ch. 12). Such theories typically constitute a computational level account of the phenomenon under consideration (see Marr, 1982; Palmer, 1989). □Gentner did not borrow a normative theory from anywhere else, she created one.

There were ideas in philosophy and the philosophy of science about the nature of analogy (see e.g., Hesse, 1966) but these ideas were quite different from Gentner's. In her initial theoretical statement Gentner maintained that parts of a base domain are mapped into a target domain such that:

- relations were transferred; hit(a b) -> hit(c d)
- attributes tend not to be transferred; red(a) -/-> red(c)
- coherent sets of relations are preferred (systematicity principle)

So, for example, in the solar system / atom analogy people map the coherent system of relations -- about the weight difference between the planets and the sun and the attraction of the planets to the sun causing the planets to revolve around the sun -- into the atom domain to suggest that the weight difference between the electrons and the nucleus, combined with their mutual attraction causes the electrons to revolve around the nucleus. And in this mapping the subjects tend not to transfer the attributes of, for example, the sun (i.e., that it is yellow and has a mass of a certain amount).

The core idea in Gentner's statement is the notion that the domains of knowledge have an inherent structure (which can be captured by propositional representations), and that the structure of knowledge plays an important role in determining the results of an analogical comparison. This idea has persisted in until the present day and has, as we shall see, been supported by numerous experiments.

The obvious competitor to Gentner's statement was the work of Gick & Holyoak (1980). At around the same time, they had proposed a theory of analogical problem solving that was rich and complex. It addressed many issues and provided the first concrete empirical support for people's analogical abilities (see next section). However, in my view, Gentner's theory takes precedence over their proposals for a number of reasons. □Gick & Holyoak (1980) had an ill-defined idea of how mapping was done; they talked about parts of both domains being "placed in correspondence" on the basis of an ill-specified matching process. They had some notion that the higher-order structure of knowledge in the domain was important but they did not have the concept of systematicity (which turns out to be very important) and they later abandoned the notion they did have about the role of this knowledge (which alluded to a type of production firing when the conditions in the initial statement of the analogues matched). So, in general, Gick & Holyoak were less explicit than Gentner about the constraints on analogical comparisons, although as we shall see in the next section they made an important theoretical contribution in clarifying the sub-processes in analogical problem solving.

Clarifying the Stages of Analogical Processing

Gentner's initial theoretical statement was a mixed blessing. On the one hand it provided a useful specific, but high-level, description of what constituted an analogical comparison. But, on the other hand, one was left wondering exactly what the psychological context was for this comparative processing. Much of this context was supplied in the early papers of Gick & Holyoak (1980, 1983), taking the form of a specification of the different stages of analogical processing. In clarifying the stages of analogical processing it became clear that Gentner's theory was predominantly about one of a number of processing sub-stages.

Distinguishing the Retrieval and Mapping Stages

Gick & Holyoak approached analogical problem solving from an empirical perspective. In their first paper in 1980, they established that the phenomenon of analogy could be demonstrated in a laboratory situation (although they may have been somewhat pre-dated by Duncker, 1926) .

Consider the following radiation problem which they used in their experiments (adapted from Duncker, 1945):

> 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.

One slightly surprising finding which emerged early on in this research was that subjects did not automatically use the story to solve the problem. Gick & Holyoak found that subjects had to be given a hint to use the story to solve the problem, otherwise they tended not to notice its relevance. In order to notice the story, one has to in some way access it in memory. Hence, prior to mapping the analogy, some retrieval process had to act to access a suitable analogue in long-term memory. The distinction between these processing stages -- of retrieval and mapping -- was an advance beyond the initial statement of analogy theory because it became clear that Gentner was specifically proposing a theory of analogical mapping.

This strong distinction between a retrieval stage and a mapping stage was subsequently demonstrated in several later studies (e.g., Gentner & Landers, 1985; Holyoak & Koh, 1987; Keane, 1987). In Keane's (1987) study, in order to parallel the sort of retrieval problems which subjects face in real-life, the story analogue was encountered in an unrelated context to the analogy experiment (some days prior to testing the story was introduced to subjects in a lecture). Furthermore, the semantic closeness of the story to the problem was systematically varied; one story was about a general using his army to attack a fortress (semantically-remote analogue), another about a general using rays to destroy a missile (semantically-intermediate) and one about a surgeon using rays to destroy a cancer (semantically-close). The basic finding, across a number of experiments, was that the semantically-closer the story is to the problem, the easier it is to retrieve (see Figure 1). However, this trend does not occur in the mapping of the story to the problem; when we look at this measure we find a constant high rate of successful mappings for the different stories irrespective of semantic closeness (see Figure 1). These findings on deep analogies are bad news, because they suggest that even though people may have a suitable analogue encoded in long-term memory, their problem solving performance may suffer because they cannot access it in memory.

