UNDERSTANDING ANALOGICAL REASONING:
Viewpoints from Psychology and Related Disciplines[1]

Rogers P. Hall

Irvine Computational Intelligence Project
Department of Information and Computer Science
University of California, Irvine, 92717

May 29, 1986

Technical Report 86–10

Abstract

Analogy and metaphor have a long history of study in linguistics, education, philosophy and psychology. Consensus over what analogy is or how analogy functions in language and thought, however, has been elusive. This paper, the first in a two part series, examines these various research traditions, attempting to bring out major lines of agreement over the role of analogy in individual human experience. As well as being a general literature review which may be helpful for newcomers to the study of analogy, this paper attempts to extract from these literatures existing theories, models and concepts which may be interesting or useful for computational studies of analogical reasoning.

[1] This work was supported in part by ONR grant N00014-85-K-0373

Contents

1. Introduction

While metaphor and analogy have a long–appreciated history in the arts, the same has not always been true in traditional academic approaches to the subject. As noted by Johnson (1981), prior to the last two decades metaphor was consistently ignored in "respectable" philosophical discussion. Since the 1960's, however, there has been an explosive growth in research on metaphor and analogy, both in philosophy and other disciplines, to the point that we are currently experiencing "metaphormania" (Johnson, 1981, p. *ix*). This paper is the first in a two part series which reviews portions of this growing literature in linguistics, education, philosophy of science and psychology. This literature is examined both to achieve a general understanding of the subject matter and to determine what these traditional research disciplines might contribute to computational studies of metaphor and analogy. By far, psychological studies of metaphor and analogy appear to be the most promising contributors, and a review of that work forms the bulk of this paper. The second paper in this series focuses on computational approaches to analogical reasoning which, by comparison to other disciplines, are just beginning.

Most studies of metaphor or analogy begin by pointing out the prevalence of figurative language and reasoning in cognitive experience. Analogical reasoning appears to occur in even the most mundane forms of human cognition. The widespread occurrence of figurative language consisting either of novel or "frozen" (so commonly accepted as to appear literal) metaphors is evident in virtually any sample from the news media (Schōn, 1979; Carbonell, 1981) and even in scientific writing (Reddy, 1979). Analogical reasoning has long been considered a key constituent ability of the human intellect (Sternberg, 1977) and has figured heavily in considerations of a variety of cognitive activities from problem solving (Polya, 1954; Dreistadt, 1969) to creative thought more generally (Boden, 1978). The use of analogy in teaching has a long history, with the presentation of new concepts by comparison with analogically related older concepts comprising the bulk of didactic material in textbooks. Hence, widespread research interest in such a commonplace form of cognition is understandable.

1.1. Terminology

To reduce confusion, this paper outlines terminological conventions before surveying different academic approaches to metaphor and analogy. As with research into most areas of language and thought, terminology in the study of analogical reasoning varies across differing research approaches. In addition, the meaning of terms in the context of research differs from commonplace usage. This is certainly true of terms like **analogy**, **metaphor**, and **reasoning** as appended to either of the first two terms. In fact, the meaning of these terms constitutes a major research question in many of the studies surveyed in this paper.

Common academic usage outside the field of artificial intelligence (AI) takes metaphor in the most general sense to refer to the juxtaposition of dissimilar domains of knowledge, with analogy being a special case of metaphor in which particular elements of these domains are explicitly placed in correspondence. Hence, metaphorical reasoning refers to the process of discovering (if necessary) and then articulating points of correspondence between domains.

As conventionally used in the AI literature, however, these terms have taken on somewhat different meanings, metaphor being associated with figurative *language use* and analogy being used more broadly to describe a type of reasoning about literally dissimilar domains of knowledge. For convenience and, hopefully, without violating custom unnecessarily, the terms **metaphor** and **analogy** will be used interchangeably, each referring to the general notion of juxtaposing domains of knowledge which, given a successfully elaborated basis of correspondence, can be viewed as being similar in some respects. More specific definitions or terminology will be introduced as necessary in describing individual studies.

With respect to the domains of knowledge involved in analogical reasoning, some additional terminology will be required. Typically one domain, referred to as the **source** or **vehicle** in this paper, is well understood relative to the other domain, termed the **target**, **tenor**, or **topic**. The correspondence established during analogical reasoning will be termed the **ground** and provides a mapping through which knowledge of the source can be imported into the target domain. The specific nature of such a mapping provides the focus of most studies surveyed here. For example, a skilled spades player might approach the card game of hearts by recognizing that many of the strategies applicable in spades (e.g., preferring hands with many cards of the spade suit and attempting to take as many tricks as possible) can be applied in an inverse fashion in the game of hearts. Knowledge of playing spades serves as the source, hearts the target, and inverted strategies the ground in this example of analogical reasoning.

2. Metaphor and linguistics

Unfortunately, figurative language has generally received uneven attention in linguistics. Typically, metaphor is explained away as unnecessary ornamentation or as an aberration in an otherwise well-behaved linguistic repertoire. Chomsky (1965), for example, characterizes metaphor as the violation of semantic selection rules arising from the literal incongruity of tenor and vehicle. In this account metaphors are comprehended through substitution of well-formed, literal sentences which conform to selection rules. Hence, figurative language is a mere substitution of non-literal for literal expression and becomes a matter of performance rather than linguistic competence.

Much recent work on metaphor, particularly from a psychological orientation (Honeck, 1980), is in direct opposition to this narrow *substitution* view. Alternative conceptions of metaphor comprehension focus more closely on discovery of the ground relating tenor and vehicle, and allow a larger role for metaphor in both language use and cognition generally. Much of what is salient in this wider consideration of analogical reasoning can be traced to Black's (1962, 1979) differentiation between comparison and interaction views of metaphor. *Comparison* views stress implicit but still literal contrast between descriptive features of topic and vehicle, while *interaction* views focus on a more broad juxtaposition of systems of knowledge associated with tenor and vehicle. In the latter case, the ground supporting metaphorical juxtaposition of tenor and vehicle may involve figurative as well as literal relations. For example, men are like wolves by virtue of predation, but this relation must be interpreted figuratively since (most) men do not hunt in packs, etc. As a result of elaborating the ground in this fashion, entirely new meanings may arise for both tenor and vehicle domains.

Estimated to occur at a rate of 4 per minute in natural language discourse (Pollio, et al. 1977) and with a frequency of 5 per 1000 words in educational texts (Ortony, 1976), metaphors form a rather bulky portion of our linguistic output to simply explain away and have prompted many to claim that any comprehensive theory of linguistic competence must include a careful treatment of figurative language use. The frequency with which figurative language occurs also suggests that metaphor may play both a powerful and necessary role in bringing language into contact with perceptual experience.

From a pragmatic perspective, metaphor entreats the listener to view or understand some topic in a particular way and hence plays a powerful role by establishing a context suggestive of the speaker's intent. Beyond facilitating individual discourse, metaphor provides a fundamental, culturally–shared perceptual framework for experience. That is, there are inter–individual metaphors that we "live by" (Lakeoff and Johnson, 1980) in the sense, for example, that positive valuation might be expressed spatially as in "We were up for the occasion." The power of metaphor is made more concrete in a provocative paper by Schön (1979) in which generally accepted cultural metaphors are shown to provide a means of "seeing–as," suggesting which aspects of a situation are salient and what sorts of solutions might apply. For example, inner–city slums are frequently described as "urban blight," thus likening unpopular social communities to an agricultural disease amenable to diagnosis and treatment, in this case by social technology aimed at eradicating an aberrant growth. In this way, a metaphor constrains consideration of its tenor, suggesting some particular interpretation while obscuring others. From an alternative perspective, inner–city slums might be viewed as natural communities, to be supported by encouraging positive community growth. With some thought this *problem-setting* function of metaphor can be detected all around us: mental illness, cognition as computation, people as property, and surging psychological forces, to name a few. It is difficult to deny that these metaphors have played, and continue to play, a powerful role in human affairs.

In terms of necessity, metaphor has been described as extending the semantic denotation of a fairly rigid vocabulary to cover our more fluid range of experience (Ortony, 1975; Black, 1979). Briefly, metaphor is seen to provide a mechanism by which common word senses are extended to include novel experiences rather than creating new words out of descriptive necessity. A simple example might be the figurative expression, "leg of a triangle," which has become so commonplace as to be considered a literal statement. Such dead or frozen metaphors are suggestive artifacts of the dynamic, polysemantic nature of our language use. Hence new experiential entities *must* be seen as more familiar entities if our symbolic representation of the world is to remain tractable.

3. Metaphor in education

Despite the prevalence of metaphor in educational writing, a heated debate has developed within the field of educational psychology as to the desirability of instruction by analogy. Ortony (1975) argues that metaphor provides three features both necessary and desirable in education (and, more generally, in language use):

1. metaphors serve as a compact set of instructions for reconstructing experience;
2. metaphors allow us to express experiences which our current language cannot literally describe;
3. and metaphors provide a vivid communicational medium closer to our perceptual, cognitive and emotive experience.

Hence metaphors allow us to speak of things which we have experienced but cannot literally express and to do so in a compact and vivid fashion. Assuming a recipient's knowledge of the vehicle and an adequately specified ground, Ortony claims that metaphors enhance understandability and memorability of learned material.

This facilitative view of educational metaphor has been challenged. Miller (1976) takes exception to Ortony's enthusiasm by arguing that metaphors in educational writing "gloss over" an insufficient command of difficult concepts on the part of the teacher or writer and provide an oftentimes emotionally manipulative basis for specious arguments. Specifically, the compactness of metaphorical expression leads to ambiguity of interpretation on the part of a heterogeneous audience, and figurative expression of the "inexpressible" leads to vagueness and loss of clarity. Ortony (1976) counters with an inventive metaphor involving misused tools, arguing that the use of screwdrivers in violent criminal acts should not preclude the use of screwdrivers in carpentry.

In an effort to subject these arguments to empirical verification, Reynolds and Schwartz (1983) have tested subject's recall for short text passages concluded with either a metaphorical summary (e.g., "the sheep followed the leader over the cliff" concluding an historical account of Nazi Germany) or its literal equivalent. Results indicate that metaphorical summarization enhances cued recall of text passages (allowing for paraphrasing during recall), although passage novelty also has an appreciable effect on memorability. The authors cautiously conclude with alternative explanations: increased recall may be attributable to additional processing during comprehension or to the formation of a specific memory on the basis of a metaphorical summary which later guides reconstruction of the original text. In either case, the argument that metaphorical explication may facilitate comprehension and memorability of new material appears to have some support, although specific suggestions for the use of metaphor or analogy in teaching are not drawn.

Beyond the pragmatic utility of metaphor in education, many have argued strongly for its necessity in learning anything which is truly novel. Elements of this view in Ortony's (1975) description of metaphor as flexibly extending discrete symbol systems to continuous life experience are expanded by Haynes (1975). She argues for an interactive level (i.e., insight or discovery) of metaphor comprehension which involves hypothesis generation and experimentation by the recipient of the metaphor. Hence metaphor is seen to provide a "creative leap" in which new insights modify existing conceptual structures of both topic and vehicle. Petrie (1979) supports this view, arguing for the necessity of metaphor in "leaping the epistemological chasm between old knowledge and radically new knowledge" (p. 440). Metaphor, through its literal incongruity, presents an anomaly to be resolved by a learner engaged in active hypothesis generation and verification.

Quite a similar debate over the utility of metaphor in educational settings has surfaced in regard to learning to use computing systems. Halasz and Moran (1982) argue that open-ended metaphors tend to confuse computing novices, particularly to the extent that multiple, incompletely specified analogies are presented to accommodate various system features. What results, according to their analysis, is a "baroque collection of special-purpose models pasted together in a more-or-less integrated and more-or-less consistent fashion" (p. 34). Instead they argue for the presentation of carefully tailored conceptual models which reflect software system design. From this viewpoint, metaphor is used at most as a stylistic device for presentation of particular aspects of the model, but is to be discarded once a concept has been learned. Unfortunately, in addition to ignoring the active, experimentation inducing properties of educational metaphors, these authors give meager advice on how to construct teachable conceptual models for existing or proposed software architectures.

In contrast, Carrol and Mack (1982) describe a role for metaphor quite similar to the interactive level mentioned earlier. Noticing that computing novices oftentimes use spontaneously generated analogies, these authors advocate designing system interfaces which take advantage of these analogies and, hence, facilitate training. Instead of presenting static conceptual models directly to the novice, they advocate the presentation of open-ended metaphors to encourage a process of active learning in which the learner incrementally extends the metaphorical ground through experimentation with the computing system. As with Halasz and Moran, the ultimate goal of learning is a valid conceptual model of some computing system. However, Carrol and Mack argue strongly for the importance of metaphor as a component in active learning.

In summary, there are two broad classes of opinion concerning the role of metaphor in education, and these views parallel the controversy described for linguistics. One view takes metaphor to be literally comparative in the sense that essential descriptive qualities could be provided by literal language as well. Although proponents of this view allow metaphor some utility for making a point vividly, the attendant dangers of being vague or misleading recommend careful literal expression rather than metaphor in teaching. As described metaphorically by Petrie, this view claims that "metaphors have all the advantage over explicit language as does theft over honest toil" (p. 438). In contrast, others describe metaphor as being not only desirable but essential. From this perspective, metaphor provides a context for active experimentation on the part of the learner and a means for accommodating truly novel experiences into existing conceptual structures.

4. Metaphor in the philosophy of science

A similar dichotomy of opinion occurs in the philosophy of science with respect to the utility of metaphor in scientific theory construction. A strong empiricist bias against figurative language as an imprecise medium for communication is counterbalanced by an appreciation for the central role that open-ended metaphors play in scientific theorizing (Hoffman, 1980). Beyond their explanatory utility, metaphors serve as components of a theory which are not yet literally expressible, and through such a *theory constitutive role*, they provide a new way of viewing the topic under investigation by focusing research activity around elaboration of the metaphorical ground.

Boyd (1979) gives a revealing discussion of the role of computing metaphors in cognitive psychology where functional relations among computational states in machines serve as a vehicle for thinking about hypothesized psychological states and the interrelationships among them. According to Boyd, such a metaphor provides "at a relatively early stage theoretical terms to refer to various plausibly postulated computerlike aspects of human cognition, which then become the objects of further investigation" (p. 370). Notions of processing, feedback or retrieval as understood in computing can be used to generate falsifiable hypotheses about cognitive constructs and, hence, provide an active research program in addition to new descriptive terminology. Interestingly, Boyd notes that this active elaboration of the metaphor can also influence knowledge of the vehicle domain, as when empirical evidence for cognitive activity suggests novel program structure.

