

①
FILE COPY

A Framework for a Theory of Mapping

Allan Collins and Mark Burstein
BBN Laboratories Inc.

for

Contracting Officer's Representative
Judith Orasanu

BASIC RESEARCH LABORATORY
Michael Kaplan, Director



U. S. Army

Research Institute for the Behavioral and Social Sciences

December 1987

Approved for public release; distribution unlimited.



AD-A191 071

U. S. ARMY RESEARCH INSTITUTE
FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Technical Director

WM. DARRYL HENDERSON
COL, IN
Commanding

Research accomplished under contract
for the Department of the Army

BBN Laboratories, Inc.

Technical review by

Dan Ragland

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



This report, as submitted by the contractor, has been cleared for release to Defense Technical Information Center (DTIC) to comply with regulatory requirements. It has been given no primary distribution other than to DTIC and will be available only through DTIC or other reference services such as the National Technical Information Service (NTIS). The views, opinions, and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other official documentation.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ARI Research Note 87-64	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD-A191071
4. TITLE (and Subtitle) A Framework for a Theory of Mapping		5. TYPE OF REPORT & PERIOD COVERED Interim Report March 86 - March 87
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Allan Collins and Mark Burstein		8. CONTRACT OR GRANT NUMBER(s) MDA903-85-C-0411
9. PERFORMING ORGANIZATION NAME AND ADDRESS BBN Laboratories Inc. Cambridge MA 02238		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2Q161102B74F
11. CONTROLLING OFFICE NAME AND ADDRESS U.S. Army Research Institute for the Behavioral and Social Sciences, 5001 Eisenhower Avenue, Alexandria, VA 22333-5600		12. REPORT DATE December 1987
		13. NUMBER OF PAGES 25
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) --		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE n/a
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release: distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) --		
18. SUPPLEMENTARY NOTES Judith Orasanu, contracting officer's representative. This material is to appear in S. Vosniadou and A. Ortony (eds.): <u>Similarity and Analogical Reasoning</u> .		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Analogy and Similarity Artificial Intelligence Cognitive Psychology Machine Learning		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This research note is divided into three main sections. The first distin- guishes the different kinds of information that are related by analogy and similarity mappings. The second section discusses the different contexts or tasks that give rise to mappings. The third section catalogues different solutions proposed or possible for each issue. The note concludes by arguing that the issue is not whether analogies are helpful or harmful, but what determines when they are helpful and when they are harmful.		

TABLE OF CONTENTS

1. INTRODUCTION	1
2. WHAT IS MAPPED	2
3. CONTEXTS IN WHICH MAPPINGS OCCUR	5
3.1 TYPE OF TASK	5
3.2 NUMBER OF ENTITIES COMPARED	7
4. ISSUES FOR A THEORY OF MAPPING	9
5. CONCLUSION	18

1. INTRODUCTION

The literature on similarity, analogy, and metaphor ranges over many different kinds of mappings. Some of the disagreements arise because researchers are talking about different kinds of mappings or the different contexts in which mappings are made. Our goal is to clarify the issues being addressed and the critical distinctions that need to be made. We will attempt to consider the entire territory over which the discussion of mapping arises, but no doubt we will miss some of the critical distinctions and issues.

We have divided the paper into three main sections. The first section distinguishes the different kinds of entities that are related by analogy and similarity mappings, and some of their more salient properties. The second section discusses the different contexts or tasks that give rise to mappings. The third section catalogues the set of issues we have identified in the literature, and identifies some of the different solutions proposed or possible for each issue. In a concluding section we briefly discuss the implications of this framework for research.

2. WHAT IS MAPPED

The hypothesis we offer is that there are three fundamentally different kinds of entities that are mapped: systems, concepts, and properties and that all the other kinds of mappings discussed in the literature are variations on one of the three

System Mapping. The mapping from the solar system to the atom that Gentner (Gentner, 1983) discusses is the classic example of a system mapping. In a system mapping it is critical to determine two types of mappings (Gentner, this volume).

1. Which components (i.e., concepts) in the source domain are mapped into which components in the target domain.
2. Which properties of each component (including relations between components) in the source domain are mapped into which properties in the target domain.

In the solar system/atom analogy, one first has to decide what components map (sun --> nucleus, planets --> electrons) and then what properties map (planets orbit the sun --> electrons orbit the nucleus).

Concept mapping. To answer the question (Collins, 1978) "Was Nixon a crook?" or to decide how likely Linda is to be a feminist bank teller (Tversky and Kahneman, 1980, Smith & Osherson, this volume) requires only a mapping across the properties of two concepts. There is no decomposition into components, as there is with a system mapping. So, in the case of Linda in Smith and Osherson's (this volume) account, you consider the properties of salary, education, and politics in the mapping process, comparing Linda and feminist bank tellers with respect to these properties.

Property mapping. The simplest kind of mapping specifies a particular property of two concepts for comparison, as when one judges whether an object 3 inches in diameter is more similar to a quarter or a pizza (Rips, this volume). (This example is actually a double mapping, discussed later under three-element comparisons, between a 3 inch object and a quarter, and between a 3 inch object and a pizza - system and concept mappings can also involve double mappings.) Property mappings differ from concept mappings in that the concepts are compared with respect to a particular property rather than with respect to many properties.