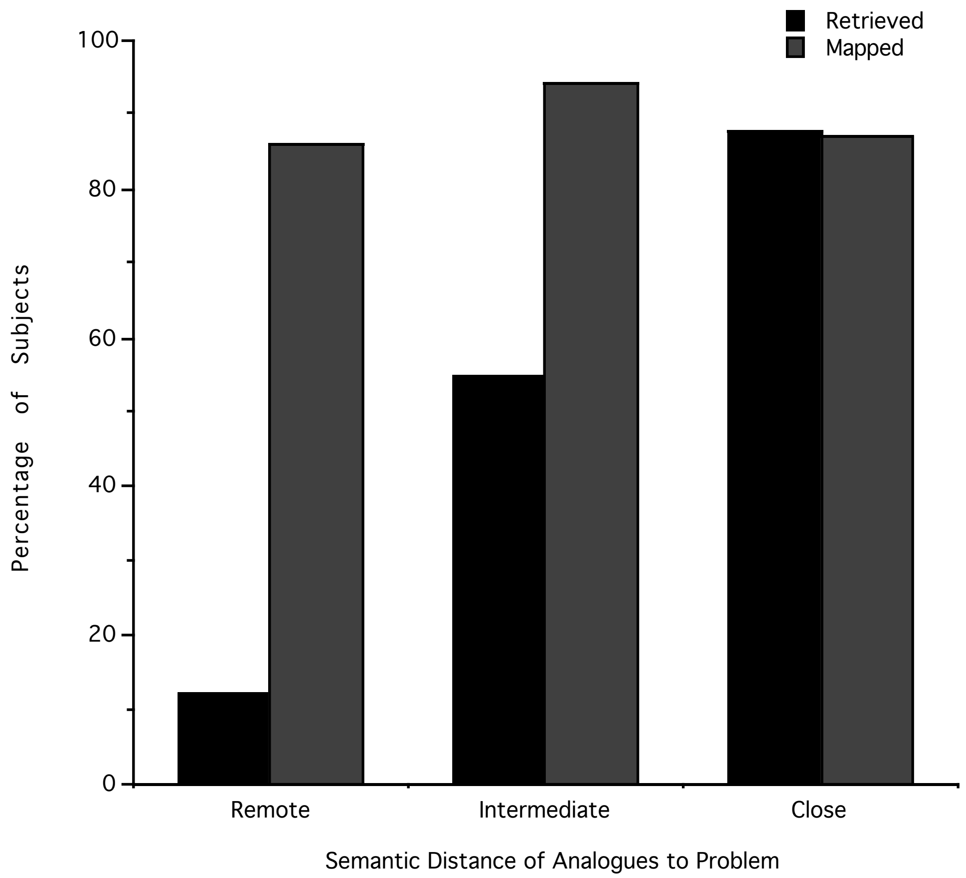

Figure 1 The Percentage of Successful Retrievals and Mappings for the Different Story Analgoues Used in Keane (1987)

Further clarification of the distinction between retrieval and mapping and of the factors affecting retrieval was found by Holyoak & Koh (1987). Using the distinction between surface and structural cues in retrieval (from a similar notion in expertise research; see e.g., Chi, Feltovich & Glaser, 1981), Holyoak & Koh showed that varying surface cues (e.g., whether the story involved the use of lasers or ultra-sound) had a greater impact on the frequency of noticing than varying structural features (e.g., changing a constraint in the story). They also found a separation between retrieval and mapping because structural features have a greater effect than surface features on a subject's ability to use the analogue once a hint has been give as to its utility. In the next section, we shall see that some controversy as to the details of mapping and retrieval were soon to emerge in the area.

Distinguishing Induction as a Stage

Following on the distinction between retrieval and mapping, Gick & Holyoak (1983) showed that analogising involved an induction stage which could have important implications for problem solving skills. In these experiments, Gick & Holyoak first asked subjects to rate the similarity of two stories, before they were asked to solve the target problem. They found that subjects, who were given two analogues to rate, produced more convergence solutions than subjects who were given a single analogue and an irrelevant story (in conditions where no hints were given). Gick & Holyoak also found that the quality of subjects abstract description of the two stories (a measure they called *schema quality*) was a good predictor of whether or not they succeeded in solving the problem by analogy.

Spencer & Weisberg (1986) have shown that these results are sensitive to minor modifications in the experimental paradigm. For example, if a delay is enforced between the acquisition stage (when the stories are presented) and the test stage (when the problem is presented) or if the encoding context of the stories is changed (i.e., if subjects are given them for apparently different reasons) then the success of subjects in the two-analogue condition disappears. These experiments demonstrated that schema quality was not sufficient to produce analogical transfer (see also Gick & McGarry, 1992). More recently, Novick & Holyoak (1991) have provided direct evidence that schema induction is a natural consequence of analogies transfer.

The Importance of Representation as a Stage

From an early stage in research on complex analogies, the processing stage that established the representation of the domains involved in the analogy was important. While this is in one sense obvious -- for example, a story analogue with every second sentence missing will not deliver a very good representation -- there is a deeper sense in which aspects of this stage are crucial. For instance, other variables like expertise could affect the goodness of subjects' representation of a story analogue (see Novick, 1988) or a story analogue might not be considered useful because the problem aspects of the story were not made clear (see Gick & Holyoak's, 1980, discussion of this point).

Much of the work on induction points to the importance of having the right representation with which to analogise. As were saw above, the quality of subjects induced schemas is a very good predictor of subsequent success in an analogical problem solving task. Novick's (1988) study of analogy in arithmetic word problems demonstrated that experts showed better spontaneous transfer than novices when two problems shared structural but not surface features, and found that novices exhibited negative transfer when the two problems shared surface but not structural features. It is thus clear that expertise affects the representation that subjects construct of the problems and, hence, subsequent problem solving performance. Other influences on representation have been shown in the effects of providing diagrams with story analogues and problems (see Beveridge & Parkin, 1987; Gick, 1985).

The Emergence of the Adaptation Stage

The final and most recent processing stage that has been added to the set of sub-processes involved in analogy is that of *adaptation* (Novick, 1988; Holyoak & Novick, 1992) or *validation* (Anderson & Thompson, 1989; Keane, 1988, 1990) or *verification* (Falkenhainer, 1987, 1990). In early accounts of analogy, the assumption had been that when an analogy was drawn between two things, most of the work was carried out by mapping alone. This view is implicit in all initial statements in the literature (see Gentner, 1983; Gick & Holyoak, 1980; Holyoak, 1985). However, it is clear that something more must be done.□

A close examination of many analogues reveals that mapping alone is not sufficient to solve the problem by analogy; that the results of the mapping need to be validated or verified and if needs be *adapted* to make them fit the target problem. For example, consider an alternative version of the General story in which the soldiers inspect their rifles before launching their convergent assault on the fortress. The mapping process would transfer this knowledge to the target, Radiation problem. However, the transfer produced -- which would suggest that the "rays inspect their rifles" -- is clearly ridiculous and inappropriate in the medical domain. So, this solution would have to be adapted by deleting the "inspects" proposition in order to make the rest of the solution workable in the target domain. Some researchers have proposed that a validation (Keane, 1990) or verification process (Falkenhainer, 1990) might spot this sort of sub-optimal solution and present it for adaptation. □

The recognition of the importance of adaptation in analogy has been very late, in sharp contrast to research on case-based reasoning in Artificial Intelligence where it was recognised early on (see e.g., Kolodner, 1988; Riesbeck & Schank, 1990; Smyth & Keane, 1993). For this reason the proposal has received little empirical substantiation. Novick & Holyoak (1991) have, to my knowledge, provided the first empirical demonstration of the distinction between mapping and adaptation.

