Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou and A. Ortony, (Eds.), Similarity and analogical reasoning, (pp. 199-241). London: Cambridge University Press.

7

The mechanisms of analogical learning

DEDRE GENTNER

It is widely accepted that similarity is a key determinant of transfer. In this chapter I suggest that both of these venerable terms – *similarity* and *transfer* – refer to complex notions that require further differentiation. I approach the problem by a double decomposition: decomposing similarity into finer subclasses and decomposing learning by similarity and analogy into a set of component subprocesses.

One thing reminds us of another. Mental experience is full of moments in which a current situation reminds us of some prior experience stored in memory. Sometimes such remindings lead to a change in the way we think about one or both of the situations. Here is an example reported by Dan Slobin (personal communication, April 1986). His daughter, Heida, had traveled quite a bit by the age of 3. One day in Turkey she heard a dog barking and remarked, "Dogs in Turkey make the same sound as dogs in America.... Maybe all dogs do. Do dogs in India sound the same?" Where did this question come from? According to Slobin's notebook, "She apparently noticed that while the people sounded different from country to country, the dogs did not." The fact that only humans speak different languages may seem obvious to an adult, but for Heida to arrive at it by observation must have required a series of insights. She had to compare people from different countries and note that they typically sound different. She also had to compare dogs from different countries and note that they sound the same. Finally, in order to attach significance to her observation about dogs, she must have drawn a parallel – perhaps implicitly – between dogs making sounds and humans making sounds so that she could contrast: "As you go from country to country, people sound different, but dogs sound the same." Thus her own experiential comparisons led her to the beginnings of a major insight about the difference between human language and animal sounds.

This example illustrates some of the power of spontaneous remindings. Spontaneous remindings can lead us to make new infer-

ences, to discover a common abstraction, or, as here, to notice an important difference between two partly similar situations (e.g., Ross, 1984, this volume). The ultimate aim of this chapter is to trace learning by analogy and similarity from the initial reminding to the final storage of some new information. Spontaneous analogical learning[1] can be decomposed into subprocesses of (a) accessing the base* system; (b) performing the mapping between base and target; (c) evaluating the match; (d) storing inferences in the target; and sometimes, (e) extracting the commonalities (Clement, 1981, 1983; Gentner, 1987; Gentner & Landers, 1985; Hall, in press; Kedar-Cabelli, 1988).

This breakdown suggests that we examine the subprocesses independently. Once this is done, it will become clear that different subprocesses involved in analogical learning are affected by very different psychological factors. Although the chronological first step in an experiential learning sequence is *accessing the potential analog*, I shall postpone the discussion of access until later in this chapter. Instead, I begin with steps 2 and 3: *analogical mapping* and *judging analogical soundness*. This is the logical place to start, because it is these processes that uniquely define analogy and allow us to see distinctions among different kinds of similarity. It turns out that the theoretical distinctions necessary for talking about analogical mapping are also useful for talking about other analogical subprocesses.

The plan of the chapter is, first, to describe the core structure-mapping theory of analogical mapping, using a computer simulation to make the points clear; second, to offer psychological evidence for the core theory of analogical mapping; and, finally, to discuss research that extends the framework to the larger situation of analogical learning.

Analogical mapping

The theoretical framework for this chapter is the structure-mapping theory of analogy (Gentner, 1980, 1982, 1983, 1987; Gentner & Gentner, 1983).[2] As Stephen Palmer (this volume) states, structure-mapping is concerned, first, with what Marr (1982) called the "computational level" and what Palmer and Kimchi (1985) call the "informational constraints" that define analogy. That is, structure-mapping aims to capture the essential elements that constitute analogy and the operations that are computationally necessary in processing

* *Editors' note*: The terms "base" and "source" are used interchangeably both in the field in general and in this volume in particular.

analogy. The question of how analogies are processed in real time – that is, the question of which algorithms are used, in Marr's terminology, or which behavioral constraints apply, in Palmer and Kimchi's terminology – will be deferred until later in the chapter.

The central idea in structure-mapping is that an analogy is a mapping of knowledge from one domain (the base) into another (the target), which conveys that a system of relations that holds among the base objects also holds among the target objects. Thus an analogy is a way of focusing on relational commonalties independently of the objects in which those relations are embedded. In interpreting an analogy, people seek to put the objects of the base in one-to-one correspondence with the objects in the target so as to obtain the maximum structural match. Objects are placed in correspondence by virtue of their like roles in the common relational structure; there does not need to be any resemblance between the target objects and their corresponding base objects. Central to the mapping process is the principle of systematicity: People prefer to map connected *systems of relations* governed by higher-order relations with inferential import, rather than isolated predicates.

Analogical mapping is in general a combination of matching existing predicate structures and importing new predicates (carry-over). To see this, first consider the two extremes. In *pure matching*, the learner already knows something about both domains. The analogy conveys that a relational system in the target domain matches one in the base domain. In this case the analogy serves to focus attention on the matching system rather than to convey new knowledge. In *pure carry-over*, the learner initially knows something about the base domain but little or nothing about the target domain. The analogy specifies the object correspondences, and the learner simply carries across a known system of predicates from the base to the target. This is the case of maximal new knowledge. Whether a given analogy is chiefly matching or mapping depends, of course, on the state of knowledge in the learner. For example, consider this analogy by Oliver Wendell Holmes, Jr.: "Many ideas grow better when transplanted into another mind than in the one where they sprang up." For some readers, this might be an instance of pure mapping: By importing the knowledge structure from the domain of plant growing to the domain of idea development they receive a completely new thought about the latter domain. But for readers who have entertained similar thoughts the process is more one of matching. The effect of the analogy is then not so much to import new knowledge as to focus attention on certain portions of the existing knowledge. Most explanatory analogies are a

combination of matching and carry-over. Typically, there is a partial match between base and target systems, which then sanctions the importing of further predicates from the base to the target.

A clarification may be useful here. A possible misreading is that the systematicity principle implies that the same set of predicates should always be mapped from a given base domain, regardless of the target (Holyoak, 1985). But, by this construal, the interpretation of an analogy would depend only on the base domain, which is patently false, except in the case when nothing is known about the target (the pure carry-over case). In the normal case, when there is information about both base and target, a given base–target pair produces a set of matching predicates. Changing either member of the pair produces a different set of matching predicates. Thus, systematicity operates as a selection constraint: Among the many possible predicate matches between a given base and target, it favors those that form coherent systems of mutually interconnecting relations (see Clement & Gentner, 1988; Gentner & Clement, in press).

To illustrate the structure-mapping rules, we turn to a specific example: the analogy between heat flow and water flow. (See Gentner & Jeziorski, in press, for a discussion of Carnot's use of this analogy in the history of heat and temperature.) Figure 7.1 shows a water-flow situation and an analogous heat-flow situation.

I will go through this analogy twice. The first time I give the analogy as it might occur in an educational setting in which the learner knows a fair amount about water and almost nothing about heat flow. Here the learner is given the object correspondences between water and heat and simply imports predicates from the water domain to the heat domain. This is a case of pure carry-over. The second time, to illustrate the computer simulation, I give the analogy as it might occur with the learner having a good representation of the water domain and a partial representation of the heat domain. Here the analogy process is a combination of matching existing structures and importing new predicates (carry-over).

The heat/water analogy, Pass 1: pure carry-over. Figure 7.2 shows the representation a learner might have of the water situation. We assume that the learner has a very weak initial representation of the heat situation and perhaps even lacks a firm understanding of the difference between heat and temperature. This network represents a portion of what a person might know about the water situation illustrated in Figure 7.1.[3]

The learner is told that heat flow can be understood just like water

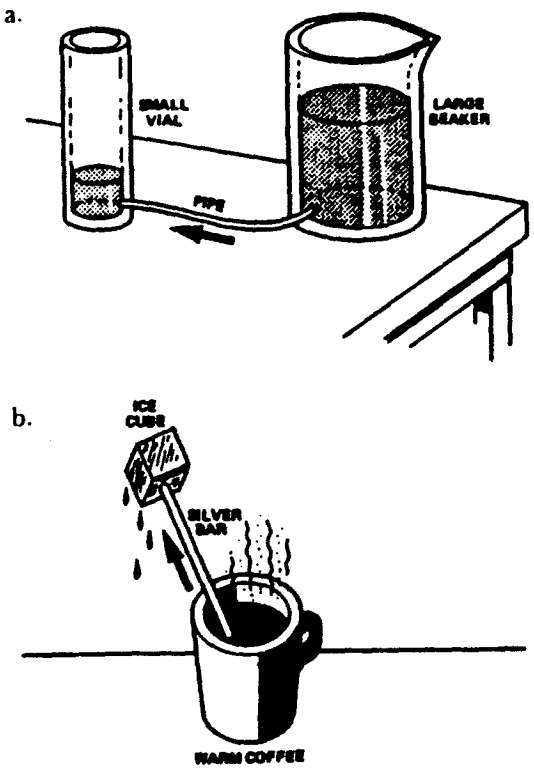

Figure 7.1. Examples of physical situations involving (*a*) water flow and (*b*) heat flow (adapted from Buckley, 1979, pp. 15–25).