The critical distinction between these three kinds of mappings is that the system mappings involve component (or object) mappings as well as property mappings, that concept mappings involve multiple property mapping, and that property mappings involve individual properties of two concepts. The distinction between system and concept mappings is not entirely straightforward. For example, one elementary text we studied (Collins, Gentner, and Rubin, 1981) explained the composition of the earth by analogy to a peach. There is the crust which is analogous to the skin, the mantle analogous to the fruit, and the core analogous to the pit. This may appear to be a concept mapping, since it is a comparison of the properties of two concepts. But in fact it is a system mapping, since it requires first decomposing the earth and peach into their components (i.e., the three layers), and then comparing the properties of each pair of components (e.g., the skin and crust are both very thin), and their relations to each other. Thus the distinction between a system mapping and a concept mapping rests upon whether there is a two-stage process of first mapping an organized set of components and then the properties of each component (i.e., a system mapping) or a single-stage process of mapping properties (i.e., a concept mapping).

To give a second example of a system mapping that may be difficult to recognize, one might hypothesize (Collins and Michalski, 1987) that a bird's pitch depends on the length of the bird's neck, which is why ducks quack and geese honk, and more generally why small birds sing and big birds squawk (Malt and Smith, 1984). This hypothesis might be generated by analogy to the fact that human pitch (e.g., children vs. adult voices) depends on the length of the windpipe. To make the inference about birds by analogy to humans requires mapping windpipe length onto neck length, and human pitch onto bird pitch. Because the analogy involves both a mapping between their components (e.g., windpipes and neck) and a mapping of some of their components' properties (relative length), it is a system mapping. In this case the property mapped (e.g., "pitch is inversely related to length") is a relational property in Gentner's (this volume) terms or a mutual dependency in Collins and Michalski's (Collins and Michalski, 1987) terms.

There are a number of other kinds of mappings discussed in the literature which we think are special cases of these three kinds of mappings. We will briefly describe each.

Procedure mapping. VanLehn and Brown (VanLehn and Brown, 1980) discuss mapping between the addition and subtraction procedures we learn in school and different addition and subtraction procedures with Dienes blocks (which are wooden blocks in three denominations, units are small squares, tens are ten unit blocks long, and hundreds are ten by ten unit blocks). Similarly, Anderson and Thompson (this volume) describes mapping between the procedure for factorial and that for summorial. Mappings of procedures are essentially system mappings, where the components of the procedure must first be mapped (e.g., unit blocks onto the numbers in the right hand column, etc.), and the manipulations on those components are subsequently mapped like properties.

Problem mapping. Ross (this volume), Holyoak & Thagard (this volume), and Carbonell (Carbonell, 1986), among others, discuss mapping between a problem you are trying to solve and an earlier problem you have solved. This kind of mapping is frequently used in science texts where students solve new problems by referring back to the sample problems worked in the text. Gick and Holyoak (Gick and Holyoak, 1980, 1983, Holyoak and Thagard, this volume) discuss the analogy between a fortress problem, where an army must split up in small units to capture a fortress, and Duncker's (Duncker, 1945) ray problem, where a ray source must be split in order to kill a tumor without destroying healthy tissue around it. Problem mappings require mapping of components first (e.g. ray --> army units, tumor --> fortress), and so they are system mappings.

Story Mappings. Gentner and Landers (Gentner and Landers, 1985) and Ross (this volume) have studied mappings between stories. These again are simply system mappings, where it is necessary first to map the characters or objects from one story to the other and then the relations or events between these entities.

There are undoubtedly other kinds of mappings that are made, but we think they will all be variations of the three kinds of mappings we have identified.

3. CONTEXTS IN WHICH MAPPINGS OCCUR

Various tasks or real world demands require different kinds of reasoning when relating entities. Our taxonomy of contexts in which mappings occur consists of two dimensions: *type of task* and *number of entities compared*. The overall structure of the taxonomy of contexts is shown in Table 1.

These two dimensions, type of task and number of elements, define a space of possible contexts in which mappings are made. There may be some cells empty in the space, but most combinations are possible.

3.1 TYPE OF TASK

Type of task breaks down into three basic categories: comparative judgements, mappings, and conceptual combinations. We will briefly describe six different kinds of comparative judgements, and then two kinds of mappings. Last, we will briefly discuss conceptual combination. The comparative-judgment types are derived primarily from the Rips (this volume) and Linda Smith (this volume) papers. This may not be a complete list of comparison judgements, but it covers the types discussed in this volume.