In a series of analogies to maths problems, they were able to show by the provision of hints to aid mapping, that successful mapping was insufficient for successful analogical transfers and that the major source of transfer difficulty was adapting an analogous solution to the target problem in hand.

Conclusion

As we have seen, one of the major trends in analogy research has been the unpacking of the various sub-processes involved in the phenomenon. At present, there is some consensus on five main stages: representation, retrieval, mapping, adaptation, and induction. One of the notable features of this line of development is the time it has taken to carry out empirical demonstrations of the various distinctions between different stages. Once these empirical distinctions have been made, further work is carried out on specifying the details of a particular stage. This line of development has been best carried out for the analogical retrieval and analogical mapping stages. It has been in these stages, that researchers have suggested counter-proposals on the nature of mapping and retrieval, counter-proposals which we will now consider.

Controversies and Counter-proposals

In the previous section, we saw that one of the first distinctions to be made between processing stages was that between retrieval and mapping. Once this distinction was made, the flow of work moved on to specifying these sub-processes in more detail. Gentner's structure-mapping theory provided an initial statement of the analogical mapping stage and a variety of ideas from memory research provided some basis for an account of retrieval. In this section, we shall see how these two stages were developed through criticisms and empirical work directed at the initial theoretical statements made.

Counter-proposals on analogical mapping

Much of the work on analogy has concentrated on the analogical mapping stage because it appears to be the core, unique processing stage in analogy. In one sense, we have already seen the first assault on Gentner's initial statement of analogy in the proposals made to distinguish various processing stages. This work, provided a processing context for analogising that was lacking in Gentner's initial statement. However, even as a statement of analogical mapping the completeness of Gentner's formulation of analogical mapping was soon challenged (see Holyoak, 1985; Keane, 1985).

Keane (1985, 1988) suggested that analogy was a function of purpose, that structure-mapping did not take the goals of the situation sufficiently into account. He pointed out that two people could take different perspectives on the same domain information and arrive at different analogies. So, even though the knowledge entering into the analogy was the same, different analogical interpretations could be derived because of the goal differences driving the analogising. Keane (1985) also showed that certain attributes -- he termed them functionally-relevant attributes -- played an important role in facilitating analogical comparisons. When subjects were given analogous convergence solutions to the radiation problem involving either troops converging or parcels converging on a central point, the former resulted in more convergence solutions to the problem than the latter. While this finding does not contradict Gentner's proposals -- that attributes tend not to be carried over in analogical comparisons -- it does show that attribute correspondences can facilitate analogical mapping.

Holyoak (1985) advanced similar arguments in specifying a *pragmatic* approach to analogy, as part of a general framework for induction (see Holland, Holyoak, Nisbett & Thagard, 1986). Again, Holyoak argued that goal-related information was important in analogising; that people mapped information that was pragmatically-important. Holyoak advanced empirical support for this view by showing that when the goal structure of a situation was modified people found it harder to form the analogy. So, for example, in Holyoak & Koh (1987), the low-structural versions of the story analogue, which people found hard to map, lacked a constraint in the original story analogues (i.e., the constraint preventing the administration of an intense force from one direction).

There are, however, several problems with this attack on structure-mapping. The main problem is that any manipulation of the goal structure of a story, will change its causal structure. In Gentner's theory, the systematicity principle says that a mapping which preserves higher-order structure should be preferred; so, a mapping involving much causal structure will be preferred over one which lacks such structure. So, Holyoak's arguments on the importance of goals can just be viewed as a special case of the systematicity principle. Furthermore, around this time, the importance of so-called structural factors (essentially the effects of systematicity) received significant empirical support. Gentner & Toupin (1986) gave two age groups of children (5-7 year olds and 8-10 year olds) a story to memorise followed by a set of new characters with whom they had to re-enact the story; hence, they had to map the story about one set of characters onto a new set of characters. The stories were varied in terms of (i) their systematicity or the explicit causal structure and (ii) the transparency of the corresponding characters in base and target. The main finding was that when the object correspondences were hard to establish (lacked transparency) older subjects exploited the causal structure of the story to perform the mapping, whereas the younger age-group failed to exhibit similar performance. More recently, research have shown that the systematic structure of an analogue also plays a significant role in the selection of parts of that domain (see Clement & Gentner, 1991; Gentner & Clement, 1988).

The second main problem with the pragmatic view is that it is really very vague; a vagueness which has been exploited to include further proposals against structure-mapping. Many of these proposals have met with more success. More recently, Holyoak & Thagard (1989) have expanded on the connotations of pragmatic importance to include cases where some aspect of the current context suggests that a particular mapping should be *presumed*. Evidence of this type of pragmatic influence is found by Spellman & Holyoak (1992) using different analogies about Gulf War and World War II or Vietnam. For example, if one of the mappings is presumed, by asking subjects the question "If Saddam is Hitler..." then various coherent mappings follow: then George Bush may be Churchill, in which case the US of 1991 is Great Britain. Furthermore, Keane (1990; also Keane, et al, in press) has stretched the meaning of pragmatic factors further by including "all factors that constitute the pragmatics of the task context" and provided further empirical support for this view. In this view, pragmatic factors may include specific task instructions, various aspects of the task materials (e.g., whether a solution in a story analogue is a successful one or not). So, although what constitutes pragmatic factors seems to be constantly in flux, there does seem to be some evidence for their effects.