This view of scientific metaphor as providing a program for research is corroborated by retrospective accounts of notable scientific activities. For example, Oppenheimer (1956) while recounting the use of analogy in physics (e.g., light "waves") states "in every case an immense amount of experience, of measurement, of observation, and of analysis has gone both to the correction of the analogies and to their confirmation" (p. 133). Such examples are numerous (Dreistadt, 1968) and underscore both the utility of metaphor in scientific explanation and its suggestive power in a theory constitutive role. Hence, not only does metaphor occupy a central position in the explanation of scientific theory, but also appears to underlie novel insights into the nature of scientific problems. As a colorful example, Dreistadt (1968) describes Kekule's sudden appreciation for the ring structure of the benzene molecule on waking from a dream in which a snake has seized its own tail and spun before his eyes. Thus metaphor can be described as having at least three roles in scientific affairs: at the individual level of scientific insight, as an explanatory mechanism, and in a larger context of providing a problem setting which gives impetus and direction to continuing experimentation.

5. Metaphor in psychology

As measured against other cognitive phenomena, the use of metaphor has received relatively little attention in psychology. Nonetheless, the psychological literature in this area is broad and of considerable relevance for computational approaches to metaphor and analogy. Hence, this section forms the bulk of this paper. Fortunately, theoretical and empirical studies of analogical reasoning have been periodically reviewed, thus allowing an examination of this literature by major theoretical perspectives, including where possible notable empirical findings. The discussion of psychological approaches to analogical reasoning will be roughly historical, with older, somewhat solidified research traditions presented first, followed by a heterogeneous assortment of more current efforts. Of particular interest will be hypothesized structures and processes which might be suggestive for research in AI.

5.1. Theoretical approaches to metaphor and analogy in psychology

Perhaps one of the earliest approaches to analogical reasoning comes from the systematic study of individual differences through the use of mental testing. Sternberg (1977) characterizes differential theories by their assumption that analogical reasoning, as measured by traditional psychometric tasks, draws on mental abilities which taken together constitute intelligence. An example of such a task would be the solution of proportional analogy items (e.g., A:B::C:[D1,D2,...,Dk]) where the subject makes a choice among D terms. The existence of these abilities can be demonstrated by the covariance of psychometric measures of analogical reasoning with hypothesized, underlying factors of intellectual ability. Empirically, tests of analogical reasoning have been proposed as measures of:

1. measures of generalized intellectual ability,
2. fluid ability to reason among complex relations in problems not amenable to recalled solutions (the latter viewed as crystallized ability),
3. and convergent production of a unifying concept by the grouping of divergent ideas.

Unfortunately, these hypothesized psychological constructs give no clear suggestions for processes or representations supporting mental abilities and may not be generalizable beyond the psychometric materials used in their demonstration.

Another early approach to the study of metaphor comes from Gestalt psychology (Billow, 1977; Honeck, 1980) in which figurative constructions are seen to arise through similarity of the "expressive qualities" of literally dissimilar events. Similarity at a sensory level is supported by a number of assumptions concerning perceptual organization:

1. perception involves all sensory modalities simultaneously and thus allows synaesthetic experiences (e.g., a "sad" landscape or "orange" warmth);
2. perceptual inputs come as undifferentiated and unanalyzable wholes;
3. and perceptual experience is inextricably embedded in a motoric and affective context.

Hence subjects are assumed to have an affective disposition towards a subject of experience which naturally gives rise to metaphorical comparisons. For example, empirical evidence has been shown for agreement across subjects in linear expression (i.e., relatively simple line drawings) of linguistic connectives (e.g., and, if, because) in isolated sentences (Werner and Kaplan, 1963). Unfortunately, the undifferentiated nature of such organismic theories leaves specific processing or representational constructs unclear.

Psychodynamic theories have also contributed to the study of metaphor (Billow, 1977). Assuming the existence of an unconscious mind in which past events originally experienced as being unpleasant or threatening remain active, similar events in current conscious experience are seen to stimulate these repressed, threatening memories. Metaphorical expressions arise in the speaker's unconscious attempt to discharge affective energy stemming from these threatening unconscious experiences. As an example, during therapy a patient's remark that "I've wandered off the point and can't find it again" is seen to reflect a similarity between the patient's apprehension regarding negative evaluation by the therapist and early childhood experiences in which the patient feared maternal rejection on the basis of difficulties during breast feeding. Metaphors are of interest to psychodynamic theorists, then, primarily in that they provide projective (in the sense of unconscious material projecting into conscious manifestations) evidence of psychic processes. Although the hypothesis of unconscious, surging affective forces meets with resistance from other psychological orientations, metaphor as a mechanism of cognitive homeostasis is interesting, particularly with respect to the evocative affective qualities which many literary metaphors enjoy. However, it is not clear that the psychodynamic perspective contributes directly to an examination of the role of metaphor and analogy in problem solving or learning as typically construed within a computational framework.

Although the bulk of psychological research assumes mediating cognitive processes, metaphor has also been examined by behavioral psychologists (Billow, 1977; Paivio, 1979; Honeck, 1980) with an eye towards excising unobservable mental events in favor of clearly observable behavioral outcomes. Theoretically, metaphor is seen to arise out of stimulus similarity such that verbal responses originally reinforced in the presence of a complex of stimuli are emitted in the presence of new stimuli sharing sensory qualities with the original complex. As an example, a child's first glass of soda water might elicit a response likening the soda to the prickling sensations typically experienced when one's foot "falls asleep." From this perspective, metaphor amounts to little more than a verbal utterance which is relatively new in the commonplace verbal repertoire. Admitting limited influence for cognitive mediation, some theorists have argued for an explanation of metaphor in terms of common verbal associations, with utterances designated metaphorical arising out of somewhat more remote associations than those occuring in normal discourse. Behavioral theories of metaphor, by excluding mediating structure and process, are fundamentally at odds with most research approaches in AI (at least information-processing approaches) and appear to offer relatively little guidance to these activities.

As a final theoretical approach to analogical reasoning, information processing theories might seem promising for AI in their explicit concentration on representational structures and processes which are amenable to computer simulation. Unfortunately, these theories do not specify a particularly coherent set of constructs (Sternberg, 1977), diverging widely over what constitutes likely structures and processes. Hypothesized cognitive structures include fixed sets of relation types, propositional networks, lists of transformation operators, and lists of measurable features. Generalized processing notions include some form of encoding or decomposition of stimulus inputs, inferential discovery of relationships forming an analogical ground, application of these inferred relationships for the current stimulus situation, and evaluation of inferred relationships with respect to possible goal specifications (e.g., finding a unique or best solution candidate for a proportional analogy item). Sternberg evaluates information processing theories harshly, concluding that they generally fail with respect to adequate empirical support, complete specification of processing mechanisms, generality across varied forms of analogies (most studies focus on proportional analogy test items), and accounts of how individuals differ in their ability to reason analogically. Nonetheless, information processing theories will be discussed at some length in coming sections of this paper, including a component–process model advanced by Sternberg.

5.2. Exemplary psychological studies

Descending from abstract theoretical orientations which are wide in scope, examination of several recent psychological studies of analogy and metaphor can give some feeling for both the depth and diversity of hypotheses currently flourishing. As a pragmatic

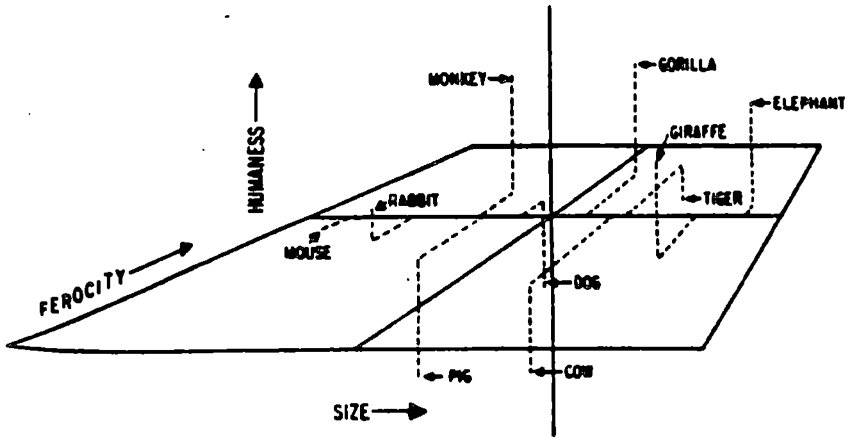

Figure 1
Concepts as points in a multidimensional space (Rumelhart and Abrahamson, 1973).

division, these studies will be described separately by focus on proportional analogical reasoning, metaphor comprehension, and problem solving by analogy. Several studies will be described within each area, followed by brief concluding remarks.

5.2.1. Studies of proportional analogies

Study of proportional analogies in the psychological literature has continued, probably owing to the use of proportional analogies as measures of intelligence in early theories of mental abilities and to the surface simplicity of these kinds of problems (i.e., target and source are explicitly presented, along with partial information regarding the ground). Furthermore, proportional analogies play a major role in many tests of general intellectual ability, and so provide often–used indicators of cognitive development. With respect to processing specificity, studies of proportional analogy are among some of the most complete in the psychological literature, although this specificity may be purchased at the cost of generality in underlying representational structure.

Analogy as conceptual proximity

Rumelhart and Abrahamson (1973) account for proportional analogical reasoning by proposing a similarity metric of inverse geometric distance in a conceptual space modelled as a multidimensional coordinate system (see Figure 1). Concepts (animals in this case) are represented as vector measurements over a fixed set of dimensions (ferocity, size and humanness). Given such a representation for **A**, **B**, **C** and **D** terms, an ideal analogy point, **I**, is calculated as ($C_k + B_k - A_k$, for each of k dimensions). Solution of the proportional analogy then amounts to selection of a **D** term which gives a minimal vector distance from **I**. Empirical results with human subjects show that both first choice and rankings of **D** terms fit those predicted under the hypothesized inverse distance similarity metric. However, item variation is appreciable, suggesting differential item difficulty not predicted by the theory. In addition, Rumelhart and Abrahamson show that arbitrary ideal analogy points are trainable and can be used as nominal concepts in later analogical reasoning.

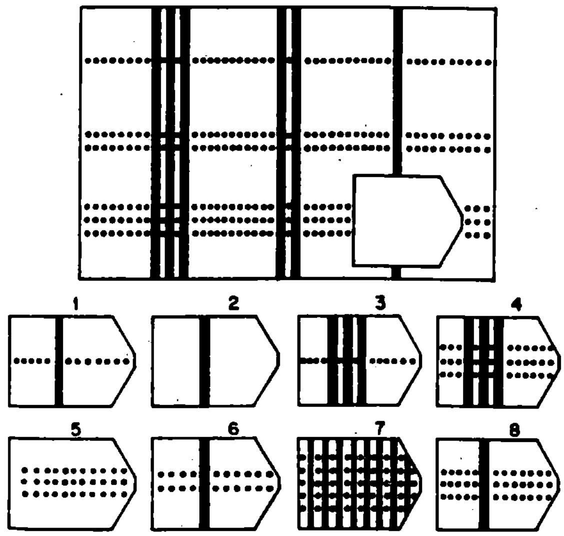

Figure 2
Sample item from Raven's Progressive Matrices (Hunt, 1973).

While these results are interesting, it is not clear that representation of concepts as points in a multidimensional space would be sufficiently expressive for more complex judgements of similarity (e.g., figurative similarity over abstract functional relations as in *Men are wolves*). Specifically, dimensions participating in such a representation reflect a rather clever condensation of information and processes contributing to similarity judgements which are largely implicit because of the spatial metaphor used in the description of represented objects (Palmer, 1978). From a symbol-processing perspective, one might wish that such information and processes were more explicitly reflected in the knowledge representation. Interestingly, this notion of similarity expressible as distance in a multi–dimensional space has been advocated recently (Parker and Kiselewich, 1984) in the guise of utilizing an "analog" representation to facilitate partial rule matching and generalization in an expert system.

Feature–matching models of proportional analogy

Hunt (1973) gives "gestalt" and "feature analytic" process models of proportional analogical reasoning for figural analogies from Raven's Progressive Matrices Test, an instrument often used as a measure of general intellectual ability (see Figure 2). Matrix items are encoded as unspecified visual figures or sets of features, respectively, for gestalt and featural models. In both representations, figures within the matrix are distinguished from background patterns, thus giving clear information regarding the salience of available stimulus materials. Gestalt processes find a solution by repeated application of continuation (e.g., continue a line through a row element) or superimposition (e.g., align an element by some arbitrary reference point) until a unique solution pattern is found which best fits a candidate answer. The feature analytic process model essentially conducts means-ends-analysis over a set of transformation operators (e.g., delete, expand or contract a matrix element) to account for differences in features for given rows. Eventually a set of required features for the missing element can be found which is a subset of the features describing a particular candidate solution. According to Hunt, hand simulation of gestalt and featural process models favors the feature analytic model, which is capable of finding correct solutions to all items on a standard form of the Progressive Matrices. An adult achieving comparable results would be judged "bright."

As with Rumelhart and Abrahamson's work, it is not clear if Hunt's feature analytic model could be extended to more complex forms of knowledge, although in fairness neither study proposes to do so. The gestalt model presented by Hunt is unfortunately difficult to evaluate owing to the indeterminate nature of "visual figures" and the operations which compare and transform them. This latter difficulty holds much in common with gestalt theoretical orientations towards analogical reasoning discussed earlier.

In a more recent contribution to the study of proportional analogical reasoning, Sternberg (1977) presents a carefully crafted component-process model which incorporates aspects of many more traditional approaches. Competence in solving proportional analogies is hypothesized to rest on some additive combination of five mandatory process components:

1. an *encoding* process identifies attributes and values for a term in the problem;
2. an *inference* process discovers a rule, X, relating A and B terms;
3. a *mapping* process discovers a rule, Y, relating A and C terms;
4. an *application* process applies the "higher order" rule, Y, to the inferred rule, X, to generate a rule, Z, which relates C and D terms;
5. and a *preparation-response* process serves as a control mechanism for monitoring other processes and making a response.