A. Comparative Judgements

1. **Similarity judgement.** Judging how similar two entities are is a common task in psychological experiments (Tversky, 1977; Rips, this volume; Smith & Osherson, this volume; Barsalou, this volume). Smith and Osherson (this volume) and Collins and Michalski (Collins and Michalski, 1987) argue that similarity judgments affects the certainty of many inferences people make. Similarity judgments obviously can apply to pairs of systems, concepts, or properties.
2. **Typicality judgments.** Typicality has been studied in psychology since Rosch (Rosch, 1975), and plays much the same kind of role in plausible reasoning as similarity (Collins and Michalski, 1987). Rips (this volume) has shown convincingly that typicality and similarity judgment are not always made in the same way, so they must be distinguished in any theory. Like similarity, typicality applies to pairs of systems, concepts, or properties.
3. **Categorization judgments.** Rips (this volume) discusses the similarity theory

of categorization, which he rejects. In any case, categorization requires a comparison between properties of two entities, the thing to be categorized and the category. Categorization only applies to systems and concepts, not to single properties of concepts, except when they are treated as concepts in their own right.

- 4 **Identity judgments.** Linda Smith (this volume) raises the issue of making identity judgments between entities – that is comparing whether all their properties are the same. Of course, no two entities are ever exactly the same (e.g. her examples of identical elements are not quite the same darkness or shape), so it is necessary to learn what degree of variability of a property can be called the same. Identity judgements therefore depend on context.
- 5 **Overlap judgments.** None of the papers in this volume mention overlap judgments (e.g. whether therapists are psychiatrists), but logically if one includes categorization and identity judgments, then overlap and difference judgments must also be included. Evaluating a "some" statement (e.g. "Some women are doctors") requires making an overlap judgment (Meyer, 1970).
- 6 **Difference judgments** The question of whether two entities are different (e.g. "Are whales fish?") also involves a comparison of properties. Like categorization, identity, and overlap judgments, difference judgments are contextually defined. For example, whales and fish are different, but both are animals and can be treated as the same in some contexts, such as grouping things as plants and animals.

The last four of these judgments, categorization, identity, overlap, and difference correspond to the four possible relations between two circles in Venn diagrams, as shown in Figure 1.

B. Mappings

The other type of task that is referred to frequently in the literature is one of mapping properties, components, or both from the source domain to the target domain.

Property mapping. Most of the work on analogy (e.g. Anderson & Thompson, this volume, Gentner, this volume, Holyoak & Thagard, this volume) concerns itself with bringing properties (including relational properties) of objects in the source domain over into the target domain. A similarity or typicality judgment between the source and target is made before mapping a property over, and affects the certainty with which the property is believed to hold for the target domain. For example, before deciding that the pitch of birds depends on their neck length, based on an analogy to the human vocal tract, a person would compare humans and birds with respect to

their similarity, particularly on those properties related to sound production (in this case, properties of the relevant components, such as vocal cords and necks). A person's certainty about whether the property holds for birds depends on this similarity judgment.

Component mapping Sometimes in the mapping of two systems, whole components are introduced by the mapping. In the earth/peach analogy, the text introduced two new components of the earth to students (the mantle and the core) in the course of explaining the analogy. This same thing can occur when people consider an analogy in their own mind (Collins and Gentner, 1980). For example, in relating the texture of foods to materials science, one might notice that chewiness corresponds roughly to elasticity, crispness to ductility, and then wonder what juiciness corresponds to. One possibility is liquid-filled porosity, a critical concept in geology.

C. Conceptual Combinations

Smith and Osherson (this volume) raise the possibility that conceptual combination (feminist + bank teller --> feminist bank teller) is another task that a theory of mapping should address. We see conceptual combination, as they have modeled it, as primarily addressing the issue of how property mappings are combined when there is prior information about the properties involved in the target system. This becomes particularly important when learning or making predictions from multiple analogies, and in interpreting descriptive metaphors.

3.2 NUMBER OF ENTITIES COMPARED

Number of entities compared is the other dimension we have identified with respect to the contexts in which analogies occur. This can range from two, as in the earth/peach mapping, to four as in analogies like wolf dog tiger.cat, and the geometric analogies considered by Evans (Evans, 1968). Slightly different constraints operate in two, three, and four-element mappings, shown in Figure 2.

Two-element comparisons Many of the mappings discussed in the literature (e.g. Gentner's (this volume) water flow to heat flow mapping, and Holyoak & Thagard's (this volume) fortress problem to ray problem) are two-element mappings. In a two-element

mapping there is a source from which properties or components are mapped onto a target. There are no other concepts, even implicitly, that are compared in a two-element comparison.

Three-element comparisons A good example of a three-element mapping is the task used by Rips (this volume) where subjects had to decide whether a three inch object was more likely to be a pizza or a quarter. In many tasks that appear to be two-element mappings, there may be a comparison element implicit that subjects generate on their own in doing the mapping. For example, if you tell a child that a whale is a mammal, they may compare whether whales are more like mammals or fish, which is a three-element comparison. Three-element comparisons, therefore, compare properties of X to those of Y vs. Z.