So, if we accept that pragmatic factors of some kind, have a unique role that is independent of structural factors, how then are these to be squared with the initial statement of Gentner's theory. Gentner (1989) provides one possible answer by proposing an architecture in which analogising occurs (see Figure 2). The basic position advanced here is that structure-mapping remains intact as a central component in this system, while the various pragmatic effects operate in a peripheral fashion and are the result of a variety of pre- and post-processors. While this account is plausible, it is not the only way to go. As we shall see in the next section, it is possible to develop computational models of analogical mapping that integrate structural and pragmatic factors in single component with no pre- or post-processors.

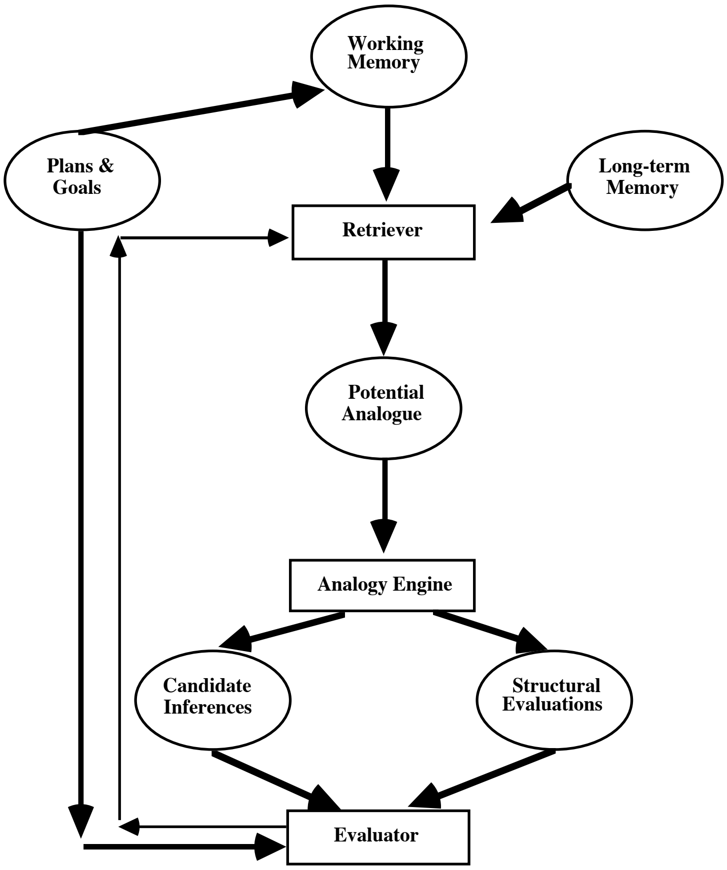

[redacted]

Figure 2 Gentner's (1989) Architecture for Analogical Processing

Controversies in Retrieval

As we have seen, retrieval processes were not a basic concern of the initial statement on analogy made by Gentner. We have already seen several studies that demonstrated that retrieval was distinct from mapping (Gentner & Landers, 1985; Holyoak & Koh, 1987; Keane, 1987); although they acknowledge that the matching of concepts which is carried out for retrieval should provide important information for subsequent mapping processes. Keane's (1987) study showed that varying just the surface features of analogues had a significant influence on people ability to retrieve the story analogue; although in his model of retrieval, which used Schank's (1982) dynamic memory, structural features were seen to have a role in retrieval too. Holyoak & Koh (1987) showed that surface features had a greater impact on retrieval than structural ones. So, both of these accounts saw retrieval as the result of both surface and structural cues.

The controversy in this area surrounds the relative contribution of surface and structural cues to retrieval. If retrieval is heavily dependent on surface cues then this is bad news for human creativity, because it suggests that the human retrieval system is biased towards retrieving superficially-similar items rather than deep, remote analogues. There was empirical research to support such a view.

Gentner & Landers (1985) had shown that in retrieval people tend to rely on surface features more than structural ones. They had subjects memorise stories which varied in their relationship to a target story in the following ways:

- true analogies: which shared the same lower-order relations and causal structure, but involved different objects
- mere appearance matches: which just shared objects and lower-order relations
- false analogies: which shared lower-order relations but had different objects and higher-order relations

They found that mere appearance matches showed the highest rates of retrieval followed by true analogies and false analogies. Recent follow-ups to this research by Gentner, Ratterman, & Forbus (in press) have substantiated these results, by showing that access to long-term memory is heavily biased towards surface features (later we will see how these effects have been modelled).

However, this picture may be more complex than it appears; the pre-eminence of superficial cues is not as straight-forward as it seems. Wharton et al. (in press) have shown that analogues may be retrieved which have few surface features in common with a target; the critical aspect is what the nature of the competition is between analogues in memory.

In conclusion, these results offer us some hope for the future creativity of the human race. While there *are* retrieval problems in finding remote analogues, we can take some succour from the Wharton et al. findings which suggest that if our memories are organized in the right ways then we are likely to happen on a suitable deep analogue which might deliver a solution to a difficult problem we face.

Conclusions on Counter-proposals

Perhaps one of the hallmarks of research on analogical mapping and retrieval is that the controversies which have arisen have not been very controversial. As we have seen much of the dispute has been about the finer details of a given sub-process rather than about radically different formulations of the processes used to achieve a particular task. There is a sense of rapprochement between different researchers in the field. Indeed, this convergence becomes striking when we consider the various computational models which have emerged in the area; models which can be structured within a unifying meta-theoretical framework.

The Emergence of Consensus in Computational Models

Thus far we have seen how the theoretical development of analogy theory has been influenced by empirical results. We have seen how distinctions between various stages were demonstrated in a series of experiments and how various ideas about analogical mapping and retrieval have been supported by experimental tests. Parallel to these developments, beginning around 1986, researchers in this area began to implement the details of various sub-processes in computational models. This move, which has given analogy research a strong Cognitive Science flavour, clarified the predictive focus of various theoretical statements and has formed the basis for an emerging consensus on the important factors underlying analogy. One is tempted to assume that had this computational move not been made, we might all still be arguing over vague ideas.