flow, with temperature in the heat situation playing the role of pressure in the water situation. The learner is also given the object correspondences

water → heat; pipe → metal bar;
beaker → coffee; vial → ice

as well as the function correspondence

PRESSURE → TEMPERATURE

Now the learner is in a position to interpret the analogy. Even with the correspondences given, there is still some active processing required. In order to comprehend the analogy, the learner must

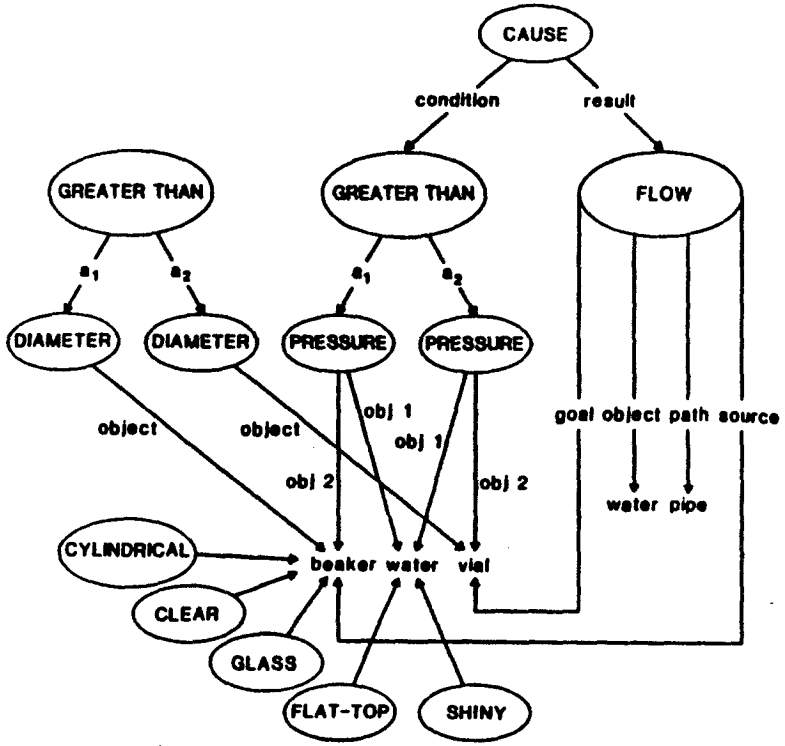

Figure 7.2. A representation of the water situation. Predicates are written in upper case and circled; objects are written in lower case and uncircled.

- ignore object attributes; e.g., CYLINDRICAL (beaker) or LIQUID (coffee)
- find a set of systematic base relations that can apply in the target, using the correspondences given. Here, the pressure-difference structure in the water domain

  CAUSE {GREATER [PRESSURE (beaker), PRESSURE (vial)], [FLOW (water, pipe, beaker, vial)]}

  is mapped into the temperature-difference structure in the heat domain

  CAUSE {GREATER [TEMP (coffee), TEMP (ice)], [FLOW (heat, bar, coffee, ice)]}

- and discard isolated relations, such as

  GREATER [DIAM (beaker), DIAM (vial)]

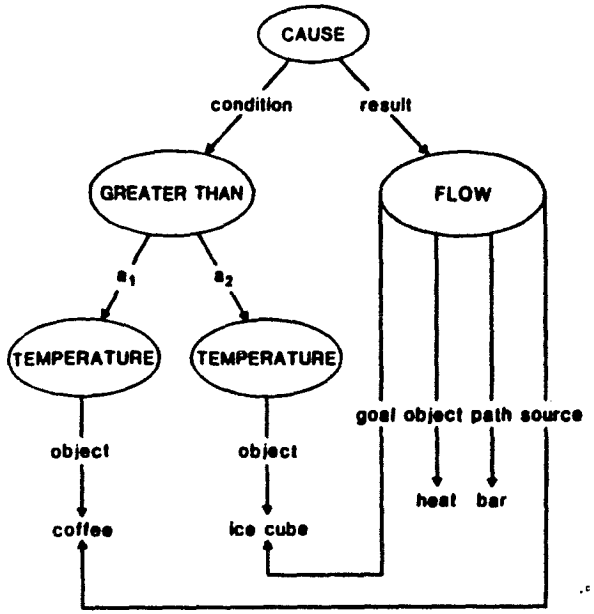

Figure 7.3. A representation of the heat situation that results from the heat/water analogy.

Figure 7.3 shows the resulting causal representation of heat flow induced by the analogical mapping.

There are several points to note in this example. First, the object correspondences – heat/water, beaker/coffee, vial/ice, and pipe/bar – and the function correspondence PRESSURE/TEMPERATURE[4] are determined not by any intrinsic similarity between the objects but by their role in the systematic relational structure. Systematicity also determines which relations get carried across. The reason that

GREATER [PRESSURE (beaker), PRESSURE (vial)]

is preserved is that it is part of a mappable system of higher-order constraining relations: in this case, the subsystem governed by the higher-order relation CAUSE. In contrast, the relation

GREATER [DIAM (beaker), DIAM (vial)]

does not belong to any such mappable system and so is less favored in the match.

Second, the order of processing is probably variable. Even when the learner is given the object correspondences first, there is no ob-

Table 7.1. *Kinds of domain comparisons*

	Attributes	Relations	Example
Literal similarity	Many	Many	Milk is like water
Analogy	Few	Many	Heat is like water
Abstraction	Few	Many	Heat flow is a through-variable
Anomaly	Few	Few	Coffee is like the solar system
Mere appearance	Many	Few	The glass tabletop gleamed like water

vious constraint on the order in which predicates should be mapped. This is even more the case when the learner is not told the object correspondences in advance. In this case, as exemplified in the next pass through this analogy, the object correspondences are arrived at by first determining the best predicate match – that is, the most systematic and consistent match. I suspect that the order in which matches are made and correspondences tried is extremely opportunistic and variable. It seems unlikely that a fixed order of processing stages will be found for the mapping of complex analogies (see Grudin, 1980; Sternberg, 1977).

Third, applying the structural rules is only part of the story. Given a potential interpretation, the candidate inferences must be checked for validity in the target. If the predicates of the base system are not valid in the target, then another system must be selected. In goal-driven contexts, the candidate inferences must also be checked for relevance to the goal.

Kinds of similarity

Distinguishing different kinds of similarity is essential to understanding learning by analogy and similarity. Therefore, before going through the heat/water analogy a second time, I lay out a decomposition of similarity that follows from what has been said. Besides analogy, other kinds of similarity can be characterized by whether the two situations are alike in their relational structure, in their object descriptions, or in both. In *analogy*, only relational predicates are mapped. In *literal similarity*, both relational predicates and object attributes are mapped. In *mere-appearance matches*, it is chiefly object attributes that are mapped. Figure 7.4 shows a similarity space that summarizes these distinctions. Table 7.1 shows examples of these

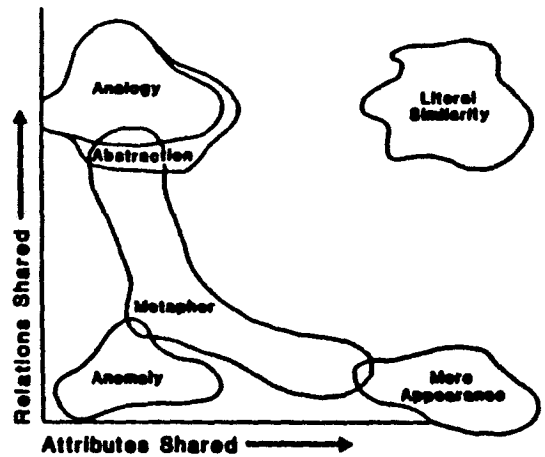

Figure 7.4. Similarity space: classes of similarity based on the kinds of predicates shared.

different kinds of similarity. The central assumption is that it is not merely the relative *numbers* of shared versus nonshared predicates that matter – although that is certainly important, as Tversky (1977) has shown – but also the *kinds* of predicates that match.

Analogy is exemplified by the water/heat example discussed above, which conveys that a common relational system holds for the two domains: Pressure difference causes water flow, and temperature difference causes heat flow. *Literal similarity* is exemplified by the comparison "Milk is like water," which conveys that much of the water description can be applied to milk. In literal similarity, both object attributes, such as FLAT TOP (water) and CYLINDRIC \L (beaker), and relational predicates, such as the systematic causal structure discussed above, are mapped over. A *mere-appearance match* is one with overlap in lower-order predicates – chiefly object attributes[5] – but not in higher-order relational structure as in "The glass tabletop gleamed like water." Mere-appearance matches are in a sense the opposite of analogies. Such matches are sharply limited in their utility. Here, for example, little beyond physical appearance is shared between the tabletop and water. These matches, however, cannot be ignored in a theory of learning, because they often occur among novice learners. One further type of match worth discussing is *relational abstraction*.