Four-element comparisons. Standard analogy tests pose questions using the syntactic form W:X:Y:Z. We view such problems as falling into two categories, based on whether the analogy's interpretation depends on one or two comparisons. True four-element comparisons depend on both sets of comparisons, as in the analogy wolf:dog:tiger:cat. The within-group comparisons (e.g. wolf:dog and tiger:cat) determine the properties or dimensions along which the pairs differ (wildness or not), and the between group comparisons identify the dimensions along which the pairs are similar (feline or canine class membership). Evans (Evans, 1968) discusses the need for both kinds of comparisons (relating the components of both W and X and W and Y) in solving some geometric analogies.

Some analogies stated in the same syntactic fashion are more properly interpreted as analogies between two systems, where W and X are related in one system, while Y and Z are related in an analogous system. For example, Johnson-Laird (this volume) discusses, the analogy Water:Sluice:Gas:Jet. Here, there is a between-system mapping of water --> gas and sluice --> jet, but comparison of water and sluice, or gas and jet is not useful. Instead, there are relational systems relating each pair (e.g. a sluice directs water and a jet directs gas). True four-element mappings relate each concept in two different mappings, but in Johnson-Laird's example, a similarity judgment is required between the functional relations in the two systems.

4. ISSUES FOR A THEORY OF MAPPING

There are a number of issues running through the papers in this volume and the literature more generally. In part they reflect the set of subprocesses outlined by Gentner (this volume), but they have wider scope. Our attempt here is simply to delineate the set of issues as best we can, and to discuss possible resolutions to them. We start with the most microscopic issues and work up to the more macroscopic issues.

How are individual properties compared?

Potentially there are two kinds of properties that a theory must take into account: discrete properties (e.g. male or female) and continuous properties (e.g. size). Tversky and Gati (Tversky and Gati, 1982) have shown how it is possible to treat all continuous properties as if they were discrete. Another possibility is to treat all discrete properties as continuous (a person is on a continuum of male/female and most people fall near one or the other ends of the continuum).

Rips (this volume) addresses the question of how continuous properties are compared for different kinds of three-element comparisons: similarity, typicality, and categorization judgments, which he finds are judged differently. His results suggest that categorization judgments are based on the relative height of the distribution - e.g. a three inch object is more likely a pizza than a quarter, because the distribution of pizzas is higher at that point. His results for similarity judgments suggest both height of the distribution and distance from the mean (or mode) come into play. Typicality judgments appear to fall in between categorization and similarity, as if some subjects treat them like categorization judgments and others like similarity judgments (or perhaps they are combination judgments).

There are many possible functions for computing any of these judgments: for example, similarity might be based on the relative distance between modes of the distribution compared; typicality judgments might be simply similarity judgments between a concept and its superconcept, as Smith & Osherson (this volume) assume. Rumelhart's (this volume) theory probably makes a prediction as to which of these functions will best fit the data, but he is not explicit on this point. Most of the other theories take no stand on this issue.

How are judgments from different properties combined?

Tversky (Tversky, 1977) proposes a combining function for similarity judgments, which Smith and Osherson (this volume) have adopted for their theory of decision making. The essence of the Tversky combination rule is that matching properties increase similarity and mismatching properties decrease similarity between concepts. Mismatching properties consist of two sets: property values of one concept that the other does not share, and property values of the other that the first does not share. Mismatching properties include properties where one concept has a known value and the other has no known value. Each of these three sets (one matching and two mismatching properties) is weighted appropriately depending on the direction of the judgment. Thus, people think North Korea is more like China than China is like North Korea, because there are many properties they know about China that do not apply to North Korea, but few properties they know about North Korea that do not apply to China (Tversky, 1977).

The Tversky rule is defined only over similarity judgments and discrete properties. If one adopts the view that all properties are continuous, then a modification of the Tversky rule is necessary. Whether it applies to other kinds of judgments (e.g. categorization judgments) is an open question. And, of course, there are an infinite number of other combination rules, some of which might still be viable given Tversky's (Tversky, 1977) data.

How do people access similar entities?

The question of access is fairly central to the papers of Ross, Gentner, Bransford et al., Barsalou, Brown and Kane, and Holyoak and Thagard (this volume). It is called "noticing" by Ross. All of these papers address the access issue for the case where the source must be found in memory. As Johnson-Laird (this volume) point out, the source is often given, as when a text explains that the earth is like a peach or the atom like a solar system. In Ross's paradigm, when one is working problems, a person may go back through a book to find a similar problem. This access may or may not be governed by the same properties as the access from memory.

Gentner (this volume) proposes that attributes (or superficial properties) govern access more than relational properties. This seems to accord fairly well with both her

data and those of Gick and Holyoak (Gick and Holyoak, 1983) Rumelhart (this volume) takes the position that access is governed by a match on all microfeatures, but in different contexts it may be attributes that match or higher-order relational features. These positions are compatible if one posits that superficial attributes are the most available, and therefore will usually dominate higher-order relations in most matches.

There is some evidence (e.g. (Chi, Feltovich and Glaser, 1981)) that part of becoming an expert is learning to pay attention to higher-order relations rather than superficial attributes. This also accords with Ross's (this volume) observation that superficial properties will mislead people if the principles underlying the problem (i.e. higher-order relations) are confusable. Brown and Kane (this volume) give evidence that functional fixedness and cognitive embeddedness of problem solving contexts are sources of diminished accessibility to potential analogs in children, as well.