Sternberg also mentions an optional *justification* process which may choose among alternative candidate solutions if none unambiguously match a choice term.

On the basis of these primitive process components, Sternberg outlines several alternative models which combine components in different ways. Latency and error measurements with human subjects are used to select one of these alternative models. The preferred model is an additive combination of component processes in which subjects:

1. *encode* all terms of a particular analogy item,
2. exhaustively *infer* transformation rules for A and B terms,
3. incrementally *map* and *apply* attributes for A, C and D terms until a unique choice option can be discriminated.

Thus, the empirically preferred model suggests that subjects thoroughly explore relations between A and B terms, but perform mapping and application processes for each attribute within terms, stopping when a discriminating choice is possible between D terms. Comparison of the proposed process model with human performance suggests that mapping and application components are necessary (i.e., rules relating A to B terms are not applied directly in the search for solution), and that encoding (i.e., mentally representing terms of the item) takes more time than any other component while correlating positively with overall performance. Grudin (1980) presents some disconfirming evidence for this model, showing that some subjects solving similar problems while thinking aloud appear to discover the relation between A and C terms, and then apply this relation to the B term in coming up with an answer.

As a mediating representation, stimulus terms are assumed to be stored in a working memory as lists of attribute–value pairs. For example, consider the representation of the concept, Lincoln, in the following analogy:

Washington:1::Lincoln:[a. 10, b. 5]

[president(16th), portrait-on-currency(5 dollars), war-hero(Civil)]

Sternberg claims this form of representation is isomorphic to Rumelhart and Abrahamson's multidimensional space model (one might assume also to Hunt's feature analytic model). However, Sternberg points out that such a simple featural representation is not a necessary component of the model, and that other representational schemes might be accommodated within the preferred component-process model.

Proportional analogical reasoning: some conclusions

In evaluating the contribution of studies of proportional analogical reasoning, it is easy to be encouraged by the specificity of processing descriptions and apparent verification with human experimentation. However, even the most careful work (probably that of Sternberg) leaves much open to speculation. For example, the inferential "discovery" of rules relating A to B terms apparently rests on some form of search through a space of candidate relations, perhaps constrained by partial similarity between representations of participating terms. Unfortunately, the inference process is not clearly described, and, given Sternberg's claim that alternative representations might be equally suitable, it seems likely that the process of discovering relations between terms is beyond detailed description as the theory stands. The discovery and use of mapping relations seems even more problematic. Perhaps as has been pointed out in another context, "an entire homunculus can be concealed within one processing box in a flow diagram" (Gentner, 1983b). Were one to ignore uneven process specification, there remains the issue of whether or not processes implicated in proportional analogical reasoning are generalizable to less constrained forms of reasoning.

Although discovery and application of analogical relationships seems likely in more complex forms of analogy (e.g., aspects of theory formation and extension), clearer specifications of knowledge representation and the integration of inferential processes are necessary. For example, a subject's knowledge of Abraham Lincoln may include information regarding the presidency, five dollar bills and the civil war, but a moment's reflection reveals seemingly innumerable other pieces of information (e.g., beardedness, humble beginnings, assassination, etc.) which might be considered in establishing some measure of similarity between one's knowledge of Lincoln and an arbitrary target. While proportional analogies introduce some measure of constraint in this search for salient similar aspects of target and source, the specifics of "inference" and "mapping" processes advocated by Sternberg would again seem to depend heavily on the representational medium in which these processes occur.

In a similar vein, the interpolative process proposed by Rumelhart and Abrahamson depends upon a scaling of experimental items such that the three dimensions (size, ferocity and humanness) do indeed capture relevant similarity among animals included in the study. Were other, less salient dimensions included (e.g., warm or cold–bloodedness), it is not clear that comparable results would be achieved. In both instances, a view of the underlying representation is given which fortuitously provides just those semantic aspects seemingly required to make sense of particular proportional analogies. Hence, although studies of proportional analogical reasoning give a useful starting point for consideration of more complex forms of reasoning (e.g., metaphor comprehension and problem solving), these studies generally trade processing specificity against representational generality, if not suffering on both counts.

5.2.2. Studies of metaphor comprehension

Studies of metaphor comprehension in the psychological literature are legion, concentrating either on nominal metaphors or similes (i.e., A is a B, A is like B) out of context or, less frequently, on metaphorical expressions in context (e.g., metaphorical summarization of a text passage). Much of this research, particularly studies concerned with developmental aspects of metaphor comprehension, has been confounded by poor controls for pre-existing knowledge, difficulty and appropriateness of task materials, level of expressive competence of subjects, and difficulty in distinguishing subjects' figurative expressions from simple categorization mistakes (Ortony, Reynolds and Arter, 1978).

To illustrate central points in the diversity of opinion surrounding metaphor comprehension, two widely divergent approaches will be described in some detail. The first approach takes metaphor as an implied *comparison* of target and source concepts, dependent on common characteristics between them. The second approach casts metaphor as an invitation to the reader (or hearer) to understand a target subject as some source subject, restructuring the target in a novel way much in keeping with traditional *interaction* views of metaphor. These approaches have obvious parallelism with views of metaphor and analogy discussed in earlier sections.

Metaphor as feature comparison

Exemplary of comparison–based approaches to metaphor comprehension, Malgady and Johnson (1980) describe a variety of studies in which mediating representations are given as sets of attributive properties or features which are "open-ended" to reflect context (linguistic and extralinguistic) and individual differences. In this approach, metaphor depends upon the similarity between target and source concepts as defined by a weighted linear function of the "salience" of common and distinctive features (Tversky, 1977). Salience in this context refers to the relative frequency of occurrence of a particular feature determined over a normative sample in which subjects are asked to give descriptive features of a concept. This allows a quantitative estimate of the typicality of various features and, by composition, the salience of concepts they define. According to Malgady and Johnson, "salience [of a concept] increases as the number of defining features increase, and as their probabilities approach a rectangular distribution" (p. 249). Thus, for example, based on a normative sample the concept of *rain* can be judged lower in salience than the concept of *tears*. Assuming that order of comparison for target and source is reflected in the weights applied to distinctive features, similarity can be asymmetric depending on the salience of compared concepts. For example, Poland would be judged more similar to the USSR than the USSR to Poland, owing to higher salience provided by the feature set attributive of the USSR.

Metaphor comprehension in this model occurs as distinctive features of the source are transferred to the target where they are integrated in some unspecified fashion. "Good" metaphors are assumed to have an intermediate level of contrast between target and source concepts, being neither bizarre nor mundane. Empirical evidence is presented that subjects' recall performance and judgements of similarity are higher for metaphors with more salient source concepts, that subjects give greater attention to distinctive features of the source than the target concept, and that modifiers included in metaphorical sentences can enhance subjects' ratings of similarity and figurative goodness. These results are given as support for the predictiveness of feature comparison models with respect to asymmetry of similarity judgements and effects of linguistic context.

Metaphor as a transfer of abstract relations

In opposition to comparison–based approaches, Verbrugge and McCarrell (1977) present a series of experiments designed to demonstrate the inadequacy of static sets of features or properties as a basis for metaphor comprehension. They examine mediating structures hypothetically resulting from metaphor comprehension, using prompted recall of figurative sentences (e.g., Billboards are warts on the landscape) as empirical evidence for subjects' inference of a ground relating target and source concepts (e.g., are ugly protrusions on a surface). These authors argue strongly for "abstract relations" as the basis for such a metaphorical ground and against "attributive concepts" such as feature sets or property lists. As a mediating representation, a "system of abstract transformational and structural invariants" (p. 526) termed a *schema* is activated by recognition of nominal terms (e.g., billboard) designating target and source domains. As an example of such a schema for the source concept of "straw" in the metaphor, "Tree trunks are straws for thirsty leaves and branches," Verbrugge and McCarrell suggest the following three representational components:

1. *structural invariants* are "a structure of relatively rigid nonporous material, of a hollow cylindrical shape, with a small diameter relative to its length;"

2. *transformational invariants* are "the vertical cylindrical space channels fluid flow from a receptacle to a destination against gravity; the goal of the fluid transport is to ameliorate thirst; the force for the flow is suction;"

3. and a *normative contextual instantiation* provides "the structure is paper or plastic, the receptacle is a bottle or cup, the destination is a person (the thirsty agent), and the source of suction is the person's mouth and lungs" (p. 526).

Verbrugge and McCarrell hypothesize that inference of the ground between target and source domains should be reflected in accurate recall of figurative sentences when a statement of abstract relations comprising the ground is provided as a recall cue. The ground of a metaphor is seen to asymmetrically favor the source domain with abstract relations extended to and instantiated in the target domain. In this way, the target domain may be "structured" in terms of relations holding in the source domain. Empirically, this hypothesis is partially supported with high levels of recall when subjects are given the relevant ground as a prompt (nearly as high as when given target and source terms as cues). Other evidence is presented to rule out an explanation of recall performance in terms of simple word associations or isolated effects of target and source concepts. However, while the evidence does make some interaction between source and target concepts seem likely, subject performance does not appear incompatible with a comprehension model based on a weighted combination of common and distinctive features as suggested by Malgady and Johnson. In fact, Verbrugge and McCarrell's argument against comparison–based models seems to rest heavily on the nature of the metaphorical grounds used as prompts in the recall task. To accept their argument fully, one is forced into the position of treating these prompts as if they were indeed systems of "structural and transformational invariants" quite distinct from common and distinctive properties. This commitment is not *a priori* compelling.

Verbrugge and McCarrell conclude by arguing that metaphor comprehension should best be seen as the imposition of specific constraints from the source domain such that the target domain is considered in a novel but restricted fashion. Furthermore, they argue that these constraints, in the form of abstract relations reflecting structural and transformational invariants with respect to a particular context, cannot be captured by conventional, fixed attributive concepts in which features or properties are simply concatenated to form meaning. To illustrate, the authors use two sentences with identical target terms:

1. *Tree trunks are straws for thirsty leaves and branches.*
2. *Tree trunks are pillars for a roof of leaves and branches.*

Structural invariants involved in the abstract relations forming the metaphorical ground are context dependent (i.e., a hollow tube and a solid column, respectively) and provide a novel way of considering the target concept. Verbrugge and McCarrell argue that such context–dependent novelty, characteristic of all but relatively trite metaphors, would be difficult to provide with a fixed set of attributive properties for tree trunks.

Although mediating structures and attendant processes are described in a translucent fashion, the general character of metaphor comprehension in Verbrugge and McCarrell's work is clear: target and source domains (not just terms) *interact* in a novel fashion with oftentimes unpredictable results. Knowledge of the source domain serves as a filter or framework for a restructuring of the target domain, and the resulting knowledge structures represent a fusion of knowledge from both domains rather than a straightforward comparison. Although implicit, abstraction of domain knowledge appears to play an important role in the discovery of relations capable of instantiation in both domains.

A recent report by Johnson (1984) makes the representation and process notions suggested by Verbrugge and McCarrell somewhat more explicit. Also studying the comprehension of isolated metaphors, Johnson proposes a set of "semantic combinators" (mapping functions) which allow for differing levels of "accommodation" of mapped information between target and source objects. As a representation, object schemata are described as frame–like structures consisting of "facets." Facets, similar to Verbrugge and McCarrell's transformational invariants, are described as functional properties or relations of the represented object which the possessor acquires during "goal–directed interaction with the object" (p. 193). As an example, the hardness facet of rocks is given as "rocks do not change shape under the application of external physical force" (p. 194). Other facets of rocks described by Johnson include immobility, greyness in color and heaviness.

A three–stage process model is described, with a mapping function (or combinator) given as an example at each stage (Johnson mentions but does not describe other mapping functions). At an initial, global level of processing, an *identity mapping* may be formed between facets of target and source. At this level, facets are transferred literally without any change in semantics. For example, when confronted with a sentence such as, "My shirt was a mirror," one might interpret the metaphor by directly transferring a reflective facet of a mirror schema to a shirt schema. At a more analytic level of processing (termed analytical–1), an *analogy mapping* can transfer a facet by a "change in sense" from source to target schemata. For example, interpreting "My sister is a rock," one might accommodate the sense of hardness contained in the source (rock) schema by replacing portions of the hardness facet with information more consonant with the target (sister) schema. Hence the hardness facet as transferred to the sister schema might read, "my sister does not change behavior under the application of psychological pressure." At the third processing level (termed analytical–2), a *predicate mapping* permits the elaboration of an analytical–1 level mapping within the target domain. For example, having transferred the hardness facet as described above for the analogy mapping, one might further infer that the speaker's relation with his sister could be characterized as being unfriendly or interpersonally distant, both following naturally from knowledge sources appropriate for the target schema but unrelated to the source schema.

Johnson gives descriptive findings as partial empirical support for this model of metaphor comprehension. Asked for multiple interpretations of six, single sentence metaphors (i.e., combinations of [My shirt, My sister] was a [rock, mirror, butterfly]), adult subjects typically give two or three distinct interpretations. These interpretations were subsequently coded by Johnson to determine the types of facets and levels of mapping used by subjects during interpretation. Johnson reports that 17% of all coded facet transfers were at the global level (e.g., using an identity or related combinator), 27% were at the analytical–1 level, and 50% were at the analytical–2 level. The frequency of such responses varied according to the type of objects given as target and source. For example, highest levels of global transfer were identified with "shirt" as the target, while "sister" as the target apparently required some change in sense if the metaphor was to be meaningfully interpreted. Although Johnson uses a coding framework drawn directly from her hypothetical model, these descriptive findings do suggest structures and processes which might underlie metaphor comprehension in a manner consonant with an interaction view of figurative language use. Subjects' interpretations consist primarily of abstract relations which are uncovered by more than direct comparison, particularly in the case of relations coded at the analytical–2 level. In summary, Johnson gives an interesting account of the sort of knowledge structures and mapping operators which might support the kind of figurative transfer implicated by Verbrugge and McCarrell.