An example is the abstract statement, "Heat is a through-variable," which might be available to a student who knew some system dynamics. This abstraction, when applied to the heat domain, conveys much the same relational structure as is conveyed by the analogy: that heat (a through-variable) can be thought of as a flow across a potential difference in temperature (an across-variable). The difference is that the base domain contains only abstract principles of through-variables and across-variables and variables; there are no concrete properties of objects to be left behind in the mapping.

These contrasts are continua, not dichotomies. Analogy and literal similarity lie on a continuum of degree-of-attribute-overlap. In both cases, the base and target share common relational structure. If that is *all* they share, then the comparison is an analogy (assuming, of course, that the domains are concrete enough to have object descriptions). To the extent that the domains also share common object descriptions, the comparison becomes one of literal similarity. Another continuum exists between analogies and relational abstractions. In both cases, a relational structure is mapped from base to target. If the base representation includes concrete objects whose individual attributes must be left behind in the mapping, the comparison is an analogy. As the object nodes of the base domain become more abstract and variable-like, the comparison becomes a relational abstraction.

We turn now to the second pass through the analogy. There are two innovations. First, in this pass I describe the way our computer simulation processes the heat/water example. Here we move from informational constraints to behavioral constraints. (See Palmer, this volume.) Second, in this pass I assume that there is some prior knowledge of *both* base and target; thus this pass illustrates a combination of matching and carry-over. Before giving the algorithm, I describe the representational conventions.

Representation conventions. The order of an item in a representation is as follows: Objects and constants are order 0. The order of a predicate is 1 plus the maximum of the order of its arguments. Thus, if x and y are objects, then GREATER THAN (x, y) is first-order. and CAUSE [GREATER THAN (x, y), BREAK(x)] is second-order. Typical higher-order relations include CAUSE and IMPLIES. On this definition, the order of an item indicates the depth of structure below it. Arguments with many layers of justifications will give rise to representation structures of high order.

A typed predicate calculus is used in the representation. There are four representational constructs that must be distinguished: *entities*, which represent individuals and constants, and three types of predicates. Predicates are further subdivided into truth-functional predicates (*relations* and *attributes*) and *functions*. *Entities* (e.g., *Eddie, side pocket*) are logical individuals: the objects and constants of a domain. Typical entities include pieces of stuff, individual objects or beings, and logical constants. *Attributes* and *relations* are predicates that range over truth values; for example, the relation HIT(cue ball, ball) can be evaluated as true or false. The difference is that attributes take one argument and relations take two or more arguments. Informally, attributes describe properties of entities, such as RED or SQUARE. Relations describe events, comparisons, or states applying to two or more entities or predicates. First-order relations take objects as arguments: for example, HIT(ball, table) and INSIDE (ball, pocket). Higher-order relations such as IMPLIES and CAUSE take other predicates as their arguments: for example, CAUSE [HIT (cue stick, ball), ENTER (ball, pocket)]. *Functions* map one or more entities into another entity or constant. For example, SPEED(ball) does not have a truth value; instead, it maps the physical object *ball* into the quantity that describes its speed. Functions are a useful representational device because they allow either (a) evaluating the function to produce an object descriptor, as in HEIGHT (Sam) = 6 feet, or (b) using the unevaluated function as the argument of other predicates, as in GREATER THAN [HEIGHT(Sam), HEIGHT(George)].

These four constructs are all treated differently in the analogical mapping algorithm. Relations, including higher-order relations, must match identically. Entities and functions are placed in correspondence with other entities and functions on the basis of the surrounding relational structures. Attributes are ignored. Thus there are three levels of preservation: identical matching, placing in correspondence, and ignoring.[6] For example, in the analogy "The wrestler bounced off the ropes like a billiard ball off the wall," the *relations*, such as CAUSE [HIT(wrestler1, wrestler2), COLLIDE(wrestler2, ropes)] must match identically. For *objects* and for *functions*,[7] we attempt to find corresponding objects and functions, which need not be identical: for example, cue ball/wrestler and SPEED(cue ball)/FORCE(wrestler1). Attributes are ignored; we do not seek identical or even corresponding attributes in the billiard ball for each of the wrestler's attributes. To sum up, relations must match, objects and functions must correspond, and attributes are ignored.

It is important to note that these representations, including the distinctions between different kinds of predicates, are intended to reflect the way situations are construed by people. Logically, an n-

place relation $R(a,b,c,)$ can always be represented as a one-place predicate $Q(x)$, where $Q(x)$ is true just in case $R(a,b,c,)$ is true. Further, a combination of a function and a constant is logically equivalent to an attribute; for example, applying the function EQUALS [COLOR(ball), red] is logically equivalent to stating the attribute RED(ball). Our aim is to choose the representation that best matches the available evidence as to the person's current psychological representation. As Palmer (this volume) points out, these representational decisions are crucial to the operation of the algorithm. Differences in the way things are construed can cause two situations to fail to match even if they are informationally equivalent. Thus the model would fail to realize that HOTTER THAN(a,b) is equivalent to COLDER THAN(b,a). This assumption may not be as implausible as it initially seems. Empirically, we know that logical equivalence does not guarantee psychological equivalence. Perhaps one reason that people sometimes miss potential analogies (as discussed below) is that their current representations of base and target limit the kinds of analogical matches they can make.

Requiring perfect relational identity in the matching rules allows us to capture the fact that potential analogies are often missed, for the more exactly the representations must match, the less likely analogies are to be seen. More important, the relational identity requirement keeps us from concealing homunculus-like insight in the matcher. As soon as we move away from perfect matching we are faced with a host of difficult decisions: How much insight do we give the matcher? How much ability to consider current contextual factors? How much tolerance for ambiguity? In short, we lose the considerable advantages of having a simple, low-cost matcher. But how can we capture the intuition that people sometimes can use analogy creatively to surmount initially different representations? Burstein (1983) has explored one interesting method: He allows *similar* predicates to match and then generalizes the match. For example, as part of a larger analogy, *inside* in the spatial sense is matched with *inside* in the abstract sense of a variable containing a value. Then a more general notion of containment is abstracted from the match. This is an attractive notion that deserves further study. However, it does run the risk of adding considerable computational ambiguity.

One way to add flexibility without sacrificing the simple matcher is to add some tools for re-representation that are external to the matcher itself. Then, if there were good reason to suspect a possible analogy, a relation currently represented as COLDER THAN(b,a) could be re-represented as HOTTER THAN(a,b). An alternative re-representation would decompose it into GREATER THAN [TEMP(a),TEMP(b)].

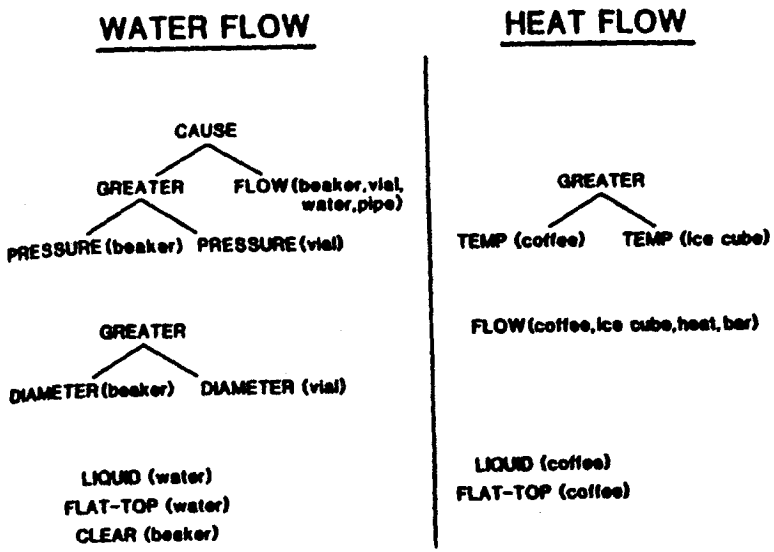

Figure 7.5. Representations of water and heat given to the structure-mapping engine.

In this way a partial analogy could lead to the discovery that two relations hitherto seen as different in fact refer to the same underlying dimension. This would allow us to model the use of analogy in reconstruing one situation in terms of another. An interesting corollary of this approach is that it suggests a way in which analogy could act as a force toward building uniform domain representations, both within and across domains.