How is knowledge about the source reconstructed?

Ross (this volume) points out that people often have to reconstruct their knowledge about the source domain after they have accessed an analogy. This reconstruction process is guided by the knowledge being sought about the target domain. For example, if people are told that heat flow is like water flow (Gentner, this volume) since they do not have a particularly good understanding of water flow (Gentner and Gentner, 1983), they must in part figure out what they know about water flow—that it flows from one container to another as long as there is a difference in the height of the water in the two containers, that the surface area of the water in the container does not matter, that the flow rate is proportional to the diameter of the connection between the containers, etc. Which properties of the source domain people think of depends on what aspects of the target they are trying to understand, as Ross (this volume) has found in his studies.

What governs which properties are transferred?

This is the central argument animating most of the discussion in the analogy literature. We will briefly delineate the different positions.

Ortony (Ortony, 1979) advocates the position that *salience imbalance* governs transfer, that is, those properties are transferred that are important in the source domain but not important in the target domain. For example, Sam is a hippopotamus transfers fatness, since that is a typical property of hippos, but not of people.

Gentner (Gentner, 1983) proposes a syntactic theory that states that, in analogies, relational properties are transferred but attributes (i.e. non-relational properties) are left behind. Furthermore, according to her systematicity principle, relational properties that are a part of a system of relations (e.g. the large mass of the sun attracts the planets into orbiting around it) are more likely to be mapped across.

Holyoak and Thagard (this volume), Johnson-Laird (this volume), Carbonell (Carbonell, 1986), and Burstein (Burstein, 1986), while there are differences in their views, take a position on mapping that appears somewhat different from Gentner. Their position is that a system (or schema) of properties is mapped over, as Gentner proposes, but with two differences. (1) attributes will be mapped if they are part of the system, and (2) the major problem is to decide which system to map over. For example, if the analogy was made between the solar system and a person tanning themselves under a sun lamp, the properties mapped would have to do with the heat being transmitted, the person rotating to cover all sides, the yellow color of the lamp, etc.

It turns out that the latter criticism may be handled by the structure mapping engine (Falkenhainer et al., 1986, Gentner, this volume) that was built recently to embody the Gentner theory. This system compares representations of two domains to decide which relations fit into a connected system that can be mapped into the target domain. Because it is effectively comparing all possible sets of relations between the objects considered, it is to some degree automatically choosing a "best system" to map. However, some pragmatic, contextual selection mechanisms will almost certainly be required as well. This is particularly true during learning, when people usually do not know enough about the target domain to pick out corresponding systems simply by matching (Burstein, 1986).

An important test of any of these computer models (Burstein, 1986, Carbonell, 1986, Gentner, this volume, Holyoak and Thagard, this volume) is whether they can select two different mappings from a source domain (e.g. the solar system) depending on what aspects of the source domain are relevant to the target domain (e.g. the atom vs. a person tanning). None of the models has, as yet addressed this central problem directly.

Whether goals and subgoals guide the selection of the system to be mapped often arises in the debate between these two positions. But that is probably because the latter researchers are all working with analogies in problem solving, whereas Gentner is dealing mainly with explanatory analogies. Certainly both sides would agree that goals are critical properties to map in problem-solving analogies and play the same central role that causal relations play in explanatory analogies.

Anderson and Thompson (this volume) rely on a set of three principles (i.e., "no function is content", "sufficiency of functional specification", and "maximal functional elaboration") to determine what is mapped. Although it is not clear to us exactly how these principles operate, they indicate the use of function as the main criteria for selecting what to map, and so would seem to fall into the latter camp.

In our view the positions of Gentner on the one hand and that of Holyoak and Thagard, Johnson-Laird, Carbonell, and Burstein on the other hand are not that far apart given the centrality of systems of properties or schemas that are mapped over. The Ortony theory is orthogonal to that issue, and could operate in conjunction with some kind of system mapping. Whether the Anderson and Thompson position is genuinely distinct, or reduces to the use of system properties as well, remains to be seen.

How are multiple mappings merged together?

This issue is raised by Burstein (Burstein, 1985, 1986, 1987), Spiro (this volume) and Collins and Gentner (Collins and Gentner, 1983). In Burstein's work, students were learning to program and were forced to combine the mappings of systems like putting things in boxes and the interpretation of arithmetic equalities in forming a mental model to understand computer statements like $A=B+I$. Collins and Gentner (Collins and Gentner, 1983, 1987) describe how subjects combined different analogies (e.g. billiard-ball analogy, a rocketship analogy, a crowded-room analogy) in understanding evaporation processes. It is clear that people frequently construct their understandings of systems by multiple mappings, and so theories will have to specify how conflicts are resolved about what properties to map from each analogy, and whether, in fact, some form of conceptual combination is required to merge related properties mapped from several different sources. In Burstein's model conflicts between mappings are usually resolved by reasoning from specific examples in

the target domain that cause one or another analogical mapping to fail. However, the hypotheses that are eventually selected must still be integrated with what had been mapped previously or was otherwise known about the target domain (Burstein, 1987).