Metaphor comprehension: some conclusions

The approaches described above give fairly convincing but strongly divergent pictures of psychological structures and processes supporting metaphor comprehension. "Attributive" representations (e.g., feature sets), through their ease of operationalization, allow quantitative measurement of the similarity of concepts which purportedly underlies metaphor comprehension. Unfortunately, parsimony in describing the basis of judgements of similarity apparently contributes little to clear processing notions of comprehension. While two concepts may be judged similar with some inter–subject reliability, it is not clear how knowledge of a source domain is extended to a given target domain in the sense, for example, of the theory constitutive metaphors discussed in an earlier section of this paper. As Ortony *et al.* (1978) point out, comparison may be a component in metaphor comprehension, but is probably insufficient as an exclusive explanation. In contrast, more complex relational representations offer a richer explanatory vocabulary for metaphor comprehension (e.g., the extension of invariant abstract relations from source to target), but falter at the point of explicitly specifying or measuring characteristics of such representational structures.

The distinction raised in preceding paragraphs between featural and structural mediating representations extends beyond the particular studies presented here. Ortony *et al.* (1978) discuss a similar distinction between propositional and schema theories of semantics. Propositional theories would include featural models such as that of Malgady and Johnson (1980) as well as proposals which represent meaning directly as some fixed set of propositions (e.g., Miller, 1979). In these models, meaning is given by a relatively static set of assertions assumed to capture "core meaning" or things which are necessarily true of the concept being defined. As an example, a cow could be represented as a mature, female member of cattle. Such representational commitments, through their inflexibility, typically require special forms of processing for figurative language. These processes must uncover literal statements of similarity implicitly underlying figurative expression. Ortony *et al.* argue that such a representation of word meaning "simply fails to permit the kind of flexibility that would be required to make sense of a metaphor" (p. 935).

Schema–based theories of semantics, on the other hand, are described as relatively open-ended collections of stereotypic knowledge about concepts being represented. In contrast with propositional accounts, a cow might be represented as a domesticated animal which gives milk, eats grass, and is in turn eaten by humans. These and other elements of practical knowledge are assumed to take the form of variables or slots with loose constraints on suitable instantiating values, all interconnected by a system of relations. Such schemata, Ortony *et al.* argue, are more flexible than their propositional counterparts, allowing figurative language to be understood in a similar fashion to literal language. Candidate schemata are applied on the basis of a more or less complete match with an input expression, and judgements of metaphoricity occur *post hoc* on the basis of a constructed understanding.

In a subsequent paper, Ortony (1979) argues that highly salient predicates involved in source schemata can be "promoted" or actually "introduced" into the target domain, resulting in a novel conceptualization of that domain. This process of matching and schema application, relatively straightforward for literal utterances, gives rise to the unpredictable, interactive nature of language understanding implicated in the instantiation of "transformational invariants" of Verbrugge and McCarrell. Rumelhart (1979) makes this argument as well, claiming that comprehension of conveyed meaning in metaphorical utterances does not depend on compositional notions of semantics, but rather that figurative and literal language are comprehended in essentially the same fashion through the selection, application and verification of "conceptual schemata." Rumelhart and Norman (1981) extend the same approach to the acquisition of procedural skills such as using a text editor.

From a distance, then, parallel themes can be discerned in the psychology of metaphor comprehension, centering around representational proposals for semantics. Propositional representations generally lead to comparison–based theories of metaphor comprehension in which figurative expression implicitly signals underlying literal communication which must be uncovered by processes distinct from those employed in literal language comprehension. Schematic representations, on the other hand, generally lead to more flexible, interaction–based theories of metaphor comprehension in which figurative expression is processed in a manner consonant with literal language comprehension. Both forms of theory continue to flourish in the psychological literature, although perhaps with different promise. According to Ortony *et al.* (1978) "one shortcoming of almost all the research ... is that the locus of a metaphor is assumed to lie in a word, or perhaps an expression, within a sentence. If research is to progress it will be necessary to adopt a less restrictive account of what a metaphor is" (p. 937). In particular, if theories of metaphorical and analogical reasoning are to extend to more far–reaching issues of learning and creative thought, more flexible, open-ended accounts of metaphor comprehension might be preferred.

5.2.3. Studies of problem solving

It is hardly surprising that analogical reasoning in problem solving has been and continues to be an area of vigorous research in cognitive psychology. The use of analogy in complex problem solving is frequently advocated (Newell, 1983), although with relatively little guidance as to how such reasoning skills are to be developed. In addition, transfer of problem solving skills beyond the temporally–limited experiences available during training forms the central purpose of most educational curricula, and a clear understanding of how to facilitate appropriate transfer while minimizing negative (or inappropriate) transfer is an obvious goal. Finally, as with continued study of proportional analogical reasoning, the use of analogies in problem solving is assumed to be an important contributor to intelligence in general. Since a large body of research has accumulated, this section will give a brief overview of studies of analogical transfer in problem solving, examining selected studies in detail to highlight current research issues. The interested reader is referred to existing reviews of the transfer literature related to problem solving, especially Novick (1985), Brown and Campione (1985), Brown, Bransford, Ferrara and Campione (1983), and Lester (1982).

From an abstract vantage point, aspects of analogical reasoning in problem solving appear quite similar to generalized components discussed in previous sections: given a target problem to solve, the subject recognizes and retrieves an appropriate source solution method, eventually applying that method to solve the original target problem. The benefits of analogy in problem solving are (perhaps) obvious: by retrieving an applicable solution approach, previous problem solving efforts can be pressed into service without the necessity of "starting from scratch" for either superficial problem variants or genuinely novel problems. Given the apparent flexibility of human reasoning, we might even assume the necessity of analogy in surviving a changing environment. The costs incurred through use of analogy in problem solving, however, may be less obvious. Effort expended in recognizing a novel target problem as an instance of a known problem class or in adapting a retrieved solution method to the idiosyncrasies of the target situation might exceed the effort required for solving the problem "from scratch." Perhaps worse, the problem solver might be led astray by devoting resources to a "false analogy," which appears promising at retrieval but proves inappropriate as a solution strategy for the target problem.

Conditions under which the recognition, retrieval and application of a source occur, unfortunately, have yet to be adequately understood. Returning to the abstract characterization developed in the last paragraph, we find that a researcher is quickly faced with providing hypothetical answers to a number of important and much more specific questions.

1. What is it that subjects understand about a "target problem?" What aspects of the problem do they attend to in forming a problem representation? This representation must play a crucial role in recognizing and then using an applicable source.
2. What constitutes a "source problem?" Do specific problem cases or more abstract solution methods (or both) serve as a sources for analogical reasoning?
3. At what point during problem solving do subjects recognize the partial similarity between their understanding of the target problem and source situation(s)? Do analogies arise at problem solving "impasses" or are they integrated more smoothly into routine reasoning?

While most researchers make some commitment to one or more of these questions, none have been answered convincingly. Successful transfer of problem solving skills appears to be sensitive to a number of factors, including characteristics of the problem solver, the problem solving task and the immediate setting in which reasoning occurs.

Early investigations of analogy use in problem solving were equivocal with regard to the frequency with which subjects spontaneously show positive transfer from previously experienced problems. Research with relatively formal "puzzle problems" (e.g., the Towers of Hanoi) suggested that positive transfer did occur (Hayes and Simon, 1977; Luger and Bauer, 1978) but was dependent upon problem difficulty and the order in which problem variants were presented. Unfortunately, despite the formal specification of reasoning tasks within an information processing framework, these studies do not arrive at convincing accounts of which problem solving processes contribute to transfer or what aspects of competence are actually transferred. Transfer in these studies is typically measured as a reduction in the time required to solve a problem or the number of problem states entered during a solution attempt. In neither case do the results show how symmetries between problem isomorphs are noticed, retrieved or utilized. Instead, they serve as evidence that some form of transfer can occur, even on problem solving tasks which subjects are assumed never to have seen before entering the experimental setting. In some ways, it is surprising that any transfer at all can be observed under conditions in which subjects can be assumed to be novices in both the source and target domains for problem solving. This is a recurring problem in much of the available literature.

Some recent studies of analogy in problem solving, while not much more specific about knowledge attributable to subjects, do give a more suggestive picture of the circumstances under which analogies between source and target domains might be recognized and used. We will examine the work of three researchers in some depth, particularly because they provide divergent theoretical accounts of analogical reasoning in problem solving.

Analogy as structure mapping

Gentner (1982, 1983a) provides a structure mapping theory for analogy (taken to subsume metaphor, simile and model) which relies on a syntactic mapping of domain structure from source (or base) to target. These theoretical concepts will be presented before a discussion of empirical findings. As a representation, Gentner assumes that domain knowledge is partitioned into objects and predicates. Objects form the basal level for a particular reasoning context, although at a more detailed level of reasoning they could be further decomposed into constituent objects and predicates in a recursive fashion. Predicates may be either unary "attributes" or n–ary "relations." Taking objects or other predicates as arguments, relations are combined in a propositional network. Gentner (1983a) makes strong claims for the importance of this structured representation, arguing that simple attributive predications (e.g., features) cannot account for judgements of nonliteral similarity crucial to analogical reasoning. In addition, the syntactic specification of this representation is not taken to be either arbitrary or logically complete. Rather, these representational assumptions are offered as a model of "the way people construe a situation" (Gentner, 1983a, p. 157).

The basis for analogical reasoning in Gentner's theory is a "non-literal similarity comparison" between domains. The comparison results in a mapping, M, of nodes from source to target domains which is constructed so that much of the relational structure of the source domain is "true" in the target domain. That is, identical correspondence at some level of relational abstraction is possible. This latter condition is qualified by a "systematicity principle" requiring that corresponding relational structure involve few attribute predicates and many "higher–order" predicates (i.e., predicates which take yet other relations as their arguments). With respect to the use of analogies in scientific theorizing and explanation (Gentner, 1982), a good explanatory analogy for scientific purposes maps few attributes and many relations (particularly those participating in higher–order or systematic conceptualizations) from source to target domains. While described syntactically with respect to a structured representation of domain knowledge, Gentner's structure mapping theory implicitly attempts to capture semantic aspects of causality and constraint which can be used in a predictive fashion in a new domain.

Gentner (1982) develops a rich descriptive vocabulary for explanatory analogies which is worth discussing here. A number of internal characteristics of an analogy are described:

1. *clarity* speaks to the precision of object correspondence, both in terms of a single object mapping to multiple objects in another domain and those multiple objects in the mapping range being relationally heterogeneous (either situation reduces clarity);

2. *richness* refers to the number of mappable relations per object;
3. *systematicity* describes the extent to which mapped relations participate in higher–order relations forming "mutually constraining conceptual systems;"
4. and *abstractness* refers to the hierarchical level (relative to contextual knowledge) from which mapped relations are drawn.

External characteristics of explanatory analogies include:

1. *base specificity* refers to the reasoner's degree of explicit knowledge of the source domain;
2. *validity* describes the verifiability of imported relations in the target domain;
3. and *scope* describes the applicability of a source analog across a variety of specific targets.

Gentner argues for a tradeoff between scope, clarity and richness in explanatory analogies, claiming that an analogy which is wide in scope while mapping a rich density of relations is difficult to achieve without allowing object and predicate correspondence to "slide about." Effective analogies for scientific explication and prediction are described as being high in scope, clarity and abstractness but low in richness. "Expressive" metaphors intended for descriptive or evocative purposes, on the other hand, are described as being high in richness for lower level relations but lower in clarity and systematicity. Empirically, subjects rate good scientific analogies as being high in clarity but low in richness, while good expressive analogies are rated high in richness. As additional supporting evidence (Gentner, 1983a), subjects' interpretations of analogies contain more relational than attributive assertions, and subjects rate as most apt those analogies richest in mappable relations.

As noted by others (Carrol and Mack, 1982), Gentner's theoretical model of analogical reasoning is sparse with respect to explicit process specification. Aside from predicting the differing utility of two analogies for explanatory purposes, the structure mapping theory stops short of providing a framework for understanding when or how analogies might be spontaneously recognized and used while problem solving. Nonetheless, some hypotheses are evident in Gentner's writing.

In relation to recognition and retrieval, an "analogical shift conjecture" is offered (Gentner, 1983a; Forbus and Gentner, 1983, 1986) as a stage model of the spontaneous comparisons that might occur as a subject acquires skill in a new domain. Central to this model is the hypothesis that accessibility of a source in memory is inversely related to its utility for reasoning in the target domain. The first stage predominantly consists of "literal similarity matches" involving correspondence of both object descriptions and relational structure. As a result of considerable overlap, access to the source is predicted to be relatively easy, perhaps through multiple paths of activation during retrieval. In contrast, utility of the retrieved source should be low since there will be too much overlap for the subject to judge the causal relevance of various aspects of the source. At the second stage, subjects are able to recognize analogies involving few object mappings but considerable relational overlap. Access to an analogous source might be more difficult since object matching does not contribute to activation, but utility of the retrieved analogy should be higher since the shared structure between source and target is "sparse enough to permit analysis" (1983a, p. 168). In the final stage, subjects are assumed to have formed "general laws" within the domain which consist almost entirely of abstract relational structures. Hence, subjects are able to notice and apply these laws when confronted with new (target) situations. Access in this third stage is not described, while utility is assumed to be high across different domains (scope). In overview, novices in a problem solving domain rely on literally similar, previously experienced problem solving episodes, while more experienced problem solvers utilize analogies and eventually are able to apply general laws. Thus recognition is primarily a function of the reasoner's experience in a problem domain.

While little empirical evidence is yet available for the stage model of learning, Gentner and Landers (1985) have examined the influence of structural similarity on recall of previously experienced stories. Subjects read a set of stories during a training condition and at test (six to eight days later) were asked to report any remindings of training stories which occurred while reading target stories. Subjects were asked to write down the stories they were reminded of as accurately as possible. The stories presented at test were designed to bear different structural similarity to each of the training (source) stories. When written accounts of recalled stories were compared with the given target stories, results showed decreasing likelyhood of access for sources which were:

1. "mere appearance" matches (similar objects and first–order predicates),
2. "true analogies" (similar first and higher–order predicates),
3. or "false analogies" (similar only in first–order predicates).

Although subjects report recalling superficially similar sources most readily, their ratings of the appropriateness of these source stories after recall favor the "true analogy" matches. Thus access to a source in memory may be facilitated by surface level similarity, but subjects are later able to recognize structurally similar remindings as being most apt, suggesting that the utility of a retrieved source may be determined by similarity at a more systematic (structural) level. While suggestive in terms of the shift conjecture, this study does not examine the effects of subjects' experience in the task domain or the relevance of problem material in the surrounding reasoning context at training or test. As will be discussed in the conclusion to this section, both are important considerations for the use of analogies in problem solving.