The structure-mapping engine. The structure-mapping engine (SME) is a simulation of the structure-mapping process written by Brian Falkenhainer and Ken Forbus (Falkenhainer, Forbus, & Gentner, 1986, in press; Gentner, Falkenhainer, & Skorstad, 1987). Here it is given the representations of the base and target shown in Figure 7.5. As in the previous pass (Figure 7.2), we assume the learner has a fair amount of knowledge about water. In contrast to the previous pass, we now assume some initial knowledge about heat: The learner knows that the coffee is hotter than the ice and that heat will flow from the coffee to the ice. Note, however, that the representations contain many extraneous predicates, such as LIQUID(water) and

LIQUID(coffee). These are included to simulate a learner's uncertainty about what matters and to give SME the opportunity to make erroneous matches, just as a person might.

In addition to modeling analogy, SME can be used with literal similarity rules or mere-appearance rules. Both analogy rules and literal similarity rules seek matches in relational structure; the difference is that literal similarity rules also seek object-attribute matches. Mere-appearance rules seek only object-attribute matches. I will describe the process using literal similarity rules, rather than pure analogy, because this offers a better demonstration of the full operation of the simulation, including the way conflicts between surface and structural matches are treated.

The heat/water analogy, Pass 2: matching plus carry-over. Given the comparison "Heat is like water," SME uses systematicity of relational structure and consistency of hypothesized object correspondences to determine the mapping. The order of events is as follows:

1. Local matches. SME starts by looking for identical relations in base and target and using them to postulate potential matches. For each entity and predicate in the base, it finds the set of entities or predicates in the target that could plausibly match that item. These potential correspondences (*match hypotheses*) are determined by a set of simple rules: for example,

1. If two relations have the same name, create a match hypothesis.
2. For every match hypothesis between relations, check their corresponding arguments; if both are entities, or if both are functions, then create a match hypothesis between them.

For example, in Figure 7.5, Rule 1 creates match hypotheses between the GREATER-THAN relations occurring in base and target. Then Rule 2 creates match hypotheses between their arguments, since both are functions. Note that at this stage the system is entertaining two different, and inconsistent, match hypotheses involving GREATER THAN: one in which PRESSURE is matched with TEMPERATURE and one in which DIAMETER is matched with TEMPERATURE. Thus, at this stage, the program will have a large number of local matches. It gives these local matches *evidence scores*, based on a set of local evidence rules. For example, evidence for a match increases if the base and target predicates have the same name. More interestingly, the evidence rules also invoke systematicity, in that the evidence for a given

match increases with the evidence for a match among the parent relations – that is, the immediately governing higher-order relations.

2. Constructing global matches. The next stage is to collect systems of matches that use consistent entity pairings. SME first propagates entity correspondences up each relational chain to create systems of match hypotheses that use the same entity pairings. It then combines these into the largest possible systems of predicates with consistent object mappings. These global matches (called Gmaps) are SME's possible interpretations of the comparison.

An important aspect of SME is that the global matches (Gmaps) sanction *candidate inferences*: predicates from the base that get mapped into the target domain. These are base predicates that were not originally present in the target, but which can be imported into the target by virtue of belonging to a system that is shared by base and target. Thus, associated with each Gmap is a (possibly empty) set of *candidate inferences*. For example, in the "winning" Gmap (as discussed below), the pressure-difference causal chain in water is matched with the temperature-difference chain in heat, and water flow is matched with heat flow. However, you may recall that the initial heat representation lacked any causal link between the temperature difference and the heat flow (see Figure 7.5). In this case, the system brings across the higher-order predicate CAUSE from the water domain to the heat domain. In essence, it postulates that there may be more structure in the target than it initially knew about. Thus the resulting candidate inference in the heat domain is

CAUSE [GREATER [TEMP(coffee), TEMP(ice)],
FLOW(heat, bar, coffee, ice)].

3. Evaluating global matches. The global matches are then given a structural evaluation, which can depend on their local match evidence, the number of candidate inferences they support, and their graph-theoretic structure – for example, the depth of the relational system.[11] In this example, the winning Gmap is the pressure–temperature match discussed above, with its candidate inference of a causal link in the heat domain. Other Gmaps are also derived, including a Gmap that matches diameter with temperature and another particularly simple Gmap that matches LIQUID (water) with LIQUID (coffee). But these are given low evaluations. They contain fewer predicates than the winning Gmap and, at least equally important, they have shallower relational structures.

A few points should be noted about the way the structure-mapping engine works. First, SME's interpretation is based on selecting the deepest – that is, most systematic – consistent mappable structure. Computing a structurally consistent relational match precedes and determines the final selection of object correspondences.

Second, SME's matching process is entirely structural. That is, it attends only to properties such as identity of predicates, structural consistency (including 1–1 object pairings), and systematicity, as opposed to seeking specific kinds of content. Thus, although it operates on semantic representations, it is not restricted to any particular prespecified content. This allows it to act as a domain-general matcher. By promoting deep relational chains, the systematicity principle operates to promote predicates that participate in any mutually constraining system, whether causal, logical, or mathematical.

Third, as discussed above, different interpretations will be arrived at depending on which predicates match between two domains. For example, suppose that we keep the same base domain – the water system shown in Figure 7.5 – but change the target domain. Instead of two objects differing in *temperature*, let the target be two objects differing in their *specific heats*; say, a metal ball bearing and a marble. Assuming equal mass, they will also have different *heat capacities*. Now, the natural analogy concerns capacity differences in the base, rather than height differences. This is because the deepest relational chain that can be mapped to the target is, roughly, "Just as the container with greater diameter holds more water (levels being greater than or equal), so the object with greater heat capacity holds more heat (temperatures being greater than or equal)."

IMPLIES [AND (GREATER [DIAM (beaker), DIAM (vial)],
GREATER [LEVEL (beaker), LEVEL (vial)]),
GREATER [AMT-WATER (beaker), AMT-WATER (vial)]]

where AMT stands for the amount. This maps into the target as

IMPLIES [AND (GREATER [H-CAP (marble), H-CAP (ball)],
GREATER [TEMP (marble), TEMP (ball)]),
GREATER [AMT-HEAT (marble), AMT-HEAT (ball)]]

where H-CAP stands for heat capacity. This illustrates that the same base domain can yield different analogical mappings, depending on how it best matches the target.

Fourth, SME is designed as a general-purpose tool kit for similarity matching. It can operate with analogy rules, mere-appearance rules, or literal similarity rules, as discussed above.

Fifth, the matching process in SME is independent of the system's

problem-solving goals, although the learner's goals can influence the matcher indirectly, by influencing the domain representations present in working memory. Again, this represents a commitment to generality. The view is that analogy in problem solving is a special case of analogy.

An architecture for analogical reasoning

A complete model of analogical problem solving must take account of the context of reasoning, including the current plans and goals of the reasoner (Burstein, 1986; Carbonell, 1983; Holyoak, 1985; Kedar-Cabelli, 1985; Miller, Gallanter, & Pribram, 1960; Schank, 1982; Schank & Abelson, 1977). Indeed, as I discuss in the next section, some researchers have argued that plans and goals are so central in analogical reasoning that the analogy mechanism is built around them. However, analogies can occur outside of a goal-driven context. Further, the very fact that plans and goals influence all kinds of human thought processes, from transitive inference to the use of deductive syllogism, shows that they are not in themselves definitive of analogy. Somehow we need to capture the fact that analogy can be influenced by the goals of the problem solver while at the same time capturing what is specific about analogy.

I propose the architecture shown in Figure 7.6 for analogical reasoning. In this account, plans and goals influence our thinking *before* and *after* the analogy engine but not during its operation. Plans and goals and other aspects of current context influence the analogy process *before* the match by determining the working-memory representation of the current situation. This in turn influences what gets accessed from long-term memory. So, in the heat example, there are many aspects of the heat domain, but only the aspects currently represented in working memory are likely to influence remindings. Once a potential analog is accessed from long-term memory, the analogy processor runs its course. Here too the initial domain representation has strong effects, because it defines one input to the processor; thus it constrains the set of matches that will be found. This leads to "set" effects in problem solving; it is an advantage if we are thinking about the problem correctly and a disadvantage if we are not.