Burstein (Burstein, 1985) and Collins and Gentner (Collins and Gentner, 1983) also raise the issue of *vertical integration* of mental models. Analogies do not always map onto the same level of description of a target system. In such cases, one cannot directly merge analogs. Instead, the mapped structures must be maintained distinctly, and rules of correspondence formed between the different views or levels of abstraction described by the different analogical models.

How are mappings refined?

After a mapping is made, some properties carried over into the target domain will not apply. How are the correct properties identified and replaced? Both Burstein (Burstein, 1986) and Anderson and Thompson (this volume) address this question in the context of mapping computer program statements. In Burstein's model, analogically mapped predictions are compared to the actual results in target domain examples. If the predictions are wrong, alternative structures are considered for mapping, either from the same or a different source domain. Anderson and Thompson discuss several examples of failures due to overgeneralization from an analogy, and suggest that they may be handled by searching for contextual features that were *not* mapped, and adding them as preconditions.

Another kind of refinement occurs when successful analogies are extended to encompass new sets of corresponding systems or related causal principles. In addition to mapping new relational properties, this kind of analogical extension can lead to the introduction of new object or concept correspondences. For example, in the kinds of demonstration physics experiments that are often used to explain the diffraction and interference behavior of sound and light by using water wave tanks, a number of experimental objects are introduced to cause different wave behaviors. Each object that is introduced in these experiments must be related to an analogous object that causes a similar kind of interference with light or sound. In this sense, each new experiment described causes the refinement of the analogy between water waves and light or sound waves, because new objects and new causal implications are placed in parallel.

What is generalized from a mapping?

This is the question of how, when, and if generalizations are made based on a mapping between two domains. For example, one hypothesis might be that the corresponding components in the two systems are replaced by their common supersets and the generalization is stored as a set of (possibly generalized) relations on these common supersets. Both Gentner (this volume), Anderson and Thompson (this volume) and Winston (Winston, 1982) have addressed this issue to some degree, although no specific claims have been made.

It is not at all clear that analogies always lead to new generalizations. Most analogies are only useful because they map one or two specific pieces of information from one domain to another. In such cases, the generation of a new general principle may not be warranted.

At the other extreme, attempting to generalize from an analogy that related radically different classes of objects by a new principle calls for a strong form of conceptual reclassification, as when sound and light are reclassified as waves. Very strong evidence of the analogy's pervasiveness may be needed for this kind of reclassification to occur. Alternatively, "bridging analogies" can be used to show why the analogy is justified. Clement (Clement, 1981, 1986) gives examples of series of bridging analogies designed to convince people of the generality of physical laws. One set of these analogies shows how the behavior of a spring is related to the longitudinal and torsional flex of a wire, by considering intermediate cases where the wire is partially bent. Clement (Clement, 1986) also discusses Newton's analogy between the moon and an apple falling from a tree, with the a sequence of bridging analogs where a cannonball is fired at greater and greater speeds until it is in orbit around the earth.

How does the process of mapping develop?

This is the central issue raised by Linda Smith's paper (this volume). In it, she proposes that development proceeds from overall resemblance matches to identity matches and finally to dimensional matches. Her proposal perhaps is best summed up by saying that children learn to make finer discriminations in their comparison processing with age.

Her thesis raises the question of how children can make overall resemblance comparisons without being able to make individual property comparisons. This is not really a paradox from the vantage point of the kind of microfeature theory proposed by Rumelhart (this volume). Overall resemblance comparison in Rumelhart's theory can be carried out by comparing two concepts with respect to all their microfeatures. This requires no identification of microfeatures with particular properties (like color) of entities in the world. Based on the kind of perceptual learning described by Bransford and his colleagues (this volume), dimensions or subgroups of the microfeatures will emerge as contrastive sets of microfeatures that inhibit each other. Making an identity match would seem to require learning how much variability is possible on any dimension so that one can assess whether the difference between two entities falls below the normal range of variability on that dimension. In any case, the papers of Smith, Rumelhart, and Bransford et al together promote a consistent picture of how similarity matching develops.

Are analogies helpful for learning?

This issue was raised by Halasz and Moran (Halasz and Moran, 1982). Their position is that if you give people explanatory analogies, such as the analogy that computer addresses are like boxes (Burstein, 1986) or that heat flow is like liquid flow (Gentner, this volume), you lead them to make more wrong mappings than helpful ones. So they argue that it is better to give people descriptions of the mechanisms involved, rather than analogies.

There are at least two arguments against the Halasz and Moran (Halasz and Moran, 1982) position. First, when people learn about novel systems, they are going to impute mechanisms to them. In order to understand any mechanistic description, they have to draw from their stock of basic mechanisms, such as Carbonell (this volume) or Collins and Gentner (Collins and Gentner, 1983) have described. So, whether you give students an analogy or not, they are going to make an analogy to some mechanism they already understand. The continuum from remembering, to reminding, to analogy that Rumelhart (this volume) describes is operating here. Subjects will pull in the mechanism they know about that matches most closely. By giving students an explicit analogy, you then accomplish two things: (a) you make sure they impute the best matching mechanism, and (b) you know what wrong inferences they are likely to draw so that you can try to counter them as you explain the mechanism.