Computational Models of Analogical Mapping

Three main models of analogical mapping have been developed in the last few years: the Structure Mapping Engine (SME; Falkenhainer, Forbus & Gentner, 1986, 1989), the Analogical Constraint Mapping Engine (ACME; Holyoak & Thagard, 1989) and the Incremental Analogy Machine (IAM; Keane & Brayshaw, 1988; Keane, Ledgeway & Duff, 1991, in press). The latter two models were strongly influenced by SME, so we will consider it first.

The Structure Mapping Engine

The Structure Mapping Engine is an implemented version of Gentner's (1983) structure-mapping theory. It was a ground-breaking program in its illustration of the importance of various structural factors in analogical mapping (see Falkenhainer, Forbus & Gentner, 1986, 1989; Forbus & Oblinger, 1990; Gentner & Forbus, 1991). SME matches up aspects of both domains and tries to find an isomorphic or one-to-one mapping between elements of both domains using the following techniques:

- *it matches the relations in both domains* relational predicates are matched with other identical relational predicates. REVOLVES(A B) in a base would be matched with REVOLVES(C□D) in a target, BUT the REVOLVES predicate would never be matched with SPINS(C, D) or with the object C.

- *it exploits 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 et al., 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.

These techniques have been shown to be very powerful even if they are very sensitive to the representations used. SME demonstrated in a very concrete way the sort of computational problems which people solve when mapping analogies. For instance, it demonstrated that even for apparently straight-forward analogies (like the atom / solar system or General story / Radiation problem analogies) there are a large number of matches to be considered and many possible analogical interpretations to be chosen between. The key insight is that systematicity turns out to be a very important principle in deciding which of these interpretations is optimal. Other statements in the initial theory have also been specified in the model. For example, Gentner originally stated that "attributes tend not to be carried over in an analogical comparison". This is realised in the model by match rules that exclude matches between attributes, unless they are placed in correspondence by the computation of structural consistency.

SME has supported more detailed empirical tests of systematicity (see e.g., Gentner & Toupin, 1988; Clement & Gentner, 1991; Gentner & Clement, 1988). It has also been used to generate predictions for people's classifications of analogical and other comparisons (see Skorstadt, Gentner & Rattermann, 1987). As such, it has become a cornerstone of analogy research and an important first implementation that has influenced all subsequent, competing models.

The Analogical Constraint Mapping Engine

The Analogical Constraint Mapping Engine by Holyoak & Thagard (1989) was designed to go beyond SME, as it takes into account pragmatic factors (and has a more flexible treatment of similarity). However, it is essentially a parallel implementation of the structural factors underlying SME (with some added functionality for pragmatic and similarity factors).

ACME uses parallel constraint satisfaction in an interactive network to find the optimal mapping between two domains. It establishes a network of units or nodes. Each node represents a legal match between two predicates in the base and target domain. Excitatory and inhibitory connections are formed between these nodes to implement the various factors that influence the mapping; for instance structural factors or *structural constraints* as they are termed in the model. So, for example, a match between REVOLVES(A, B) and REVOLVES(C, D) involves nodes representing the matches between REVOLVES=REVOLVES, A=C, A=D, B=C and B=D. To implement structural consistency, there are excitatory links between the REVOLVES=REVOLVES node and the A=C and B=D nodes. To enforce a one-to-one mapping there are inhibitory links between the REVOLVES=REVOLVES and other competing match nodes like REVOLVES=ATTRACTS nodes, and A=D and B=C. Node-matches which are semantically similar are given extra activation from a separate *semantic unit* outside the network and node-matches that are pragmatically-important are also given extra activation from a *pragmatic unit*. When the network has been constructed, it is run until the activations of the nodes settle into a stable state. The nodes whose activation exceed a threshold correspond to the optimal set of matches between the two domains. So, in one sense, ACME produces just one mapping, a single optimal interpretation.

ACME was an important contribution to analogy research for a number of reasons. First, Holyoak & Thagard have shown that aspects of the model's performance correspond well to the empirical data; the number of cycles the network goes through before it reaches the correct mapping corresponds well to the frequencies of correct solutions found among subjects in a wide range of analogical problem solving experiments (e.g., Holyoak & Koh, 1987). Second, it was an important demonstration that a phenomenon which was generally classified as "thinking", could be modelled in a parallel fashion. Third, it showed that pragmatic factors could be integrated into an analogy model in a central fashion, without pre- and post-processors (as Gentner, 1989, had suggested). Recently, some of the effects observed in ACME have been captured in a more neurally plausible network which uses dynamic binding (see Hummel & Holyoak, 1992).

The Incremental Analogy Machine

The Incremental Analogy Machine (Keane & Brayshaw, 1988; Keane, 1990; Keane et al., 1991, in press) was designed as a counter-point to aspects of SME and ACME. Its starting point is the observation that people are limited-capacity processors, partly as a result of working memory limitations. Therefore, it is unlikely that they have the capacity to entertain many possible alternative interpretations of a mapping. So, IAM uses a heuristic type of mapping that reduces the amount of processing needed to be carried out.

IAM uses serial constraint-satisfaction that takes into account structural, pragmatic and similarity factors. 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 which will tend to be the optimal interpretation. However, if it has to, IAM can consider several alternative mapping 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 (see Keane et al., 1991, for more details).