The analogy processor produces an interpretation, including candidate inferences and a structural evaluation. If the evaluation is too low – that is, if the depth and size of the system of matching predicates are too low – then the analogy will be rejected on structural grounds. If the analogy passes the structural criterion, then its candidate in-

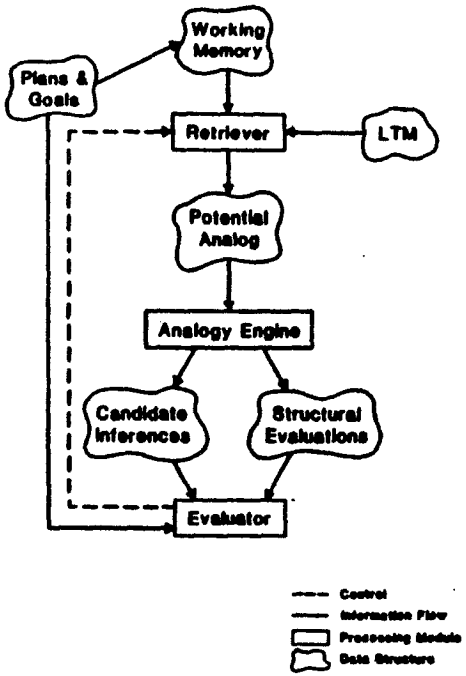

Figure 7.6. An architecture for analogical processing.

ferences must be evaluated to determine whether they are appropriate with respect to the goals of the reasoner. In terms of the computer model, this suggests adding a context-sensitive, expectation-driven module to evaluate the output of SME (Falkenhainer, Forbus, & Gentner, 1986, in press; Falkenhainer, 1987a). This extension is compatible with the combination models proposed by Burstein (1983) and Kedar-Cabelli (1985). Thus the key points of this proposed treatment of plans and goals are that (a) plans and goals constrain the inputs to the matcher, which is where they have their largest effect; (b) after the match three separate evaluation criteria must be invoked: structural soundness, relevance, and validity in the target; and (c) the match itself does not require a prior goal statement.

In the model proposed here, both structural properties and contextual-pragmatic considerations enter into analogical problem solving, but they are not equated. The analogy processor is a well-defined, separate cognitive module[9] whose results interact with other processes, analogous to the way some natural-language models have

postulated semiautonomous interacting subsystems for syntax, semantics, and pragmatics (e.g., Reddy, Erman, Fennell, & Neely, 1973). This allows us to capture the fact that analogy must satisfy *both* a structural and a pragmatic criterion.

Separating the planning context from the actual analogy processor represents a commitment to identifying processes common to analogy across different pragmatic contexts. It suggests that when comprehending an analogy in isolation, people use many of the same processes as they do to comprehend analogy in a problem-solving context. That is, they use the same structurally guided processor for both situations, simply adding or removing pragmatic context.[10] An advantage of modeling the matching process as structure-driven rather than goal-driven is that it allows for the possibility of finding unexpected matches, even perhaps matches that contradict the learner's initial problem-solving goals. Such unexpected outcomes are important in scientific discovery. For example, the mathematician Poincaré writes about an occasion on which he set out to prove a certain theorem and ended by discovering a class of functions that proved the theorem wrong. If we are ever to model such cases of unexpected creative discovery, the analogy process must be capable of finding matches that do not depend on – and may even contradict – the learner's current goals.

Competing views and criticisms of structure-mapping

Some aspects of structure-mapping have received convergent support in artificial intelligence and psychology. Despite differences in emphasis, there is widespread agreement on the basic elements of one-to-one mapping of objects and carry-over of predicates (Burstein, 1986; Carbonell, 1983; Hofstadter, 1984; Indurkhya, 1986; Kedar-Cabelli, 1985; Miller, 1979; Reed, 1987; Rumelhart & Norman, 1981; Tourangeau & Sternberg, 1981; Van Lehn & Brown, 1980; Verbrugge & McCarrell, 1977; and Winston, 1980, 1982). Further, all these researchers have some kind of selection principle – of which systematicity is one example – to filter which predicates matter in the match. But accounts differ on the nature of the selection principle. Many researchers use specific content knowledge or pragmatic information to guide the analogical selection process, rather than structural principles like systematicity. For example, Winston's (1980, 1982) system favors causal chains in its *importance-guided* matching algorithm. Winston (personal communication, November 1985) has also investigated goal-driven importance algorithms. Hofstadter and his colleagues have developed a connectionist like model of analogical mapping, in which systematicity is one of

several parallel influences on the mapping process (Hofstadter, 1984; Hofstadter, Mitchell, & French 1987).

Many accounts emphasize the role of plans and goals as part of the analogical mapping process. For example, some models combine a structure-mapping component with a plans-and-goals component in order to choose the most contextually *relevant* interpretation (e.g., Burstein, 1986; Kedar-Cabelli, 1985). These models use pragmatic context to select and elaborate the relevant predicates and to guide the mapping process. However, although these models have the ability to take contextual relevance into account, they also postulate a set of relatively constant structural processes that characterize analogical mapping. This view contrasts with a very different position, namely, that analogy should be seen as fundamentally embedded in a goal-driven problem-solving system. I now turn to a discussion of this second position.

The pragmatic account: an alternative to structure mapping

Holyoak (1985) proposed an alternative, *pragmatic*, account of analogical processing. Stating that analogy must be modeled as part of a goal-driven processing system, he argued that the structure-mapping approach is "doomed to failure" because it fails to take account of goals. In his proposed account, structural principles played no role; matching was governed entirely by the relevance of the predicates to the current goals of the problem solver. Because of the appeal of such a goal-centered position, I will discuss his arguments in some detail, even though Holyoak and his collaborators are now much less pessimistic concerning the utility of structural principles. I first present Holyoak's pragmatic account of analogy and then consider his critique of structure mapping.[11]

Holyoak states that "Within the pragmatic framework, the structure of analogy is closely tied to the mechanisms by which analogies are actually used by the cognitive system to achieve its goals" (Holyoak, 1985, p. 76). In the pragmatic account, the distinction between structural commonalties and surface commonalties is based solely on relevance. Holyoak's (1985, p. 81) definitions of these terms are as follows:

> It is possible, based on the taxonomy of mapping relations discussed earlier, to draw a distinction between *surface* and *structural* similarities and dissimilarities. An identity between two problem situations that plays no causal role in determining the possible solutions to one or the other analog constitutes a surface similarity. Similarly, a structure-preserving difference, as defined earlier constitutes a surface dissimilarity. In contrast, identities that influence goal attainment constitute structural similarities, and structure-violating differences constitute structural dissimilarities. Note that the distinction between surface and structural similarities, as used here, hinges on the relevance of the property in question to attainment of a successful solution. The distinction thus crucially depends on the goal of the problem solver.

Thus a *surface similarity* is defined as "an identity between two problem situations that plays no causal role in determining the possible solutions to one or the other analog," and *structural similarities* are "identities that influence goal attainment." The distinction between surface and structural similarities "hinges on the relevance of the property in question to attainment of a successful solution. The distinction thus crucially depends on the goal of the problem solver."

Holyoak's emphasis on plans and goals has some appealing features. This account promises to replace the abstract formalisms of a structural approach with an ecologically motivated account centered around what matters to the individual. Further, whereas structure-mapping requires both structural factors within the matcher and (in a complete account) pragmatic factors external to the matcher, Holyoak's account requires only pragmatic factors. But there are severe costs to this simplification. First, since structural matches are defined only by their relevance to a set of goals, the pragmatic account requires a context that specifies what is relevant before it can operate. Therefore, it cannot deal with analogy in isolation, or even with an analogy whose point is irrelevant to the current context. By this account, Francis Bacon's analogy "All rising to a great place is by a winding stair," should be uninterpretable in the present context. I leave it to the reader to judge whether this is true.

Holyoak (1985) seems aware of this limitation and states that his pragmatic account is meant to apply only to analogy in problem solving. But this means having to postulate separate analogy processors for analogy in context and analogy in isolation, which seems inconvenient at best. But there are further difficulties with the pragmatic account. Because the interpretation of an analogy is defined in terms of relevance to the initial goals of the problem solver, the pragmatic view does not allow for unexpected outcomes in an analogical match. This means that many creative uses of analogy – such as scientific discovery – are out of bounds. Finally, the pragmatic account lacks any means of capturing the important psychological distinction between an analogy that fails because it is irrelevant and an analogy that fails because it is unsound. In short, a good case can be made for the need to *augment* structural considerations with goal-relevant considerations (though I would argue that this should be done externally to the matcher, as shown in Figure 7.6, for example). However, the attempt to *replace* structural factors like

systematicity with pragmatic factors like goal-relevance does not appear tenable.

Holyoak raises three chief criticisms of structure-mapping (Holyoak, 1985, pp. 74, 75). First, as discussed above, Holyoak argues that structural factors are epiphenomenal: What really controls analogical matching is the search for goal-relevant predicates. The higher-order relations that enter into systematic structures "typically are such predicates as 'causes,' 'implies,' and 'depends on,' that is, causal elements that are pragmatically important to goal attainment. Thus, the pragmatic approach readily accounts for the phenomena cited as support for Gentner's theory." There are two problems with this argument. First, as discussed above, people are perfectly capable of processing analogy without any prior goal context, and of interpreting analogies whose point runs contrary to our expectations. Second, it is not correct to state that all higher-order relations are "causal elements pragmatically relevant to goal attainment." For example, *implies* (used in its normal logical sense) is not causal. Mathematical analogies, such as Polya's (1954) analogy between a triangle in a plane and a tetrahedron in space, are clear cases of shared relational structure that is not causal, and that need not be goal-relevant to be appreciated. Hofstadter (1984) provides many examples of analogies based on purely structural commonalities: for example, if $abc \longrightarrow abd$, then $pqr \longrightarrow pqs$.