A second argument against the Halasz and Moran (Halasz and Moran, 1982) position is that the power of analogies for teaching derives from the fact that they provide a well-integrated structure that can be assimilated all at once. This structure may have acquired over a long period of time, as Vosniadou (this volume) shows for the solar system. So by telling someone the atom is like a solar system, they have a well-integrated structure acquired over many years that they can map as a whole in order to understand the atom. Thus they do not have to recapitulate the same long learning process for the atom. Analogies are particularly powerful where there is a competing structure already in place that the teacher is trying to dislodge.

The Halasz and Moran (Halasz and Moran, 1982) position, however, has to be correct if the analogy introduces too many wrong mappings. Therefore, we would argue that the issue is not whether analogies are helpful or harmful, but what determines when they are helpful vs. when they are harmful for learning.

5. CONCLUSION

Most researchers are working in a little corner of this framework, which is fine. One use of the framework is to help them see what the rest of the territory looks like in order to help them extend their theory to cover the whole territory. By trying to extend their theory in this way, it puts additional constraints on theory construction, which will help researchers refine their theories. Furthermore, as theories are extended to cover the whole domain, they will bump up against other theories in more ways which will lead to fruitful controversies and issues to be settled empirically. Psychology and artificial intelligence have a tendency to construct task-based theories and need to enforce on their theorists the desirability of constructing more global theories.

Table 1

Contexts in which mappings occur

I. Type of Task

1. Comparative judgments

- a. Similarity judgments
- b. Typicality Judgements
- c. Categorization judgments
- d. Identity judgments
- e. Overlap judgments
- f. Difference judgments

2. Mappings

- a. Property mappings
- b. Component mappings

3. Conceptual combinations

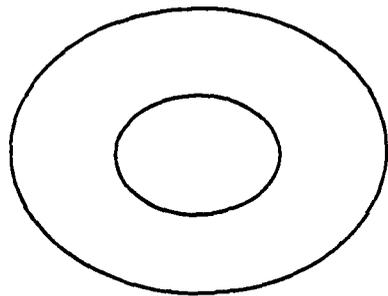
II. Number of Entities Compared

- 1. Two-element mappings
- 2. Three-element mappings
- 3. Four-element mappings

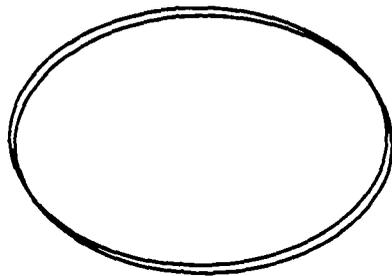
References

- Burstein, Mark H. *Learning by Reasoning from Multiple Analogies*. Doctoral dissertation, Yale University, 1985.
- Burstein, Mark H. Concept Formation by Incremental Analogical Reasoning and Debugging. In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Ed.), *Machine Learning, Volume II*. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1986. Also appeared in the *Proceedings of the Second International Machine Learning Workshop*, Champaign-Urbana, IL, 1983.
- Burstein, Mark H. Incremental Learning from Multiple Analogies. In *Proceedings of Analogica-85*. Boston, MA: Pitman, 1987. Forthcoming.
- Carbonell, Jaime G. Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In Michalski, R. S., Carbonell, J. G. and Mitchell, T. M. (Ed.), *Machine Learning, Volume II*. Los Altos, CA: Morgan Kaufmann Publishers, Inc., 1986.
- Chi, M., Feltovich, P., and Glaser, R. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981, 5(2), 121-152.
- Clement, J. Analogy generation in scientific problem solving. In *Proceedings of the Third Annual Conference of the Cognitive Science Society*. Berkeley, CA: University of California, 1981.
- Clement, J. Methods for evaluating the validity of hypothesized analogies. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Amherst, MA: University of Massachusetts, 1986.
- Collins, Allan. Fragments of a Theory of Human Plausible Reasoning. In D. L. Waltz (Ed.), *Theoretical Issues in Natural Language Processing*. Urbana-Champaign, IL: University of Illinois, 1978.
- Collins, Allan and Gentner, Dedre. A Framework for a Cognitive Theory of Writing. In L. W. Gregg and E. Steinberg (Eds.), *Cognitive processes in writing: An interdisciplinary approach*. Hillsdale, NJ: Erlbaum, 1980.
- Collins, Allan and Gentner, Dedre. Multiple Models of Evaporation Processes. In *Proceedings of the Fifth Annual Conference of the Cognitive Science Society*. Rochester, NY: Cognitive Science Society, 1983.
- Collins, Allan and Gentner, Dedre. How People Construct Mental Models. In N. Quinn and D. Holland (Eds.), *Cultural Models in Thought and Language*. Cambridge, UK: Cambridge University Press, 1987. In press.
- Collins, A. and Michalski, R. The Logic of Plausible Reasoning: A Core Theory. Submitted to *Cognitive Science*.
- Collins, A., Gentner, D. and Rubin, A. *Teaching Study Strategies* (Tech. Rep. Report No. 4794). Bolt Beranek and Newman Inc., 1981.
- Duncker, K. On problem solving. *Psychological Monographs*, 1945, Vol. 58(270).