Given a source analog, Gentner's structure mapping theory places several implicit constraints on the use of that analog in problem solving. Presumably supporting analogy, the mapping process crucial for identification of effective structural correspondence is described as what appears to be a parallel yet interdependent identification of object and relational mappings. To avoid circularity in identifying an appropriate correspondence mapping, Gentner restricts predicates to an identical match between source and target. Nonetheless, the process of extending an initial correspondence mapping given a particular source and target is not clearly specified. Within the prescriptive qualities of good explanatory analogies, we might expect that subjects would somehow prefer more systematic aspects of the source analog for extension to the target. Neither is the evaluation of an existing mapping explicitly addressed, although it is clear that a subject's familiarity with the source domain (base specificity) and goals in the target domain would each play an important role in evaluation.

Use of a given analogy in accordance with the structure mapping theory has been empirically studied. Gentner and Gentner (1983) examine problem solving with the intention of settling the issue of whether structure mapping activities are central to cognitive processes in analogical reasoning (termed the "generative analogy hypothesis") or epiphenomenal in retrospective descriptions of analogical reasoning (dubbed the "surface terminology hypothesis"). In this study, subjects are presented with either a water reservoir ("flowing waters") or crowd of mice ("teeming crowd") analogy as the basis of instruction about current flow in electrical circuits. The authors' expectation is that these analogies will differentially aid in solving progressively more difficult electrical circuit problems.

Hypothesized facilitation of battery problems by the reservoir analogy and resistance problems by the crowd analogy are partially confirmed. However, a "generalized strength-attribute" seems to influence many subjects' performance, suggesting that preexisting knowledge brought to the experimental task may influence performance as much as analogies encountered during experimentation. In addition, subjects receiving the flowing waters model appear to have difficulty in reasoning about how water behaves under gravitational forces in a closed system, suggesting low "base specificity." In summary, despite the authors' enthusiasm for the centrality of analogy in problem solving, the reported experimentation does not clearly support or disconfirm their hypothesis. Hence, although Gentner's structure mapping theory yields promising concepts relating structural aspects of knowledge representation to a plausible basis for judging explanatory analogies, the implications of this theory for analogical reasoning in problem solving have yet to be fully pursued along theoretical and empirical lines.

Analogy as pragmatic induction

An alternative theoretical account of analogical reasoning in problem solving comes from Holyoak's (1984a; 1984b; 1985; Holland, Holyoak, Nisbett and Thagard, 1985) treatment of pragmatic aspects of analogical transfer. In opposition to Gentner's syntactic theory of structure mapping guided by systematicity, Holyoak argues strongly for the necessary influence of the reasoner's goals in exploring an analogy. The gist of Holyoak's argument is that analogy or any other inductive reasoning mechanism used in problem solving must constrain the space of possible inferences by using knowledge of the purpose to which the reasoning mechanism is put. In the case of analogical reasoning, an acontextual consideration of higher-order relations in the source would be insufficient as a basis for deciding which of many possible relations should transfer. Instead, Holyoak suggests that goals and constraints within the target domain, in addition to the surrounding reasoning context, must be used as a guide for transfer.

Not surprisingly, Holyoak (1984a) sketches a model of problem solving as the surrounding context in which analogies are recognized and used. Problem solving depends upon being able to construct a *mental model* of the problem situation, refining that model successively until a concrete solution can be found. The problem model is an abstraction that preserves properties (e.g., goals, constraints and operators) of the real world problem which are causally relevant to the modelling purpose. Problems can be solved to the extent

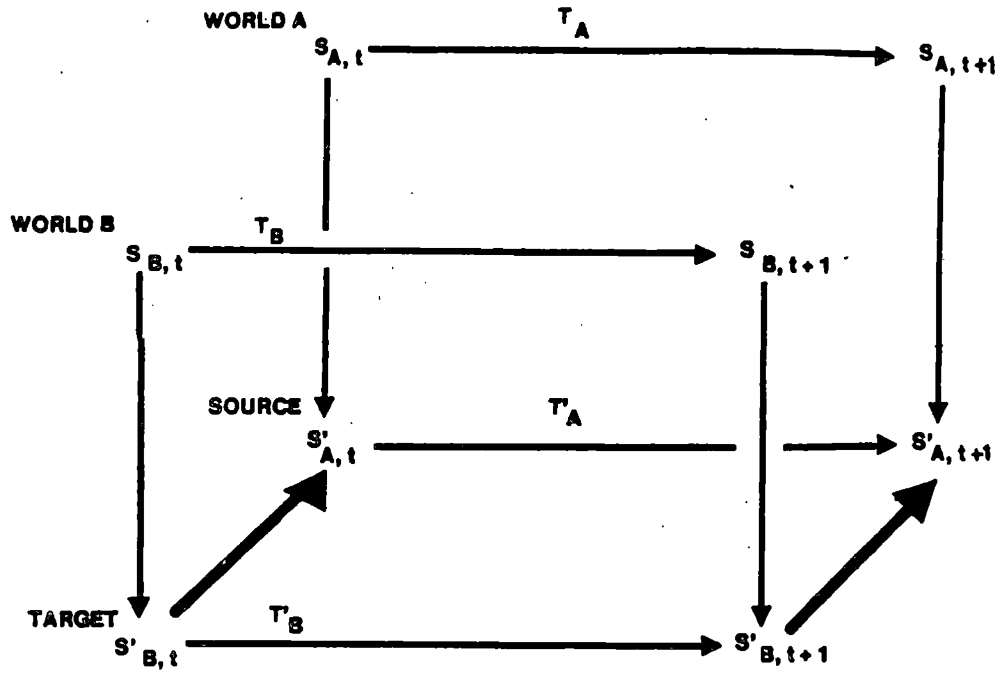

Figure 3
Analogy as a process of second–order modelling (Holyoak, 1985).

that the problem model enables useful predictions about state changes in the represented world (i.e., the real world problem). Expressed more formally, the problem model is a homomorphism between represented and representing worlds, and this mapping should commute if successful predictions are to be obtained. In actuality, models tend to be quasimorphisms in which mappings between objects, relations and actions in both worlds are incomplete. Quasimorphisms are often useful, if not completely accurate in their predictions. Problem solving, then, is a process of refining the problem model so that a solution plan can eventually be constructed and applied.

Given a genuinely ill-defined problem, the problem solver is likely to be unable to generate a solution plan which can be directly applied. When routine problem solving activities (e.g., forward or backward chaining given a set of problem–specific operators) reach an impasse, the role of analogy is to provide a refinement of the target problem model so that a solution plan can be constructed and applied. Analogical reasoning is described as a second–order modelling process in which a model of the source problem is retrieved and placed in correspondence with the incompletely specified model of the target problem (see Figure 3). State descriptions (S') and a transition function (T') in the target model are juxtaposed with corresponding elements in the source model by an analogical mapping (dark arrows at bottom of Figure 3). In the ideal, this process results in an isomorphism between source and target models, with the solution method effective for the retrieved source mapped into a solution method for the target problem. In practice, the mapped solution method may require further refinement.

To be pragmatic in a theoretical sense, Holyoak's model of analogy in problem solving should provide hypothetical frameworks for two questions:

1. How are source models recognized and retrieved;
2. How are components of the source model extended into the target to provide model refinement?

These questions are identical to those considered for Gentner's structure mapping theory discussed in preceding paragraphs.

For recognition and retrieval, Holyoak proposes a summative activation mechanism in which elements of a "retrieval cue" extracted from the target problem model serve as origins of activation for elements of a candidate source model. The process of cue extraction is not specified, but appears to rely on detection of what Holyoak terms an "implicit schema" reflecting relevant similarities between the target problem and a retrieved source. Particularly for analogies across problem solving domains (e.g., insight problems, discussed in a moment), Holyoak argues that retrieval may be quite difficult owing to considerable superficial dissimilarity between target and source models. He further hypothesizes that more abstract problem schemata induced over multiple instances of problem solutions will be more easily retrieved, since these schemata will eliminate many elements of superficial dissimilarity. Distinctions between surface (superficial) and structural (causally–relevant) similarities between target and source are used uniformly in Holyoak's writing, but appear to be context and goal–dependent for particular problem solving episodes. This is an important point given the pragmatic orientation of Holyoak's approach to analogy in problem solving.

Retrieval depends upon the cumulative activation of concepts involved in the source model by pathways originating at corresponding concepts in the target model. Since little is likely to be known in the target problem model at the outset of problem solving, Holyoak assumes that initial and goal states are likely to be the primary origins of activation, particularly as their constituent concepts are inferentially elaborated (e.g., values for properties are determined). Activation resulting from an assumed focus on the goal specification of the target model insures that most origins of activation in retrieval will be goal–relevant. Retrieval of a source model occurs when attention of the problem solver turns to a sufficiently activated source model. Attentional control, then, appears to be related to some sort of activation threshold.

Holyoak is less specific about the manner in which aspects of the source model are extended into the target model. Given that the retrieval process has revealed aspects of the "implicit schema" between source and target models, an initial mapping will be available after retrieval, and development of the analogy proceeds by "unpacking" the source model in a top–down fashion. That is, retrieval provides an initial, high–level mapping between models (i.e., correspondence between goals and constraints in the problem models) which guides the incremental extension of information from the source to the target. Extensions are subject to transformations implied by the mapping between models. Holyoak describes four types of mapping relations:

1. *Identities* are the same in both models, apparently forming part of the implicit schema. These should correspond to surface and structural similarities between models.

2. *Indeterminate correspondences* have yet to be mapped.

3. *Structure-preserving differences* do not impair causal structure in either model and do not impede construction of operators in the target model. These correspond to surface dissimilarities between models.

4. *Structure-violating differences* do impair operator construction in the target model and may require introduction of new objects or relations not present in either model. These correspond to structural dissimilarities between models.

In summary, mapping and extension processes in Holyoak's theory appear to be guided by the goal–relevance of various model components in both the target and source models. Refinement of existing operators in the target model serves as the impetus for retrieval through spreading activation, and attempts to map and transfer applicable operators from source to target models appear to dominate the process of analogical reasoning during extension (hence the term, "unpacking"). Analogical reasoning is but one mechanism in the larger context of problem solving by model refinement, and may "run dry" as structure–violating differences between models begin to predominate in the incrementally extended mapping. This would appear to leave open the possibility of other, supporting inductive mechanisms in problem solving, a possibility discussed more fully in Holland, *et al.* (1985).

Holyoak and others (Gick and Holyoak, 1980, 1983; Holyoak, Junn and Billman, 1984; Holyoak and Koh, 1985) have undertaken ambitious empirical studies of analogical transfer in solving "insight" problems. As an example, consider the following problem, originally proposed by Duncker (1945).

> Suppose you are a doctor faced with a patient who has a malignant tumor in his stomach. It is impossible to operate on the patient, but unless the tumor is destroyed the patient will die. There is a kind of ray that at a sufficiently high intensity can destroy the tumor. Unfortunately, at this intensity the healthy tissue that the rays pass through on the way to the tumor will also be destroyed. At lower intensities the rays are harmless to healthy tissue, but will not affect the tumor either. How can the rays be used to destroy the tumor without injuring the healthy tissue?

Such problems have a long history in the psychological literature, and are typically assumed to require some amount of creativity in achieving a problem solution. The basic experimental paradigm employed by Holyoak and his associates is to present novice problem solvers with a variety of source problems in the context of training (generally a story comprehension task), and then to measure the extent to which subjects are able to transfer problem solving strategies from these training materials to a novel problem presented during later testing. Thus, if the problem shown above were presented at test, subjects might have seen a structurally similar problem during training which involved a military leader wishing to overcome a fortress but unable to send his entire army over mined roads leading to the fortress. A solution in both cases would be to subdivide the attacking forces (rays or army) and have them *converge* on the central target (tumor or fortress).

Although Holyoak and associates present the results of numerous experimental manipulations within this basic paradigm, only selected results will be discussed here. Few subjects (10%) give a convergence solution to the radiation problem alone (without training); 30% spontaneously give the convergence solution having seen the military problem during training; and a total of 75% give the convergence solution having seen the military problem at training and given a hint to use the training problem in generating a solution. These ascending figures are cumulative over conditions (no training, training but no hint, training and hint). Thus, as with transfer experiments mentioned earlier, unassisted transfer from a given source problem appears infrequently, but subjects are able to use the source once it is explicitly pointed out to them.

Of more interest in the experiments reported by Holyoak and associates are explicit manipulations intended to alter the facility with which subjects can access and use a source analog. In terms of access, two general forms of manipulation are reported. First, the problem materials presented at training and test are systematically varied by their surface and structural similarity to one another. Holyoak and Koh (1985) describe experiments in which variations of a "laser and light bulb" problem are used as training (source) materials for the subsequent solution of the radiation problem at test. Briefly, the source problem requires that the damaged filament of an expensive light bulb be repaired by a laser without damaging the surrounding, fragile glass bulb. Surface similarity was varied by constructing problems in which a laser or an ultrasound device were the instrument of repair. A laser was assumed superficially more similar to the radiation device mentioned in the target problem than was the ultrasound generating device. Structural similarity was varied by introducing different forms of constraint in the source problem: either a fragile surrounding medium (glass) for the bulb or multiple generating devices, each without sufficient power to effect repair. The surrounding medium constraint was assumed more similar to the radiation problem (fragile surrounding tissues) than was the insufficiency of a single generating device. Crossing surface and structural variations of similarity (each high or low) gives four source problems for training purposes. These were given to four different groups of subjects, each of which was subsequently asked to solve the radiation problem.

Results of this experiment are consonant with Holyoak's model of retrieval discussed earlier: the high surface and structural similarity source is spontaneously used in a successful convergence solution by 69% of experimental subjects; transfer is lowered if either surface or structural similarity is reduced (38% and 33%, respectively); and with low similarity at both surface and structural levels only 13% of subjects achieve a convergence solution. Thus, any reduction in similarity (whether surface or structural) appears to interfere with spontaneous retrieval and use of an available source.