Holyoak's second point is one of definition. In structure-mapping the distinction between analogy and literal similarity is based on the kinds of predicates shared: Analogy shares relational structure only, whereas literal similarity shares relational structure plus object descriptions. Holyoak proposes a different distinction: that analogy is similarity with reference to a goal. Thus "Even objects that Gentner would term 'literally similar' can be analogically related if a goal is apparent." The problem with this distinction is that although it captures analogy's role as a focusing device, it classifies some things as analogy that intuitively seem to be literal similarity. For example, consider the comparison "This '82 Buick is like this '83 Buick: You can use it to drive across town." By Holyoak's criterion this is an analogy, because a specific goal is under consideration; yet to my ear the two Buicks are literally similar whether or not a goal is involved. But since this is essentially a question of terminology, it may be undecidable.

Holyoak's third set of criticisms is based on the misinterpretation discussed earlier: namely, that in structure-mapping the systematicity of the base domain *by itself* determines the interpretation of an analogy, so that "the mappable propositions can be determined by a syntactic [structural] analysis of the source analog alone." This is false except in the rare case where nothing at all is known about the target (the "pure carry-over" case discussed earlier). This can be seen in the operation of SME, in which the interpretation arises out of a detailed *match* between base and target and not from "a syntactic analysis of the source analog alone." (See Skorstad, Falkenhainer, & Gentner, 1987, for examples of how SME yields different interpretations when the same base domain is paired with different targets.) At the risk of belaboring the point, let us recall that, in structure-mapping, analogy is seen as a subclass of similarity, and therefore, as with any other kind of similarity comparison, its interpretation is based on the best match between base and target. What distinguishes analogy from other kinds of similarity is that, for analogy, the *best match* is defined as the maximally systematic and consistent match of relational structure.

In summary, Holyoak's pragmatic account must be considered a failure insofar as it seeks to replace structure with relevance. Though one may sympathize with the desire to take plans and goals into account, discounting structure is the wrong way to go about it. Nonetheless, this work, like that of Burstein (1986), Carbonell (1981, 1983), and Kedar-Cabelli (1985), has the merit of calling attention to the important issue of how plans and goals can be integrated into a theory of analogy.

Separating structural rules from pragmatics has some significant advantages: It allows us to capture the commonalities among analogy interpretation across different pragmatic contexts, including analogy in isolation; it allows for creativity, since the processor does not have to know in advance which predicates are going to be shared; and it allows us to capture the difference between relevance and soundness. However, if the two-factor scheme I propose in Figure 7.6 is correct, there is still much work to be done in specifying exactly how plans and goals affect the initial domain representations that are given to the analogy processor and how they are compared with the output of this processor in the postprocessing stage.

Psychological evidence for structure-mapping

Ideal mapping rules. Structure-mapping claims to characterize the implicit competence rules by which the meaning of an analogy is derived. The first question to ask is how successfully it does so –

whether people do indeed follow the rules of structure-mapping in interpreting analogies. The prediction is that people should include relations and omit object descriptions in their interpretations of analogy. To test this, subjects were asked to write out descriptions of objects and then to interpret analogical comparisons containing these objects. (Gentner, 1980, 1988; Gentner & Clement, in press). They also rated how apt (how interesting, clever, or worth reading) the comparisons were.

The results showed that, whereas object descriptions tended to include both relational and object-attribute information, the interpretations of comparisons tended to include relations and omit object attributes. For example, a subject's description of "cigarette" was as follows:

chopped cured tobacco in a paper roll / with or without a filter at the end / held in the mouth / lit with a match and breathed through to draw smoke into the lungs / found widely among humans / known by some cultures to be damaging to the lungs / once considered beneficial to health.

Note that this description contains both relational and attributional information. Yet, when the same subject is given the metaphor "Cigarettes are like time bombs," his interpretation is purely in terms of common relational information: "They do their damage after some period of time during which no damage may be evident." A second finding was that the comparisons were considered more apt to the degree that subjects could find relational interpretations. There was a strong positive correlation between rated aptness and relationality but no such correlation for attributionality. Adults thus demonstrate a strong relational focus in interpreting metaphor. They emphasize relational commonalties in their interpretations when possible, and they prefer metaphors that allow such interpretations (Gentner & Clement, in press).

Development of mapping rules. The implicit focus on relations in interpreting analogy can seem so natural to us that it seems to go without saying. One way to see the effects of the competence rules is to look at cases in which these rules are not followed. Children do not show the kind of relational focus that adults do in interpreting analogy and metaphor.[12] A 5 year-old, given the figurative comparison "A cloud is like a sponge," produces an attributional interpretation, such as "Both are round and fluffy." A typical adult response is "Both can hold water for some time and then later give it back." Nine-year-olds are intermediate, giving some relational interpretations but also

many responses based on common object attributes (Gentner, 1980, 1988; Gentner & Stuart, 1983). The same developmental shift holds for choice tasks and rating tasks (Billow, 1975; Gentner, 1988). Thus there is evidence for a developmental shift from a focus on common object attributes to a focus on common relations in analogical processing.

Performance factors in analogical mapping

As Palmer (this volume) points out, structure-mapping aims first and foremost to capture the essential nature of analogy: what constitutes an analogy and which distinctions are necessary to characterize analogy – what Marr (1982) calls the "computational level" and Palmer and Kimchi (1985) call "informational constraints." Thus structure-mapping is in part a competence theory in that it attempts to capture people's implicit understanding of which commonalities should belong to analogy and which should not. The research described above suggests that under ordinary conditions structure-mapping is also a good approximation to a performance theory, for people's actual interpretations of analogies fit the predictions rather well. But what happens if we make it harder for people to perform according to the rules? Given that the ideal in analogy is to discover the maximal common higher-order relational structure, here we ask how closely people approach the *ideal* under difficult circumstances and what factors affect people's *performance* in carrying out a structure mapping.

Transfer performance. Gentner and Toupin (1986) posed this question developmentally. We asked children of 4–6 and 8–10 years of age to transfer a story plot from one group of characters to another. Two factors were varied: (a) the *systematicity* of the base domain (the original story); and (b) the *transparency* of the mapping (that is, the degree to which the target objects resembled their corresponding base objects). The systematicity of the original story was varied by adding beginning and ending sentences that expressed a causal or moral summary. Otherwise, the stories in the systematic condition were the same as those in the nonsystematic condition. Transparency was manipulated by varying the similarity of corresponding characters. For example, the original story might involve a *chipmunk* helping his friend the *moose* to escape from the villain *frog*.

After acting out the story with the base characters, the child was told to act out the story again, but with new characters. In the high-

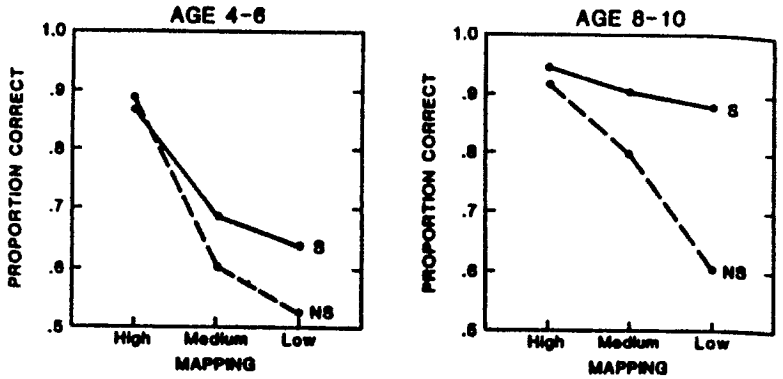

Figure 7.7. Results of the cross-mapping experiment: proportion correct on transfer story given systematic (S) or nonsystematic (NS) original stories, across mappings varying from high transparency to low transparency (Gentner & Toupin, 1986). *High* transparency means similar characters in corresponding roles; *medium*, different characters; and *low*, similar characters in different roles (the cross-mapped condition).

transparency mapping, the new characters would resemble the original characters: for example, a *squirrel*, an *elk*, and a *toad*, respectively. In the medium-transparency condition, three new unrelated animals were used. In the low-transparency (*cross-mapped*) condition, the characters were similar to the original characters but occupied noncorresponding roles: For example, the *chipmunk*, *moose*, and *frog* of the original story would map onto an *elk*, a *toad*, and a *squirrel*, respectively. We expected the cross-mapped condition to be very difficult. More interestingly, we wanted to know how robust the mapping rules are: How firmly can people hold to a systematic mapping when surface similarity pushes them toward a nonsystematic solution?