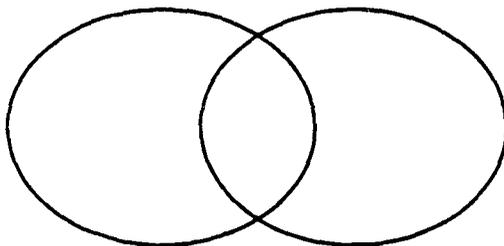
- Evans, Thomas G. A Program for the Solution of Geometric Analogy Intelligence Test Questions. In Marvin L. Minsky (Ed.), *Semantic Information Processing*. Cambridge, Massachusetts: MIT Press, 1968.
- Falkenhainer, B., Forbus, K., and Gentner, D. The Structure-Mapping Engine. In *Proceedings of AAAI-86*. Los Altos, CA: Morgan Kaufman, 1986.
- Gentner, Dedre. Structure-Mapping: A theoretical framework for analogy. *Cognitive Science*, 1983, 7(2), 155-170.
- Gentner, D. and Gentner, D. R. Flowing waters or teeming crowds: Mental models of electricity. In Gentner, D. and Stevens, A. L. (Eds.), *Mental Models*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1983.
- Gentner, D. and Landers, R. Analogical reminding: A good match is hard to find. In *Proceedings of the International Conference on Systems, Man and Cybernetics*. Tucson, AZ: University of Arizona, 1985.
- Gick, L. Analogical problem solving. *Cognitive Psychology*, 1980(12), pp. 306-355.
- Gick, L. Schema induction and analogical transfer. *Cognitive Psychology*, 1983(15), pp. 1-38.
- Halasz, Frank and Moran, Thomas P. Analogy Considered Harmful. In *Proceedings of the Human Factors in Computer Systems Conference*. Gaithersburg, MD, 1982.
- Malt, Barbara C. and Smith, Edward E. Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior*, 1984, 23, 250-269.
- Meyer, David E. On the representation and retrieval of stored semantic information. *Cognitive Psychology*, 1970, 1, 242-300.
- Ortony, Andrew. Beyond literal similarity. *Psychological Review*, 1979, 87, 161-180.
- Rosch, E. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 1975, 104, 192-233.
- Tversky, A. Features of similarity. *Psychological Review*, 1977, 84, 327-352.
- Tversky, A. and Gati, I. Similarity, separability, and the triangle inequality. *Psychological Review*, 1982, 89, 123-154.
- Tversky, A. and Kahneman, D. Causal schemas in judgments under uncertainty. In M. Fishbein (Eds.), *Progress in social psychology*. Hillsdale, NJ: Erlbaum, 1980.
- VanLehn, K. and Brown, J. S. Planning Nets: A representation for formalizing analogies and semantic models of procedural skills. In Snow, R. E., Federico, P. and Montague, W. E. (Eds.), *Aptitude, Learning and Instruction, Volume 2*. Hillsdale, NJ: Erlbaum, 1980.
- Winston, P. H. Learning new principles from precedents and exercises. *Artificial Intelligence*, 1982, 19, 321-350.



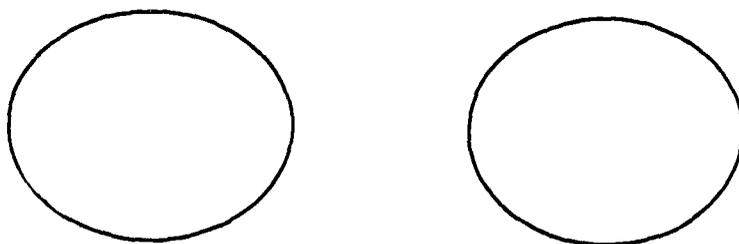
Categorization



Identity

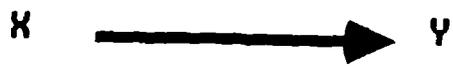


Overlap

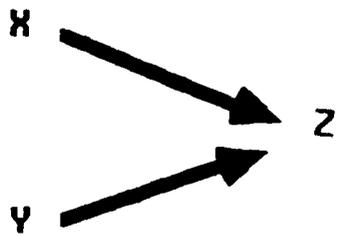


Difference

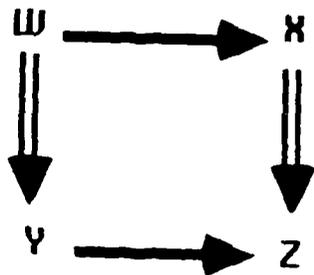
Figure 1 Four relations in comparing concepts.



Two-element comparison



Three-element comparison



Four-element comparison

Figure 2 The structure of different comparisons of concepts.