However, while surface dissimilarity might inhibit spontaneous retrieval, the theoretical account of retrieval discussed earlier suggests that subjects should best be able to retrieve a source when surface level aspects have been abstracted out of the encoded memory (e.g., by eliminative induction) and only structural aspects remain for summative activation. This hypothesis is partially confirmed in another series of experiments. As a second manipulation intended to influence retrieval, Gick and Holyoak (1983) required subjects to give summaries describing important similarities among multiple source problems presented at training (e.g., multiple problems solvable with a convergence strategy). In the process of summarization, it was hypothesized that subjects would generate an abstract problem schema, commit that schema to memory, and then be more likely to retrieve the schema than the original source problems from which the schema was induced. In actuality, summarization over two sources resulted in 45% of subjects giving a convergence solution without receiving a hint to use the source problems or schema. Further aids to summarization (a spatial diagram or a verbal summary given by the experimentor) gave a slight increase in spontaneous transfer (57% and 62% respectively). Furthermore, when subject's summaries were rated for quality, 91% of those subjects giving "good" summaries were able to produce the convergence solution without a hint. Thus, the summarization condition over two analogs appeared to increase the frequency with which subjects spontaneously transfer above baseline levels mentioned earlier (30% for a single source problem). Interestingly, the authors found little benefit in giving a summarization directly to subjects without requiring them to construct their own.

The two manipulation strategies described above also appear to have predictable effects on the use of a source, given that its retrieval has been assured by providing the subject a hint to use it. In Holyoak and Koh's (1985) study, for example, the percentage of subjects achieving a convergence solution given a hint to use the source problem was 78% for the high structural similarity condition (fragile surrounding glass) but only 54% for the low structural similarity condition (insufficient intensity of a single generating device). These percentages are collapsed across the surface similarity condition which has a negligable effect on achieving a correct solution given a hint. The authors conclude, in accordance with Holyoak's theoretical treatment, that surface similarity may play a more important role in spontaneous retrieval of a source than in its application once retrieved.

In overview, empirical support for Holyoak's pragmatic theory of analogical transfer in problem solving is substantial although in some ways nonspecific. Spontaneous generation of expected solution types for target insight problems having been exposed to structurally similar source problems appears to occur infrequently. However, the frequency with which subjects are able to generate such solutions can be affected by a variety of manipulations in both the problem solving task (e.g., requiring that subjects summarize multiple source problems at training) and in problem materials (i.e., systematic variation of surface and structural similarity between target and source). In studies with 4–to–6 year–old and 11 year–old children using age–appropriate insight problems (e.g., Miss Piggy transporting some precious jewels), Holyoak, Junn and Billman (1984) have found similar results for spontaneous transfer and effects attributable to variations in similarity between target and source problems. Although primarily quantitative results have been discussed above, Holyoak and associates also provide convergent qualitative evidence based on examination of verbal and written protocols collected while subjects participated in the experiments. Taken together, these varied results are encouraging, although the theory and supporting evidence fall short of specifying the precise nature of retrieval, mapping or application of a source problem model to achieve solution of a target problem. Construction

A WEIGHT IS HUNG ON A SPRING. THE ORIGINAL SPRING IS REPLACED WITH A SPRING

--MADE OF THE SAME KIND OF WIRE,
--WITH THE SAME NUMBER OF COILS,
--BUT WITH COILS THAT ARE TWICE AS WIDE IN DIAMETER.

WILL THE SPRING STRETCH FROM ITS NATURAL LENGTH, MORE, LESS, OR THE SAME AMOUNT UNDER THE SAME WEIGHT? (ASSUME THE MASS OF THE SPRING IS NEGLIGIBLE COMPARED TO THE MASS OF THE WEIGHT.) WHY DO YOU THINK SO?

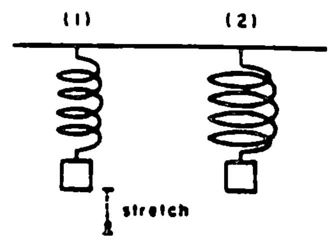

Figure 4
A target problem for experts (Clement, 1983)

of a second–order isomorphism as a means of model refinement provides an interesting abstraction for problem solving by analogy, but much work remains.

Analogies used by expert problem solvers

In contrast with relatively naive subjects used in the previous two projects, Clement (1981, 1982) describes the use of analogy in problem solving by a small sample (10 subjects) selected on the basis of creative problem solving skills in technical fields (5 physicists, one a Nobel laureate; 3 mathematicians; 2 computer scientists). Subjects were presented with a difficult physics problem (see Figure 4) and were encouraged to "think aloud" while attempting a solution. In contrast with retrieved analogies implicated in previous studies, it appears that Clement's subjects spontaneously *generate* many source problems, perhaps under heuristic control. Spontaneous use of analogies occurs in 70% of these subjects. For example, many subjects report spontaneously generating either simpler versions of the spring problem (e.g., uncoiling the spring and testing its elasticity by bending) or retrieving structurally similar problems in quite different domains (e.g., complicated chemical structures). Similar reasoning strategies have been observed in experts' solution attempts on a relatively difficult mathematical problem (Clement, 1984). On the basis of extensive protocol analysis, Clement argues for a process model of analogical reasoning consisting of four components:

1. *generating* an analogous problem,

2. *confirming* key relations between analogous problems,
3. *comprehending* the generated analog at a level of detail enabling predictions,
4. and *transferring* both predictions and general knowledge (e.g., key relations or a method of attack) from source to target problems.

The latter three of these processes are described as occurring in any order and, rather than leading to a solution instantaneously, problem solving using analogies is described as a difficult, time-consuming process akin to scientific theory formation and verification. As with the previous two research projects, Clement's work can be examined in terms of access to an analogical source (item 1, above) and its subsequent use (items 2 through 4, above).

Clement (1983) describes three mechanisms for access to a source analog, each of which is seen with differing frequency in his sample of expert problem solvers. These include: "generative transformation" of the existing problem representation to give a novel but simpler problem, retrieval of an analogy by "associative leaps" to a more familiar domain, and identification of an analogous problem via a shared "abstract principle." Detailed examination of analogies evident in verbal protocols shows that subjects most often generate an analogous case (59% of salient analogies were of this type), less frequently retrieve an analogous case through an "associative leap" (25% of this type), and infrequently retrieve an abstract principle which subsumes an analogous case (3% of this type).

Clement's characterization of these analogy types is interesting in that it documents purposeful or strategic generation of analogous cases, in contrast to the implicit or unintentional model of access to analogies evident in the accounts of Gentner or Holyoak (see previous sections). The deliberate character of recognition or access is quite clear in subjects' verbalizations during problem solving, even for analogies retrieved by associative leaps:

> I feel as though I'm reasoning in circles and I think I'll make a deliberate effort to break out of the circle somehow... what else stretches... like rubber bands, molecules, polyesters... (Clement, 1983, p. 13).

Clement does offer hypothetical explanations of how the intentional search for a source problem might occur. Intentional retrieval by an associative leap is described as requiring the subject to attend to salient aspects of the target problem (e.g., relations which are involved in causal constraints on the solution, like stretching). In this case, heeded relations should serve as origins of activation within memory, and candidate source cases which share these relations may reach an active state. Generation of a source problem, in contrast, is described as requiring the subject to attend to and intentionally alter presumably fixed aspects of the target problem (e.g., the number of coils). In this case, the altered problem is often simpler than the original. Aspects of the altered problem then serve as the origin of memory activation which may result in known solution methods for the simpler problem being retrieved by virtue of reaching an active state.

While interesting, Clement's proposals for access in the various analogy types are underspecified: he suggests increasing activation as a retrieval mechanism but falls short of any commitment to the representational structures held by expert problem solvers in his sample. Thus, distinctions between abstractions, experienced or generated problem cases, and associated solution strategies are difficult to evaluate. It is far from clear what subjects can be assumed to know about in the task domain, or whether different forms of knowledge might be differently accessible.

Clement also describes the manner in which expert problem solvers might use an existing analogy. Two subprocesses contribute to confirmation of an analogy (item 2, above). First, "key relationships" between the source and target are determined by an unspecified matching process which is apparently sensitive to problem-relevant areas of structural correspondence. Second, these relationships can be evaluated by using "bridging analogies" intermediate between the target and source problems. As an example of such a bridge, Clement (1982) describes a subject who evaluates the analogous stretching relation between the spring (target) and rod (source) by generating a third situation in which the spring has square coils (a bridging analogy). This allows the subject to qualify the partial similarity of stretching in the rod and the helical spring by noticing the importance of twisting as a contributor to how the spring actually works.

Comprehension of an existing analogy (item 3) may be facilitated by generation and confirmation of "extension analogies" or "extreme cases." In either case, the subject is faced with a new problem (the source) which resists understanding. Extension analogies are a recursive application of analogical reasoning to aid understanding of a poorly understood source, effectively treating that source as a new target problem; extreme cases are simplifications of a troublesome source by considering key problem components in the limiting case. As an example of an extreme case to aid in comprehension of the bending rod analogy, one of Clement's subjects generated a case in which the bending rod was very short, reasoning through physical intuition that an extremely short rod would bend very little. These findings suggest that multiple analogies are regularily used by expert problem solvers, forming chains of analogical inference which lend support to a line of reasoning which develops over time. Hence, problem solutions by analogy are not derived instantaneously but must be achieved with some effort by subjects.

In summary, Clement presents a description of analogical reasoning in problem solving that is compelling not only for suggestive process specification but also for its basis in non-trivial problem solving at expert levels of performance. As suggested by Gentner (1983a), increasing expertise in a problem solving domain may be strongly interrelated with subjects' abilities to access and use analogies for problem solving. Juxtaposed with studies of problem solving by assumed novices in which spontaneous successful use of analogies may be infrequent (Gick and Holyoak, 1980, 1983; Gentner and Gentner, 1983), Clement's finding that a majority of experts make liberal use of varied analogies is enticing. Since Clement's subjects were presumably experts at the time of testing, it would be very interesting to examine the role of analogies in problem solving as that expertise was acquired. This is the gist of Gentner's (1983a) analogical shift conjecture.

Analogies in problem solving: some conclusions

Looking back over introductory materials and the three research programs discussed in preceding sections, a natural division can be drawn between access to candidate analogical sources and their use. In turn, access and use appear to be sensitive to individual and situational factors which have been partially addressed by different research projects. We are now in a position to attempt a summarization of psychological studies of analogical transfer in problem solving.

Regarding access to or recognition of a source analog, it is important to distinguish between accounts of *how* retrieval takes place and *what* is actually retrieved. In the abstract, studies examined above concur over how successful retrieval of an analogical source takes place: heeded aspects of the target problem serve to activate similar aspects of candidate sources in long–term memory, resulting in one or more of these sources reaching an active state for consideration in further problem solving. Conceivably, the memory mechanisms which underly such retrieval could be considered invariant over subjects, tasks and reasoning contexts. Variety in problem solving performance, then, would be attributable to knowledge possessed by subjects or task characteristics, rather than to underlying retrieval mechanisms. This seems a reasonable assumption, although even within invariant memory mechanisms a clearer specification of supporting memory structures (e.g., the representational grain of a convergence case or schema) and the processes which act upon them (e.g., spreading activation) would be both interesting and useful.

With regard to what in the target is heeded by subjects as the process of retrieval begins, consensus among researchers abruptly stops. A major distinction, made particularly clear by Clement's work, is between passive and active accounts of attending to the target problem statement. Gentner's structure mapping account appears the most passive, with any aspect of the target problem representation potentially contributing to activation of a candidate source. This passive account of heeding appears to have some support from the prompted recall experiments with related stories mentioned earlier (Gentner and Landers, 1985). Holyoak's focus on the pragmatics of analogy in problem solving provides an intermediate point of contrast on this passive/active continuum for heeding. According to Holyoak's summative activation account of retrieval, consideration of goal–relevant aspects of the target problem representation gives disproportionate activational "weight" to contextually salient aspects of candidate sources. Thus, the problem solving context itself provides constraints on what will be heeded in the target. This differs from what Holyoak describes as a context–free examination of problem representations in Gentner's structure mapping approach. As a final point of contrast, Clement documents what appears to be purposeful consideration of key relations (e.g., stretching) in the target problem. This include alteration of fixed relations (e.g., the number of coils in a spring) in an effort to recall previously experienced problems or to generate a simpler problem variation for which a known solution method might apply. It is important to understand that the distinction being made here is between active or passive consideration of the target problem with or without the intention of retrieving a candidate source: the actual retrieval mechanism in all three accounts (spreading activation) is the same.

As mentioned before, variation in what subjects attend to in the target problem representation likely depends upon their sophistication in the problem solving domain. This sophistication may take many forms, including the general store of knowledge about the domain, a conceptual vocabulary which supports effective abstraction of problem-specific details, or even problem solving confidence. This is made clear by contrasting Clement's description of expert problem solvers with findings from other research with relatively naive problem solvers. Additional evidence is provided by studies which contrast behaviors of experts and novices in problem categorization and problem solving for a particular domain. For example, studies in the domains of physics (Chi, Feltovich and Glaser, 1981; Larkin, 1982) and mathematics (Schoenfeld and Herrmann, 1982) have shown systematic differences in the representations used by subjects with different abilities. Experts appear to use more abstract representational terms which are associated with general solution strategies, while novices appear more attentive to superficial aspects of a presented problem. Thus, a plausible but abstract explanation for why expert problem solvers make more frequent and flexible use of analogies is that they are better able to attend to relevant information in the target problem and represent it in a fashion which makes activation of appropriate sources in memory more likely. This is somewhat in contrast with Holyoak's advocacy of schema induction to facilitate access and largely consistent with Gentner's analogical shift conjecture. Unfortunately, little is known about how representational differences between novice and expert problem solvers develop.

Having outlined a framework for how access might occur at test, we can now consider *what* a problem solver might be able to retrieve. Unfortunately, the issue of what is retrieved is often blurred in the studies examined above. For example, finding that spontaneous noticing of a military source presented at training occurs infrequently, Gick and Holyoak (1983) require a summarization task of their subjects and report facilitation of spontaneous retrieval at test. While some facilitative effect seems clear, it is not clear what subjects are actually retrieving in their problem solving efforts after the summarization task: one of the original military stories, an abstracted "schema" based on those stories, or even a reasoning set which encourages abstraction over recalled concrete instances. Plausible accounts of transfer could be constructed for any of the three situations. Clement's descriptive results (1983), although collected under different circumstances (i.e., no training manipulation) and within a different subject population, suggest that varied materials may be retrieved from long-term memory, if not generated directly through heuristic reasoning strategies.