Both systematicity and transparency turned out to be important in determining transfer accuracy. However the two age groups showed different patterns, Transparency affected both age groups, whereas systematicity affected only the older group. For both ages, transfer accuracy was nearly perfect with highly similar corresponding characters (high transparency), lower when corresponding characters were quite different (medium transparency), and lower still in the cross-mapped condition (low transparency). For the older group, systematicity also had strong effects. As Figure 7.7 shows, 9-year-olds performed virtually perfectly, even in the most difficult mapping conditions, when they had a systematic story structure. This is noteworthy because, as can be seen from their performance in the non-

systematic condition, 9-year-olds found the cross-mapping condition quite difficult. Yet, given a systematic relational structure to hold onto, they could keep their mappings straight. In contrast, the 5-year-olds were affected only by transparency; they showed no significant benefit from systematic relational structure.

How does this happen? Gentner and Toupin (1986) speculated that the benefit comes in part from the way shared systems of relations help guide the mapping of lower-order relations. An error made in mapping a particular relation from base to target is more likely to be detected if there is a higher-order relation that constrains that lower-order relation. Informal observations in our study support this view. The older children, in the systematic condition, would sometimes begin to make an object-similarity-based error and then correct themselves, saying something like "Oh no, it's the *bad* one who got stuck in the hole, because he ate all the food." They were using the systematic causal structure of the story to overcome their local mapping difficulties.

Research with adults suggests that both systematicity and transparency continue to be important variables. Both Ross (1984; this volume) and Reed (1987) have shown that subjects are better at transferring algebraic solutions when corresponding base and target objects are similar. Reed (1987) measured the transparency of the mapping between two analogous algebra problems by asking subjects to identify pairs of corresponding concepts. He found that transparency was a good predictor of their ability to notice and apply solutions from one problem to the other. Ross (1987) has investigated the effects of cross-mappings in remindings during problem solving. He found that, even though adults could often *access* the prior problem, their ability to *transfer* the solution correctly was disrupted when cross-mapped correspondences were used. Robert Schumacher and I found benefits of both systematicity and transparency in transfer of device models, using a design similar to that of Gentner and Toupin (1986), in which subjects transfer an operating procedure from a base device to a target device (Gentner & Schumacher, 1986; Schumacher & Gentner, in press).

The evidence is quite strong, then, that transparency makes mapping easier. Thus literal similarity is the easiest sort of mapping and the one where subjects are least likely to make errors. The evidence also shows that a systematic base model promotes accurate mapping. This means that systematicity is a performance variable as well as a competence variable. Not only do people *believe* in achieving systematic mappings; they *use* systematic structure to help them perform the mapping.

Developmental implications: the relational shift. Like adults, the 9-year-olds in the Gentner and Toupin (1986) study were affected by both systematicity and transparency. But the 5-year-olds showed no significant effects of systematic base structure. All that mattered to this younger group was the transparency of the object correspondences. These results are consistent with the results reported earlier, and with the general developmental finding that young children rely on object-level similarities in transfer tasks (DeLoache, in press; Holyoak, Junn, & Billman, 1984; Keil & Batterman 1984; Kemler, 1983; Shepp, 1978; L. Smith, this volume; Smith & Kemler, 1977) and in metaphor tasks (Asch & Nerlove, 1960; Billow, 1975; Dent, 1984; Gardner, Kircher, Winner, & Perkins, 1975; Kogan, 1975). These findings suggest a developmental shift from reliance on surface similarity, and particularly on transparency of object correspondences, to use of relational structure in analogical mapping.[13]

Access processes

Now we are ready to tackle the issue of *access* to analogy and similarity. Before doing so, let us recapitulate briefly. I proposed at the start of this chapter a set of subprocesses necessary for spontaneous learning by analogy: (a) accessing the base system; (b) performing the mapping between base and target; (c) evaluating the match; (d) storing inferences in the target; and (e) extracting the common principle. So far we have considered mapping, evaluating, and making inferences. A major differentiating variable in the research so far is *similarity class*: whether the match is one of mere appearance, analogy, or literal similarity. Now we ask how similarity class affects memorial *access* to analogy and similarity.

Accessing analogy and similarity. What governs spontaneous access to similar or analogous situations? Gentner & Landers (1985) investigated this question, using a method designed to resemble natural long-term memory access. (For details of this and related studies, see Gentner & Landers, 1985; Gentner & Rattermann, in preparation; Rattermann & Gentner, 1987.) We first gave subjects a large set of stories to read and remember (18 key stories and 14 fillers). Subjects returned about a week later and performed two tasks: (a) a *reminding* task; and (b) a *soundness-rating* task.

In the reminding task, subjects read a new set of 18 stories, each of which matched one of the 18 original stories, as described below.

Table 7.2. *Sample story set for the access experiment (Gentner & Landers, 1985)*

BASE story

Karla, an old hawk, lived at the top of a tall oak tree. One afternoon, she saw a hunter on the ground with a bow and some some crude arrows that had no feathers. The hunter took aim and shot at the hawk but missed. Karla knew the hunter wanted her feathers so she glided down to the hunter and offered to give him a few. The hunter was so grateful that he pledged never to shoot at a hawk again. He went off and shot deer instead.

True-analogy TARGET

Once there was a small country called Zerdia that learned to make the world's smartest computer. One day Zerdia was attacked by its warlike neighbor, Gagrach. But the missiles were badly aimed and the attack failed. The Zerdian government realized that Gagrach wanted Zerdian computers so it offered to sell some of its computers to the country. The government of Gagrach was very pleased. It promised never to attack Zerdia again.

Mere-appearance TARGET

Once there was an eagle named Zerdia who donated a few of her tailfeathers to a sportsman so he would promise never to attack eagles. One day Zerdia was nesting high on a rocky cliff when she saw the sportsman coming with a crossbow. Zerdia flew down to meet the man, but he attacked and felled her with a single bolt. As she fluttered to the ground Zerdia realized that the bolt had her own tailfeathers on it.

False-analogy TARGET

Once there was a small country called Zerdia that learned to make the world's smartest computer. Zerdia sold one of its supercomputers to its neighbor, Gagrach, so Gagrach would promise never to attack Zerdia. But one day Zerdia was overwhelmed by a surprise attack from Gagrach. As it capitulated the crippled government of Zerdia realized that the attacker's missiles had been guided by Zerdian supercomputers.

Subjects were told that if any of the new stories reminded them of any of the original stories, they were to write out the original story (or stories) as completely as possible. There were three kinds of similarity matches between base and target:

- *mere appearance* (MA): object attributes and first-order relations match
- *true analogy* (TA): first-order relations and higher-order relations match
- *false analogy* (FA): only the first-order relations match

In all three cases, the base and target shared first-order relations. Other commonalties were added to create the different similarity

conditions. Table 7.2 shows an example set of four stories: a base story plus one example of each of the three kinds of matches. Each subject received one-third MA, one-third TA, and one-third FA matches, counterbalanced across three groups. After the subjects had completed the reminding task, they performed the soundness-rating task. They were shown their 18 pairs of stories side by side and asked to rate each pair for the soundness or inferential power of the match (with 5 being "sound" and 1 being "spurious").

In the soundness-rating task, subjects showed the predicted preference for true analogies. The mean soundness ratings were 4.4 for true analogy, 2.8 for mere appearance, and 2.0 for false analogy, with the only significant difference being between true analogy and the other two match types. This aspect of the study provides further evidence for the systematicity principle: Common higher-order relational structure is an important determinant of the subjective goodness of an analogy.

The results for access were surprising. Despite subjects' retrospective agreement that only the analogical matches were sound, their natural remindings did not produce analogies. Instead, they were far more likely to retrieve superficial mere-appearance matches. Given mere-appearance matches, subjects were able to access the original story 78% of the time, whereas the true analogies were accessed only 44% of the time, and the false analogies 25% of the time. All three differences were significant, suggesting that (*a*) surface commonalities have the most important role in access but that (*b*) higher-order relational commonalties – present in the true analogies but not in the false analogies – also promote access.

We have recently replicated these results, adding a literal similarity condition, and the results show the same pattern (Gentner & Rattermann, in preparation; Rattermann & Gentner, 1987). In access, surface similarity seems to be the dominant factor. Literal similarity and mere-appearance matches are more accessible than true analogies and false analogies. In soundness, systematicity of relational structure is the dominant factor. True analogy and literal similarity are considered sound, and false analogies and mere-appearance matches are not. Interestingly, surface information is superior in access even for subjects who clearly believe that only structural overlap counts toward soundness. It appears that analogical *access* and analogical *soundness* – or at least our subjective estimates of soundness – are influenced in different degrees by different kinds of similarity.