IAM is a model that tries to take the behavioural limitations of people into account, in a manner that is not captured by SME and ACME. First, it produces a single mapping based on a heuristic best guess, thus reducing the processing that needs to be done to make the analogy (see Forbus & Oblinger, 1990, for a similar modification to SME). Second, it tries to approximate people's phenomenal experience of being able to try one analogical interpretation and then another in a serial fashion. Clearly, the parallel nature of ACME does not capture this, while SME could capture it but has no explicit mechanism for doing so. Furthermore, there is direct empirical evidence to support this type of model using the following mapping problem (from Holyoak & Thagard, 1989) :

A	B
Bill is tall.	Fido is hungry.
Bill is smart.	Blackie is friendly.
Tom is tall.	Blackie is frisky.
Tom is timid.	Rover is hungry.
Steve is smart.	Rover is friendly.

In this task, subjects are asked to say which things in list A correspond to which things in list B (ignoring the meaning of the words). Essentially, subjects have to discover a one-to-one mapping between all the individuals and attributes in list A and list B. The unique one-to-one mapping which solves the problem is to match Steve and Fido, Bill and Rover, Tom and Blackie, smart and hungry, tall and friendly, and timid and frisky. IAM predicts that the order in which the information is given can affect the ease of analogical mapping. This prediction is borne out by empirical tests which show that when one of the sentences in the list A ("Steve is smart") is in first as opposed to the last position in the list, the mapping problem is solved twice as fast by subjects.

Consensus on Analogical Mapping

The development of computational models of analogical mapping has been a very useful advance for analogy research. First, it has helped to specify the ideas that were previously stated in purely linguistic terms. Second, it has allowed more specific predictions to be made about the nature of analogising, many of which have been empirically tested and confirmed. Third, it has resulted in the emergence of a certain level of consensus in the area. That is, there is now considerable agreement about what factors are important in analogising; the dispute is really about the finder details of how analogical mapping is carried out. This consensus has been crystallised in a meta-theoretical framework proposed by Keane et al. (in press) which illustrates the unity of the various theories.

A Meta-theoretical Framework of Analogical Mapping

In recent years, cognitive scientists have come to appreciate the need to stratify their theoretical proposals within a meta-theoretical framework. Marr's (1982) framework distinguishes between three levels of theory (i.e., the computational, algorithmic, and hardware levels). In several areas of cognitive science, most notably vision research, it has proved possible to elaborate theories at each of these levels (see also Johnson-Laird & Byrne, 1992; Palmer, 1989).

Palmer (1989) has pointed out that any adequate theory of analogical mapping will have to operate at several levels of description. At the highest level, one needs to characterise the *informational constraints* implied by 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 is one of *behavioural constraints* which 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* which aim to capture the neurological primitives of analogical thought (Marr's hardware level).

Theories of analogy can be divided into statements about informational constraints and behavioural constraints. Holyoak & Thagard's (1989) description of the constraints involved in ACME capture the informational constraints:

- *structural constraints* - all those constraints that ensure a one-to-one mapping is achieved (based on structural consistency, matching of types and systematicity)
- *similarity constraints* - which determine concept matching on the basis of semantic similarity or identity
- *pragmatic constraints* - which determine concept matching on the basis of pragmatic importance or task demands which affect "the pragmatics of the situation"

Keane et al. (in press) list the following behavioural constraints:

- working memory limitations - which enforce some reduction of the processing load on analogising
- background knowledge - if a set of transfers are consistent with background knowledge then they will be made more easily

This meta-theoretical framework illustrates nicely the overall differences between analogy theories and models. SME implements the structural and, to a limited extent, the similarity constraints and pragmatic constraints (for the latter see Forbus & Oblinger, 1991). ACME and IAM make explicit use of all three informational constraints. IAM is the only model at present that takes behavioural constraints into account. Although, modifications are currently being made to SME to take these constraints into account (Forbus, 1983).

Computational Models of Retrieval

We have already seen that the tension in the empirical findings on analogue retrieval hinged on the extent of the relative influence of surface and structural cues. Two salient results recur in the literature. First, the deep analogues, which are the basis of creative discoveries, that have many structural cues and few surface cues tend to be retrieved infrequently (see e.g., Keane, 1987). Second, superficially-similar items, which bear no deep relationship to a target problem, that shared many surface features tend to be retrieved quite easily (see e.g., Gentner, Ratterman & Forbus, in press; Ross, 1987). Third, literally similar items which share both surface and structural cues tend to be retrieved the most (see Gentner, Ratterman & Forbus, in press).

Computational models have been developed that attempt to capture these phenomena and to explain how such apparently paradoxical results can occur. The two main models proposed are the ARCS model (Thagard, Holyoak, Nelson & Gochfeld, 1990) and the MAC/FAC model (Gentner & Forbus, 1991).

ARCS (Thagard, et al., 1990)

The Analogical Retrieval by Constraint Satisfaction (ARCS) model by Thagard et al. (1990) is an outgrowth of the ACME mapping model. The close relationship of the two models is partly a response to the intuition mentioned earlier that retrieval processes must establish correspondences which are used in the mapping process. ARCS uses the same three constraints as ACME -- structural, similarity and pragmatic constraints -- but their relative importance changes. In particular, semantic similarity constraints are seen to be of greater importance in retrieval.

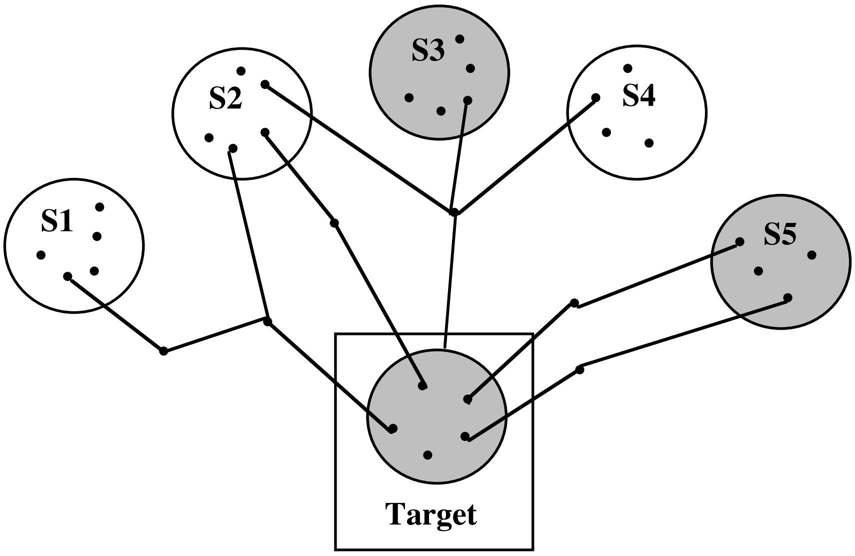

Figure 3 Schematic View of the Semantic Links in ARCS Between a Target Structure and Other Structures Stored in LTM (darker structures have greater activation)