Again, sophistication in the problem solving domain offers some explanation for individual differences in accessing a relevant source. In addition to heeding and representing relevant problem information, it appears likely that proficient problem solvers have a wider store of better organized problem solving experiences available in memory. Although this assertion is something of a truism (i.e., more knowledge leads to better performance), there is evidence for the psychological reality of "problem schemata" possessed by proficient problem solvers and the influence of these schemata on subjects' problem solving performance. For example, Hinsley, Hayes and Simon (1977) and Mayer, Larkin and Kadane (1984) give evidence for the psychological importance of problem schemata in the domain of algebra "word" or "story" problems typical of the mathematics curriculum in secondary schooling.

In brief, these studies show that proficient problem solvers can reliably group problems by type; categorization of a problem by type occurs before solution; solution methods are associated with types and are used to solve problems (this association constitutes a schema); and some subjects are able to use schemata to solve superficially dissimilar problems. Problem schemata, while certainly not the only possibility for an effective analogical source, offer what could be a bridging knowledge structure developmentally intermediate between specific problem solving cases and more abstract principles. This possibility was briefly explored in the preceding discussion of Gentner's analogical shift conjecture.

As a final contributor for understanding effective access to source analogs, task specific factors should be considered. The distinction between training and test in many of the studies described in previous sections has implications which extend beyond the experimental manipulations used in those studies. Since candidate sources are encoded at training time (or during instruction or practice in a real–world setting), the extent to which the situation at test is comparable to training conditions should play an important role in facilitating or inhibiting effective access. In addition to evidence already discussed (Holyoak and Koh, 1985; Gentner and Landers, 1985), Ross (1982) has shown both positive and negative transfer as a result of systematic variation in task materials at training and test. Training and test comparability can be determined at many levels, including not only representational similarity of presented problems, but also similarity in the reasoning and motivational contexts across situations. Thus it seems reasonable to assume that instating a story comprehension task at training but a problem solving task at test might yield different spontaneous transfer than if a similar task were used in both situations. Systematic variations in reasoning contexts were not undertaken in the studies examined in previous sections. In general, effective access at test must overcome whatever encoding specificity is retained from the training context (Tulving and Thomson, 1973; Norman and Bobrow, 1979). Again, subjects' sophistication in the problem domain(s) might influence what they attend to in the problem statement (at training or test) as well as what might be available for access at test. Clement's finding that experts strategically manipulate problem materials in an apparent effort to recall applicable cases is suggestive for proficient problem solving at test; the effectiveness of Gick and Holyoak's requirement that subjects summarize multiple source problems is suggestive for similar strategies during training. Thus, the undesirable effects of encoding specificity might be overcome either at test (during retrieval) or at training (during encoding).

Again, considering access from the vantage of *how* retrieval occurs and *what* is retrieved, a general perspective on access to source analogs is possible. Whether a problem solver does retrieve a source and what that source consists of will depend upon several issues.

1. The *mechanism* through which retrieval arises plays an important, but possibly invariant role in supporting the influence of other factors.

2. How the subject *attends* to the target problem statement will determine the focus for a memory retrieval process. Heeding irrelevant information in the target or representing heeded information in an overly specific fashion reduces the likelihood of retrieving a relevant source – a dilemma which may be avoided by more sophisticated subjects.

3. What is *available* for retrieval in the subject's long–term memory obviously delimits what might be retrieved. Availability depends upon the scope of a subject's previous experience and the manner in which those experiences were committed to memory.

4. The degree to which *reasoning contexts* match at training and test can exert a facilitative or inhibitory influence on effective access. Subjects' reasoning strategies at either time may serve to mediate the influence of context matching.

Assuming that subjects can access a source analog, we still must account for the conditions under which subjects are able to *use* the retrieved source. As with accounts of access, there seems to be general agreement over how a retrieved analog might be used: access yields an initial partial mapping between source and target, and this mapping is incrementally extended subject to the demands of the target problem. This extension could be described as a process of hypothesis formation and confirmation. A relation known to be true in the source is hypothesized to exist in the target, and inferential activities in the target serve to confirm or disconfirm this hypothesis. Again, the distinction between *how* a source is used and *what* that source consists of is important and will be mediated by subject and situational factors.

In considering how extension and confirmation of a mapping might occur, Gentner's (1982) notion of base specificity is clearly a relevant subject factor. If systematicity (or causal–relevance) is to dominate analogical reasoning, subjects must possess systematic knowledge of the source domain and (as in Holyoak's pragmatic orientation) be able to choose aspects of this knowledge for extension to the target domain. Thus, subjects must have an adequate understanding of the source if transfer is to be expected: retrieving another ill–specified problem when faced with difficulties in an existing target problem will not be very productive. Some evidence for this situation may be found in Gentner and Gentner's (1983) description of subjects' attempts to use either a fluid or crowd analogy in reasoning about electrical circuits. Without clearly understanding how fluids behave in a closed system which is influenced by gravitational forces, it is not surprising that some subjects find it difficult to transfer predicted behaviors from the fluids domain into the target circuit problems. Clement's (1981, 1982) description of experts' use of bridging analogies, extension analogies and extreme cases to support their understanding of relations in the source domain suggests that a clear understanding of the source requires effort even from quite sophisticated problem solvers.

Lest understanding of the source be over–emphasized, we should also note the importance of having a sufficient understanding of the target problem domain so that particular relations in the source can be selected as relevant hypotheses and then confirmed. This corresponds roughly to Holyoak's notion of purpose in using an analogy: consideration of many possibly relevant relations in the source is constrained by attention to relevant but troublesome aspects of the target problem. Beyond providing selective bias, some understanding of the target domain must also be assumed if extended relations (hypotheses) are to be confirmed. Thus a relation which has some explanatory or predictive force in the source problem must be verified within the confines of the target problem, a process which presupposes a subject's ability to judge the validity of imported relations. In

Holyoak's terminology (1984a), this corresponds to subjects' capacities for distinguishing between mapped relations which are identities, structure–preserving differences or structure–violating differences.

This leads to a paradoxical situation: subjects must have a reasonable understanding of their goals in the target problem and the constraints surrounding those goals if a retrieved source is to be effectively used. We might wonder why, if subjects must understand the target problem so well, analogical reasoning is required at all? This apparent paradox is less troublesome after considering the circumstances under which analogies are used in problem solving. As Holyoak (1984a) suggests, problem solvers may reach an impasse in their attempts to generate a reasonable solution in the target problem, even though they may possess sufficient knowledge to recognize or verify a reasonable solution plan if it were available. Considering the extent of subjects' knowledge of the target problem domain, Clement's work is once again instructive by showing that even highly sophisticated problem solvers make frequent use of analogies when confronted by a sufficiently difficult target problem. It is clear that Clement's subjects understand the spring problem in sufficient detail that they can identify salient problem components (e.g., "stretching" or the number of coils in a spring). They use these judgments as the basis for an active process of analogical reasoning: they produce hypotheses about qualitative and quantitative aspects of the target by considering better understood relations in a source, and then attempt to confirm or disconfirm those hypotheses by determining whether supportive reasoning in the source is also true in the target.

If analogical reasoning in problem solving is to be understood as a process of hypothesis generation and testing, *what* is retrieved as a source is again of primary importance. Hypotheses available through context–relevant consideration of a retrieved source will vary in abstraction and possible utility (Gentner's notions of scope and validity) as a function of what the reasoner retrieves: problem–specific cases (either from the same or different conceptual domains), problem schemata or abstractions developed over experience with many specific cases, or principles which may be domain independent. According to Holyoak (1984b) there may be some tradeoff between accessibility and utility of a source, with abstract schemata being more easily accessible than specific cases but possibly less powerful when applied within the confines of a specific target problem. Gentner (1983a) argues from the other side of the abstraction spectrum that specific cases may be too complicated to permit "analysis" (hypothesis generation), particularly for novice problem solvers. Both accounts prefer some optimal level of abstraction for the retrieved source if it is to be useful in achieving a target solution. Where that optimal level can be found remains an open question.

Beyond utility of a retrieved source in terms of abstraction, it appears that qualitatively different information may be made available in different instances of problem solving by analogy. According to the previous account of heeding problem relevant information, what is retrieved in analogical reasoning may be determined by the circumstances which lead a problem solver to an impasse in the target problem. Thus if a subject is "stuck" while trying to find an operator whose application would yield movement towards a goal state, she might be reminded of previous problem solving episodes which bear strong similarity by virtue of having similar goal states, constraints or types of operators. This is the gist of Holyoak's (1985) contextual–relevance argument. Alternately, impasse may occur in achieving an adequate representation of a target problem, well before consideration of how to apply available operators. Clement's documentation of problem solvers feeling "stuck" and concentrating on key relations in the problem statement may be indicative of the latter situation. In either case, we might expect that qualitatively different forms of analogy could occur: in the first case (impasse over operators) an actual solution plan from a previous problem solving experience might be retrieved and used; in the latter case (representational impasse) an entirely different problem model might be retrieved and used (e.g., that of a bending saw or compressible foam – both reported by Clement's subjects [1983]). Unfortunately, distinctions between problem solving cases, problem schemata, problem models, and abstract principles remain elusive.

Several generalizations are possible regarding use of a retrieved analogy.

1. *How* sources are used in problem solving can be described as an incremental process of hypothesis generation and confirmation. Problem solution by analogy, as documented by Clement, may not be instantaneous.
2. A subject's knowledge of the source should be seen as providing an *upper bound* on what may be transferred into the target. An acontextual consideration of the source probably cannot predict what *will* be transferred in practice.
3. A subject's knowledge of the target goes further in predicting what will be transferred by:
   a. allowing context–relevant consideration of the source in generating hypotheses, and
   b. allowing validation or confirmation of these hypotheses.
4. *What* is available for transfer may vary depending upon the circumstances under which access occurs (problem impasses) and the abstractness of a retrieved source. Optimal abstractness balances ease of access against utility.

6. Contributions of psychology and related disciplines

Before examining computational approaches to analogical reasoning, and in particular the use of analogy in learning, it seems appropriate to ask what the various disciplines discussed in previous sections of this paper can contribute to work in AI. In particular, four issues can be extracted from these other literatures which are of significance for AI:

1. the importance and prevalence of analogical reasoning in human thinking generally,
2. likely cognitive structures underlying analogical reasoning,
3. implications of these structural arrangements for process models of analogical reasoning,
4. and the role which analogical reasoning plays in learning.

With regard to the importance of analogical reasoning, opposing arguments have been described for each of the academic disciplines discussed in previous sections: linguistics, education, philosophy and psychology. Opponents of analogy and metaphor see figurative expression as, at best, ornamentation of literal expression and, at worst, as a misleading disguise for poor understanding of the topic under consideration. Proponents, on the other hand, find analogy or metaphor an indispensible and ubiquitous constituent of human cognition. In parallel with this controversy, comparison and interaction based–theories of figurative expression allow metaphor differing access to center stage in cognition.

Out of this varied landscape, an emerging path can be found, although by no means would all participants follow calmly. First, metaphor does indeed appear ubiquitous in human cognition; second, metaphor can be seen to provide a *powerful problem setting* in cultural, socio-political and scientific contexts; and third, metaphor may well provide an *epistemological bridge* both for maintaining some form of linguistic economy and for acquiring novel information, possibly through a mechanism of hypothesis generation and verification very similar to that described for scientific theory formation. Central to each of these observations is the notion that metaphor and analogy are not cognitive or expressive oddities but are important in many aspects of cognition, perhaps underlying the very manner in which we perceive and reason about our environment.

Whether taken as a general ability, synaesthesia, homeostatic transformation, stimulus generalization or "information processing," analogical reasoning has long fascinated psychologists. It is from this academic discipline that one might expect significant contributions to research in AI. However, the psychological approaches to analogy and metaphor discussed in preceding sections raise more questions than they answer. Theoretical orientations towards the subject matter are varied and may yield consensus on relatively little: analogical reasoning depends upon some sort of similarity, from the stimulus level through to general intellectual abilities. Exemplary psychological studies are likewise varied, but provide some interesting contrasts: processing specificity can often be purchased at the cost of representational generality; relatively acontextual, comparison–based views of concept similarity may be insufficient for convincing accounts of metaphor and analogy; and effective transfer of problem solving skills between domains depends in a complicated fashion upon both subject and situational factors. While not conclusive, psychological studies of metaphor and analogy provide interesting suggestions and constraints for computational efforts which are just beginning. Contributions are evident both for knowledge representations and processes supporting analogical reasoning.

A clear trend is evident towards interest in more flexible, "schematic" or "structured" forms of representation for knowledge accessed during analogical reasoning. As argued convincingly by Ortony *et al.* (1978) nearly a decade ago, static featural or propositional notations appear implausible as underlying supports for flexible language use, typically leading researchers to postulate exceptional processing mechanisms to include figurative capabilities within their models. However, the goal of an integrated account of literal and figurative cognitive behaviors could still prove elusive for both psychologists and AI researchers: there may be no reason to assume that the advocated schema or frame–based knowledge representation techniques will make researchers more insightful. In practice, concepts developed using less flexible (and less complex) propositional representations could still prove useful even if such representational commitments are not pursued in AI research.

For processes involved in analogical reasoning, there appears to be general agreement over the importance of *recognition* and retrieval of potential source analogs, *elaboration* of the ground or mapping which relates the domains in question, transfer and *evaluation* of confirmed or suspected corresponding relations, and *consolidation* of information gained during successful analogical reasoning in a form which will be subsequently available. While these processes depend in some measure on the unsettled issue of a representation for the memory structures involved, much of the recent work on analogical reasoning in problem solving utilizing structured (or schematic) representations seems particularily promising. In particular, Gentner's (1982) descriptive taxonomy of varied metaphors and advocacy of "systematicity" (Gentner, 1983a) in structure mapping provide useful notions of what constitutes a "good analogy" for varied purposes and what sorts of relational information might be preferred when extending such an analogy. This notion of partial similarity at some level of representational abstraction is supported by Holyoak's (1984a) discussion of the importance of causal relevance and the necessity of forming abstract problem solving schemata if analogies are to be effectively recognized and applied. Finally, Clement (1982, 1983) provides useful insights into the manner in which expert problem solvers utilize analogical reasoning in arriving at solutions to nontrivial problems. Chains of supporting analogies, some of which are actively generated by simplification of a problem at hand, are suggestive of how analogical reasoning may fit into an ongoing problem solving context. As may be evident in concluding remarks for the section on analogy in problem solving, there are many fruitful avenues of exploration opened by these studies to which computational approaches might make considerable contributions.