ARCS carries out retrieval in two steps. First, the concepts in the target problem are activated and the initial search process finds pathways from these concepts to similar concepts in candidate source analogues in long-term memory (see Figure 3). The semantic similarity between these concepts is determined from a separate semantic data base called WordNet (see Miller, Fellbaum, Kegl, & Miller, 1988) organized in terms of synonyms, antonyms, superordinates, subordinates and part-whole relations. In the second processing step, the influence of pragmatic and structural constraints are computed for each putative analogue and a higher-order mapping unit is established to represent the correspondence between the target and a particular analogue. These units thus capture the competition between different putative analogues to the target, a competition which is resolved when the highest activation "winning" unit emerges. It is assumed that the emergence of the winner corresponds to subjects' phenomenological experience of consciously accessing a stored base analogue.

The way in which the different constraints are applied in ARCS ensures that semantic similarity has a very strong effect on retrieval but that structural features can also play a role in discrimination between candidate retrievals (as was shown in Holyoak & Koh, 1987). However, ARCS does not fare well on the specific ordering of different types of retrievals found in the experimental literature (see Gentner & Forbus, 1991). Related research on competition in retrieval has been carried out a model in an alternative connectionist model called REMIND (see Wharton et al. 1991, 1992, in press).

The MAC/FAC Model (Gentner & Forbus, 1991)

The main competitor to the ARCS model is the Many Are Called/Few Are Chosen (MAC/FAC) model proposed by Gentner & Forbus (1991; see also Ratterman, Gentner & Forbus, in press). It also has a two-step process in which possible correspondences between a target and source analogue are computed by matching flat, content-vector representations of the target and base domains in memory. These flat representations are vectors that include the predicates in a domain but nothing of the inter-relating structure between these predicates. This first step isolates a number of possible candidate analogues (the MAC stage) which are subsequent chosen between by carrying out a structure-mapping between the target and each candidate (the FAC stage; as this is computed by SME, see previous section).

The outputs of this model correspond well to the sorts of performance people have manifested in numerous studies. Furthermore, the MAC stage is computationally more plausible as an account than the extensive search carried out by ARCS. It is also more accurate in its predications than ARCS.

Conclusions on Computational Models of Retrieval

As in the case of the model of analogical mapping, the development of models of retrieval has been an important advance in analogical mapping. These models have allowed researchers to move to finer levels of detail in their theoretical formulations and hence to produce more constrained experiments. Evidence of this can be gleaned by a comparison of the sorts of manipulations carried out in earlier and later empirical work in the area. The models have also helped to explain apparent contradictions in the empirical phenomena. Finally, there is now a corpus of test simulations in the literature which any new model has to meet. Therefore, we have a clear basis for assessing change and progress in the future.

Final Comments & Future Directions

The previous sections of this paper have attempted to chart the progress of research on analogy from the initial statements of theory in the area to the construction of detailed, competing, computational models of key phenomena brought to light in empirical tests. □As a relatively new and circumscribed area, analogy research provides an interesting example of the type of development that can occur in modern psychological research. From a philosophy of science perspective, we would suggest that it provides an interesting template against which to understand developments in Cognitive Science. From a methodological perspective it could act as a guide to researchers facing other new topic areas. In the current flux of disciplines which is called Cognitive Science, clear methodological signposts do not exist to guide researchers. I hope that the story told here might remind other researchers of the current stage of development of their areas and suggest future lines of advance.

Having said this analogy research is not to be held up as a shining example of "getting everything right". While some progress has been made since 1979, and some productive consensus has been reached between researchers there is still a lot to be done. In the following sub-sections, I will reveal something of the scruffy underbelly of the area, of what remains to be done in the area.

Competitive Tests between Models

In some respects, analogy research in an unusual area in the sense that there are several very specific computational models available. This raises the possibility of making competitive tests between them. However, this is not as straight-forward a matter as it first seems. The different models often use different measures as indicators of human performance, they use different representations in the examples they use and have radically different computational architectures (see Keane et. al., in press, on these issues). So, carrying out competitive tests between them is not clear cut. Indeed, what is needed to some sort of methodology for performing competitive tests and some agreed upon grounds for representing examples and measuring the outputs of the different models. One initial attempt at doing this is given in Keane et al. (in press).

Furthermore, we also need some way of assessing the relative goodness of different models. As more and more empirical tests are carried out, I am sure that no one model will capture all of the phenomena. Some consensus will have to emerge on what are the most important phenomena and how these are to be weighted in the evaluation of different models.

Other Stages Remain to be Understood and Modelled

While much research has been carried out on analogical mapping and analogue retrieval, a lot less has been carried out on the other sub-stages of analogy: representation, adaptation and induction. Indeed, at present, we have really only demonstrated that these stages occur. Detailed statements on how they are carried out are as developed as proposals on mapping were in 1980. Some initial steps have been taken in this direction, but more need to be done (see e.g., Anderson & Thompson, 1989; Falkenhainer, 1987, 1990; Keane, 1990).

Furthermore, we do not know what the relative importance of the different stages on analogising in different task situations. For example, current accounts place all the importance on analogical mapping, yet adaptation could be the critical stage at which an analogy succeeds or fails. Novick & Holyoak (1991) have demonstrated some cases where adaptation is more important to the successful solution of a problem than mapping.

Strategic Aspects of Analogy

A final problem which has been raised by Keane et al. (in press) is that of strategic aspects of analogy. In the attribute-mapping problem (mentioned earlier) it is clear that subjects try one analogical mapping and then, if it fails, try another. Most models, apart from IAM, do not have any mechanism for dealing with this sort of strategic mapping. Keane et al. demonstrate one case where this occurs but we do not have a good idea when or how often people map analogies in this way.

References

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.

Beveridge, M. & Parkins, E. (1987). Visual representation in analogical problem solving. *Memory & Cognition, 15*, 230-237.

Carbonell, J.G. (1983). Learning by analogy: Formalizing and generalising plans from past experience. In M.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach.* Amsterdam: Springer-Verlag.

Chi, M.T.H., Feltovich, P.J. & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121-152.

Clement, C.A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science, 15*, 89-132.

Duncker, K. (1926). A qualitative (experimental and theoretical) study of productive thinking (solving of comprehensible problems). *Journal of Genetic Psychology*, *68*, 97-116.

Duncker, K. (1945). On problem solving. *Psychological Monographs*, *58* (Whole No. 270).

Eysenck, M.W., & Keane, M.T. (1990). *Cognitive Psychology: A Student's Handbook*. London: Lawrence Erlbaum.

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. (1983). Personal communication.

Forbus, K.D. & Oblinger, D. (1990). Making SME greedy and pragmatic. *Twelfth Annual Conference of the Cognitive Science Society*. Hillsdale: Erlbaum.

Gentner, D. (1979). The role of analogical models in learing scientific topics. *BBN Technical report*.

Gentner, D. (1982). The role of analogy in science. In D.S. Miall (Ed.), *Metaphor: Problems and Perspectives* Sussex: Harvester Press.

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., & Clement, C.A. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G.H. Bower (Ed.), *The Psychology of Learning and Motivation, 22*. New York: Academic Press.

Gentner, D., & Forbus, K.D. (1991). MAC/FAC: A Model of similarity-based retrieval. *Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale: Erlbaum.

Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man and Cybernetics* , Tucson, Arizona, November.

Gentner, D. & Rattermann, M.J., & Forbus, K. (in press). The role of similarity in transfer. *Cognitive Psychology*.

Gentner, D. & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*, 227-300.

Gick, M.L. (1985). The effect of a diagram retrieval cue on spontaneous analogical transfer. *Canadian Journal of Psychology*, *39*, 460-466.

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.

Gick, M.L., & McGarry, S.J. (1992). Learning from mistakes. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *18*, 623-639.

Hesse, M. (1966). *Models and Analogies in Science*. Notre Dame: University of Notre Dame.

Holland, J.H., Holyoak, K.J., Nisbett, R.E. & Thagard, P. (1986). *Induction: Processes in inference, learning and discovery*. Cambridge, MASS: MIT Press.

Holyoak, K.J. (1985). The pragmatics of analogical transfer. *The Psychology of Learning and Motivation*, *19* , 59-87.

Holyoak, K.J. & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition* , *15*, 332-340.

Holyoak, K.J., & Thagard, P.R. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13,* 295-355.

Hummel, J.E. & Holyoak, K.J. (1992). Indirect analogical mapping. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Bloomington, Indiana.

Johnson-Laird, P.N. & Byrne, R.M.J. (1992). *Deduction*. London: Lawrence Erlbaum.

Keane, M.T. (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. (1988a). *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., & 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*.

Kolodner, J. (1988) (Editor). *Proceedings of the Case-Based Reasoning Workshop*,San Mateo, Calif: Morgan Kaufmann.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

Miller, G.A., Fellbaum, C., Kegl, J., & Miller, K. (1988). WordNet: An electronic lexical reference system based on theories of lexical memory. *Princeton University Cognitive Science Laboratory Report 11*.

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.

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.

Riesbeck, C. K., and Schank, R. C. (1989). *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates.

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 & Cognition, 13*, 629-639.

Schank, R.C. (1982). *Dynamic Memory* . Cambridge: Cambridge University Press.

Skorstadt, J., Falkenhainer, B., and Gentner, D. (1987) Analogical processing; A simulation and empirical corroboration. *Proceedings of AAAI'87*. American Association of Artificial Intelligence.

Smyth, B., & Keane, M.T. (1993). Retrieving adaptable cases. *Proceedings of the First European Workshop on Case-Based Reasoning*, Kaiserslauten, Germany.

Spellman, B.A., & Holyoak, K.J. (1992). If Saddam is Hitler then who is George Bush ?: Analogical mapping between systems of social roles. *Journal of Personality & Social Psychology, 62*, 913-933.

Spenser, R.M., & Weisberg, R.W. (1986). Context-dependent effects on analogical transfer. *Memory & Cognition*, *14*, 442-449.

Thagard, P, Holyoak, K.J., Nelson, G. & Gochfeld, D. (1990). Analogue retrieval by constraint satisfaction. *Artificial Intelligence, 46*, 259-310.

Weitzenfeld, J. & Klein, G. (1979). Analogical reasoning as a discovery logic. Klein Associates, TR-SCR-79-5, Yellow Springs, Ohio.

Wharton, C.M., Holyoak, K.J., Downing, P.E., Lange, T.E. & Wickens, T.D. (1991). Retrieval competition in memory for analogies. In K.J. Hammond & D. Gentner (Eds.),*Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Wharton, C.M., Holyoak, K.J., Downing, P.E., Lange, T.E. & Wickens, T.D. (1992). The story with reminding: Memory retrieval is influenced by analogical similarity. In J.K. Kruschke (Ed.), *Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Wharton, C.M., Holyoak, K.J., Downing, P.E., Lange, T.E. & Wickens, T.D., & Melz, E.R. (in press). Below the surface: Analogical similarity and retrieval competition in reminding, *Cognitive Psychology*.