Although the role of analogical reasoning in learning has already been mentioned, the enthusiasm in machine learning circles for "learning by analogy" (Mostow, 1983) recommends a close look at how analogy might contribute to learning. In fact, much of the psychological literature described above suggests that analogy simply provides grist for a relatively straightforward inductive mill (e.g, Holyoak's eliminative induction), a suggestion hardly worthy of revel in AI circles where work on inducing concepts from examples has a long history (Dietterich and Michalski, 1983; Angluin and Smith, 1984). However, it may be that learning in the context of analogical reasoning provides different possibilities than traditional computational studies of induction. These studies generally assume a well-behaved (e.g., noise free) environment attended by a benevolent teacher where the learner receives both explicit (in the form of teacher suggestions) and implicit (in the form of a carefully chosen conceptual vocabulary) guidance. Research problems in which the learner is assumed able to spontaneously recognize and elaborate analogies between previous and current experiences will likely force a wider consideration of how the learner's reasoning is focused. These issues will be discussed more fully in the conclusion to the second part of this survey. In summary, perhaps the notion of metaphor and analogy as a problem setting which constrains hypothesis generation and suggests illuminating experimentation will prove one of the more important contributions of the disciplines surveyed above for computational approaches to analogical reasoning.

References

Angluin, D. and Smith, C.H. (1983) Inductive inference: theory and methods. *Computing Surveys 15*(3), 237–269.

Billow, R.M. (1977) Metaphor: a review of the psychological literature. *Psychological Bulletin 84*(1), 81–92.

Black, M. (1962) *Models and metaphors*, Cornell University Press: Ithaca, New York.

Black, M. (1979) More about metaphor. In A. Ortony (Ed.) *Metaphor and thought*, Cambridge University Press: Cambridge, MA.

Boden, M. (1980) *Artificial intelligence and natural man*, Basic Books: New York.

Boyd, R. (1979) Metaphor and theory change: what is "metaphor" a metaphor for? In A. Ortony (Ed.) *Metaphor and thought*, Cambridge University Press: Cambridge, MA. 356–408.

Brown, A.L., Bransford, J.D., Ferrara, R.A., and Campione, J.C. (1983) Learning, remembering, and understanding. In P.H. Mussen (Ed.), *Handbook of Child Psychology (Fourth Edition): Vol. III. Cognitive Development*, 77–166. New York: John Wiley and Sons.

Brown, A.L. and Campione, J.C. (1985) Three faces of transfer: Implications for early competence, individual differences, and instruction. In M. Lamb, A.L. Brown, and B. Rogoff (Eds.), *Advances in Developmental Psychology*, Hillsdale, New Jersey: Lawrence Erlbaum and Associates.

Carbonell, J.G. (1981) Invariance hierarchies in metaphor interpretation. *Proceedings of the third annual meeting of the Cognitive Science Society*, 292–295.

Carrol, J.M. and Mack, R.L. (1982) Metaphor, computing systems, and active learning. RC–9636(42545), IBM Watson Research Center: Yorktown Heights, New York.

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

Chomsky, N. *Aspects of the theory of syntax.* MIT Press: Cambridge, Mass.

Clement, J. (1981) Analogy generation in scientific problem solving. In *Proceedings of the third meeting of the Cognitive Science Society*, 137–140.

Clement, J. (1982) Analogical reasoning patterns in expert problem solving. In *Proceedings of the fourth meeting of the Cognitive Science Society*, 79–81.

Clement, J. (1983) Observed methods for generating analogies in scientific problem solving. Presented at the annual meeting of the American Educational Research Association, Montreal, Canada.

Clement, J. (1984) Non-formal reasoning in experts' solutions to mathematics problems. Presented at the annual meeting of the American Educational Research Association, New Orleans.

Dietterich, T.G. and Michalski, R.S. (1983) A comparative review of selected methods for learning from examples. In R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.) *Machine learning: an artificial intelligence approach*, 41–82.

Dreistadt, R. (1968) An analysis of the use of analogies and metaphors in science. *The Journal of Psychology 68*, 97–116.

Dreistadt, R. (1969) The use of analogies and incubation in obtaining insights in creative problem solving. *The Journal of Psychology 71*, 159–175.

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

Forbus, K.D. and Gentner, D. (1983) Learning physical domains: towards a theoretical framework. In *Proceedings of the Machine Learning Workshop*, 198–202.

Forbus, K.D. and Gentner, D. (1986) Learning physical domains: Toward a theoretical framework. In R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach, Volume II.* California: Morgan Kaufmann Publishers, Inc., 311–348.

Gentner, D. (1982) Are scientific analogies metaphors? In D. Miall (Ed.) *Metaphor: problems and perspectives.* Humanities Press: New Jersey, 106–132.

Gentner, D. (1983a) Structure-mapping: a theoretical framework for analogy. *Cognitive Science 7*, 155–170.

Gentner, D. (1983b) Ninety years of mental metaphors. In *Proceedings of the Fifth Annual Meeting of the Cognitive Science Society.*

Gentner, D. and Gentner, D.R. (1983) Flowing waters or teeming crowds: mental models of electricity. In D. Gentner and A.L. Stevens (Eds.) *Mental models.* Lawrence Erlbaum Associates: Hillsdale New, Jersey, 99–129.

Gentner, D. and Landers, R. (1985) Analogical reminding: A good match is hard to find. In *Proceedings of the International Conference on Systems, Man and Cybernetics*, Tucson, Arizona, 1–7.

Gick, M.L. and Holyoak, K.J. (1980) Analogical problem solving. *Cognitive Psychology 12*, 306–355.

Gick, M.L. and Holyoak, K.J. (1983) Schema induction and analogical transfer. *Cognitive Psychology 15*, 1–38.

Grudin, J. (1980) Processes in verbal analogy solution. *Journal of Experimental Psychology: Human Perception and Performance 6*, 67–74.

Halasz, F. and Moran, T.P. (1982) Analogy considered harmful. In *Proceedings of the Conference on Human Factors in Computer Systems*, March 15–18, National Bureau of Standards, Gaithersburg, MD.

Hayes, J.R. and Simon, H.A. (1977) Psychological differences among problem isomorphs. In N.J. Castellan, Jr., D.B. Pisoni and G.R. Potts (Eds.), *Cognitive Theory*, v. 2, New Jersey: Lawrence Erlbaum Associates, 21–41.

Haynes, R. (1975) Metaphor as interactive. *Educational Theory 25*, 272–277.

Hinsley, D.A., Hayes, J.R., and Simon, H.A. (1978) From words to equations: meaning and representation in algebra word problems. In P.A. Carpenter and M.A. Just (Eds.) *Cognitive processes in comprehension.* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Hoffman, R.R. (1980) Metaphor in science. In Honeck and Hoffman (Eds.) *Cognition and figurative language.* Lawrence Erlbaum Associates: Hillsdale, New Jersey, 393–423.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., and Thagard, P. (1985) *Induction: Processes of inference, learning, and discovery.* (in preparation).

Holyoak, K.J. (1984a) Mental models in problem solving. In J.R. Anderson and S.M. Kosslyn (Eds.) *Tutorials in Learning and Memory.* New York: W.H. Freeman and Co., 193–218.

Holyoak, K.J. (1984b) Analogical thinking and human intelligence. In R.J. Sternberg (Ed.) *Advances in the psychology of human intelligence 2.* Lawrence Erlbaum Associates: Hillsdale, New Jersey, 199–230.

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

Holyoak, K.J., Junn, E.N. and Billman, D.O. (1984) Development of analogical problem–solving skill. *Developmental Psychology 55*, 2042–2055.

Holyoak, J. and Koh, K. (1985) Surface and structural similarity in analogical transfer. (in preparation, described in Holyoak, 1985).

Honeck, R.P. (1980) Historical notes on figurative language. In R.P. Honeck and R.R. Hoffman (Eds.) *Cognition and Figurative Language.* Lawrence Erlbaum Associates: Hillsdale, New Jersey, 25–46.

Hunt, E. (1973) Quote the raven? Nevermore! In L.W. Gregg (Ed.) *Knowledge and cognition.* Lawrence Erlbaum Associates: Hillsdale, New Jersey.

Johnson, J. (1984) Levels of processing in metaphor comprehension. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, 193–197.

Johnson, M. (1981) *Philosophical Perspectives on Metaphor*, Minneapolis: University of Minnesota Press.

Lakeoff, G. and Johnson, M. (1980) *Metaphors we live by.* Chicago University Press: Chicago, Ill.

Larkin, J.H. (1982) The role of problem representation in physics. In D. Gentner and A. Stevens (Eds.) *Mental models.* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Lester, F.K. (1982) Building bridges between psychological and mathematics education research on problem solving. In F.K. Lester and J. Garofalo (Eds.) *Mathematical Problem Solving: Issues in Research*, Philadelphia: The Franklin Institute Press, 55–85.

Luger G.F. and Bauer M.A. (1978) Transfer effects in isomorphic problem situations. *Acta Psychologica 42*, 121–131.

Malgady, R.G. and Johnson, M.G. (1980) Measurement of figurative language: semantic feature models of comprehension and appreciation. In R.P. Honeck and R.R. Hoffman (Eds.) *Cognition and Figurative Language*, Lawrence Erlbaum Associates: Hillsdale, New Jersey, 239-258.

Mayer, R., Larkin, J., and Kadane, J. (1984) A cognitive analysis of mathematical problem–solving ability. In R. Sternberg (Ed.) *Advances in the Psychology of Human Intelligence 2.*

Miller, G.A. (1979) Images and models, similes and metaphors. In A. Ortony (Ed.) *Metaphor and Thought*, Cambridge University Press: Cambridge, MA. 202-250.

Miller, R.M. (1976) The dubious case for metaphors in educational writing. *Educational Theory 26*, 174-181.

Mostow, J. (1983) 1983 international machine learning workshop: an informal report. *SIGART*, 24-31.

Newell, A. (1983) The heuristic of George Polya and its relation to AI. In R. Groner, M. Groner and W.F. Bischof (Eds.) *Methods of Heuristics*. LEA, 195-243.

Norman, D. and Bobrow D.G. (1979) Descriptions: an intermediate stage in memory retrieval. *Cognitive Psychology 11*, 107–123.

Novick, L.R. (1985) Analogical transfer in expert and novice problem solvers. Dissertation in progress, Stanford University.

Oppenheimer, R. (1956) Analogy in science. *American Psychologist 11*, 127-135.

Ortony, A. (1975) Why metaphors are necessary and not just nice. *Educational Theory 25*, 45-53.

Ortony, A. (1976) On the nature and value of metaphor: a reply to my critics. *Educational Theory 26*, 395-398.

Ortony, A., Reynolds, R.E. and Arter, J.A. (1978) Metaphor: theoretical and empirical research. *Psychological Bulletin 85*, 919-943.

Paivio, A. (1979) Psychological processes in the comprehension of metaphor. In A. Ortony (Ed.) *Metaphor and Thought*, Cambridge University Press: Cambridge, MA. 150-171.

Palmer, S.E. (1978) Fundamental aspects of cognitive representation. In E. Rosch and B.B. Lloyd (Eds.) *Cognition and categorization.* Lawrence Erlbaum Associates: Hillsdale, New Jersey, 259–303.

Parker, R.E. and Kiselewich, S.J. (1984) The modelling of human cognitive decision processes in the intelligent machine model (TIMM). Presented at the annual meetings of the American Control Society, June, 1984, San Diego, CA.

Petrie, H.G. (1979) Metaphor and learning. In A. Ortony (Ed.) *Metaphor and Thought*, Cambridge University Press: Cambridge, MA. 438-461.

Pollio, M.R., Barlow, J.M., Fine, H.J. and Pollio, M.R. (1977) *Psychology and the poetics of growth: figurative language in psychology, psychotherapy, and education.* Lawrence Erlbaum Associates: Hillsdale, New Jersey.

Polya, G. (1945) *How to solve it.* Princeton University Press: Princeton, NJ.

Reddy, M. (1979) The conduit metaphor: a case of frame conflict in our language about language. In A. Ortony (Ed.) *Metaphor and thought.* Cambridge University Press: Cambridge, England.

Reynolds, R.E. and Schwartz, R.M. (1983) Relation of metaphoric processing to comprehension and memory. *Journal of Educational Psychology 75*(3), 450-459.

Rumelhart, D.E. and Abrahamson, A.A. (1973) A model for analogical reasoning. *Cognitive Psychology 5*, 1-28.

Rumelhart, D.E. (1979) Some problems with the notion of literal meanings. In A. Ortony (Ed.) *Metaphor and Thought*, Cambridge University Press: Cambridge, MA, 78-90.

Rumelhart, D.E. and Norman, D.A. (1981) Analogical processes in learning. In J.R. Anderson (Ed.) *Cognitive skills and their acquisition*, Lawrence Erlbaum Associates: Hillsdale, New Jersey, 335-359.

Schoenfeld, A.H., and Herrmann, D.J. (1982) Problem perception and knowledge structure in expert and novice mathematical problem solvers. *Journal of Experimental Psychology 8*, 484–494.

Schön, D.A. (1979) Generative metaphor: a perspective on problem- setting in social policy. In A. Ortony (Ed.) *Metaphor and Thought*, Cambridge University Press: Cambridge, MA, 254-283.

Sternberg, R.J. (1977) *Intelligence, information processing and analogical reasoning: the componential analysis of human abilities.* Lawrence Erlbaum Associates, Publishers: Hillsdale, New Jersey.

Tulving, E. and Thomson, D.M. (1973) Encoding specificity and retrieval processes in episodic memory. *Psychological Review 80*, 352–373.

Tversky, A. (1977) Features of similarity. *Psychological Review 84*(4), 327-352.

Verbrugge, R.R. and McCarrell, N.S. (1977) Metaphoric comprehension: studies in reminding and resembling. *Cognitive Psychology 9*, 494-533.

Werner, H. and Kaplan, B. (1963) *Symbol Formation: An Organismic- Developmental Approach to Language and the Expression of Thought.* John Wiley and Sons: New York.