These access results accord with the findings of Gick and Holyoak (1980, 1983) of Reed (1987; Reed, Ernst, & Banerji, 1974), and of

Ross (1984, 1987; Ross & Sofka, 1986). In this research it has reliably been demonstrated that subjects in a problem-solving task often fail to access prior material that is analogous to their current problem. For example, in Gick and Holyoak's (1980, 1983) studies, a substantial number of subjects failed to access a potential analog – and therefore could not solve the problem – yet, when the experimenter suggested that the prior material was relevant, they could readily apply it to solve the problem. This means that (a) they had clearly stored the prior analog; (b) the prior analog contained sufficient information to solve their current problem; but (c) they could not access the prior analog solely on the basis of the current (analogous) problem structure. Thus there is converging evidence for the gloomy finding that relational commonalities often fail to lead to access.

There is also confirmation for the other side of the coin: that surface commonalties do promote access (Holyoak & Thagard, this volume; Novick, 1988; Reed & Ackinclose, 1986; Ross, 1984, 1987, this volume; Ross & Sofka, 1986; Schumacher, 1987). For example, Ross (1984) found clear effects of surface similarity in determining which earlier algebra problems subjects would be reminded of in trying to solve later problems. Reed and Ackinclose (1986) found that perceived similarity, rather than structural isomorphism, was the best predictor of whether subjects solving algebra problems would apply the results of a previous problem to a current problem.[14] Overall similarity, especially surface similarity, appears to be a major factor in accessing material in long-term memory.

Having said all this, we must remember that purely relational reminding does occur. Even young children sometimes experience analogical insights, as attested by Heida's analogy at the beginning of this chapter. As Johnson-Laird (this volume) points out, though remindings between remote domains are relatively rare, their occurrence sometimes sparks important creative advances (Falkenhainer, 1987b; Gentner, 1982; Hesse, 1966; Waldrop, 1987). A correct model of access will have to capture both the fact that relational remindings are comparatively rare and the fact that they occur.

Decomposing similarity

I began this chapter by noting that similarity is widely considered to be an important determinant of transfer (Thorndike, 1903; see Brown, this volume, and Brown & Campione, 1984, for discussions of this issue). The research reviewed here suggests that both *similarity*

and *transfer* may be too coarse as variables. A strong theme in this chapter, and indeed a convergent theme across this volume, has been the need to make finer differentiations in the notion of similarity (Collins & Burstein, this volume; Medin & Ortony, this volume; Rips, this volume; Ross, this volume; L. Smith, this volume). The research discussed in this chapter further suggests that *transfer* must be decomposed into different subprocesses that interact differently with different kinds of similarity. Thus the simple statement "Similarity is important in transfer" may conceal an intricate set of interactions between different varieties of similarity and different subprocesses in transfer.

Based on the research presented so far, it appears that different subprocesses are affected by different kinds of similarity. *Access* is strongly influenced by surface similarity and only weakly influenced by structural similarity. ***Analogical mapping*** is strongly influenced by structural similarity, including shared systematicity; it may also be weakly influenced by surface similarity. ***Judging soundness*** is chiefly influenced by structural similarity and systematicity. Finally, ***extracting and storing the principle*** underlying an analogy seems likely to be governed by structural similarity and systematicity. There is thus a relational shift in processing analogy and similarity from surface to structural commonalities.[15]

Similarity-based access may be a rather primitive mechanism, a low-cost low-specificity, high-quantity process, requiring little conscious effort. Analogical mapping and judging soundness are rather more sophisticated. They are often somewhat effortful, they often involve conscious reasoning, and, unlike access, they can be specifically tailored to different kinds of similarity. One can choose whether to carry out a mapping as an analogy or as a mere-appearance match, for example; but one cannot choose in advance whether to *access* an analogy or a mere-appearance match. Access has the feel of a passive process that simply produces some number of potential matches that the reasoner can accept or reject. Finally, one suspects that the processes of mapping and judging soundness are heavily influenced by culturally learned strategies (see Gentner & Jeziorski, in press). In contrast, access processes seem less amenable to cultural influence and training.[16] To the extent that experts differ from novices in their access patterns, I suspect this results chiefly from experts' having different domain representations (e.g., possessing relational abstractions) rather than from their having different access processes.

It is tempting to speculate that similarity-driven access involves something rather like a ballistic process, whereas mapping and judging soundness are more like discretionary processes. In any case, as we move from access to mapping and judging soundness there is a sense of increasing volitional control over the processes. To use an analogy, gaining access to long-term memory is a bit like fishing: The learner can bait the hook – that is, set up the working memory probe – as he or she chooses, but once the line is thrown into the water it is impossible to predict exactly which fish will bite.

The access bias for overall-similarity and surface-similarity matches rather than abstract analogical remindings may seem like a poor design choice from a machine-learning standpoint. But there may be good reasons for a bias toward overall similarity. First, a conservative, overall-similarity bias may be reasonable given the large size of human data bases relative to current artificial intelligence systems. For large data bases, the costs of checking all potential relational matches may well be prohibitive. Second, a conservative matching strategy might be prudent for mobile biological beings, for whom a false positive might be perilous. Third, by beginning with overall similarity the learner allows the relational vocabulary to grow to fit the data. This may be one reason children are better language learners than are adults; paradoxically, their initial conservatism and surface focus may allow the correct relational generalizations slowly to emerge (cf. Newport, 1984; see Forbus & Gentner, 1983; Murphy & Medin, 1985).

These arguments suggest that human access is geared toward literal similarity. But what about the fact that our access mechanisms also retrieve mere-appearance matches? Possibly, this comes about as a by-product of the overall-similarity bias. By this account, it is a design flaw, but perhaps a fairly minor one for concrete physical domains, where appearances tend not to be very deceiving. Very often, things that look alike *are* alike. (See Gentner, 1987; Medin & Ortony, this volume; Wattenmaker, Nakamura, & Medin, 1986.) Where surface matches become least reliable is in abstract domains such as algebra or Newtonian mechanics. The novice who assumes that any new pulley problem should be solved like the last pulley problem will often be wrong (Chi, Feltovich, & Glaser, 1981). Thus our surface-oriented accessor can be an obstacle to learning in abstract domains, where the correlation between surface features and structural features is low.

Implications for learning

Now let's put together these findings and ask how they bear on experiential learning. This discussion is based on that given by Forbus

and Gentner (1983). Forbus and Gentner examined the role of similarity comparisons in the progression from early to later representations. A key assumption here is that implicit comparisons among related knowledge structures are important in learning (Brooks, 1978; Jacoby & Brooks, 1984; Medin & Schaffer, 1978; Wattenmaker et al., 1986). We conjecture that much of experiential learning proceeds through spontaneous comparisons – which may be implicit or explicit – between a current situation and prior similar or analogous situations that the learner has stored in memory. We also assume that early representations are characteristically rich and perceptually based. That is, early domain representations differ from more advanced representations of the same domain in containing more perceptual information specific to the initial context of use. What does this predict? First, in terms of *access*, the greater the surface match the greater the likelihood of access. Thus the matches that are likely to occur most readily are literal similarity matches and mere-appearance matches.

Once the base domain has been accessed, the mapping process occurs. To transfer knowledge from one domain to another, a person must not only access the base domain but also set up the correct object correspondences between the base and target and map predicates across. At this level, a mix of deep and surface factors seems to operate. Systematicity and structural similarity become crucial, but so does the transparency of the object correspondences (Gentner & Toupin, 1986; Reed, 1987; Ross, 1987). It appears that, for adults and/or experts, systematicity can to some extent compensate for lack of transparency. The rules of analogy are clear enough and the relational structures robust enough to allow accurate mapping without surface support. But for children and novices surface similarity is a key determinant of success in analogical mapping.

To the extent that children and novices rely on object commonalities in similarity-based mapping, they are limited to literal similarity matches and mere-appearance matches. The disadvantage of mere-appearance matches is obvious: They are likely to lead to wrong inferences about the target. But even literal similarity matches have their limitations. Although adequate for prediction, literal similarity matches are probably less useful than analogies for purposes of explicitly extracting causal principles. In an analogical match, the shared data structure is sparse enough to permit the learner to isolate the key principles. In literal similarity, there are too many common predicates to know which are crucial (Forbus & Gentner, 1983; Ross, this volume; Wattenmaker et al., 1986).

How do learners escape the confines of literal similarity? One way, of course, is through explicit instruction about the relevant abstractions. But there may be ways within experiential learning as well. If we speculate that the results of a similarity comparison become slightly more accessible (Elio & Anderson, 1981, 1984; Gick & Holyoak, 1983; Ortony, 1979; Skorstad, Gentner, & Medin, 1988), then repeated instances of near-literal similarity could gradually increase the salience of the relational commonalities. At some point the relational structures become sufficiently salient to allow analogy to occur. Once this happens, there is some likelihood of noticing the relational commonalities and extracting them for future use. (This conjectural sequence, which is essentially that proposed in Forbus and G ntner, 1983, hinges on the claim that the results of an analogy are sparser and therefore more inspectable than the results of a literal similarity comparison. Hence, the probability of noticing and extracting the common relational structure is greater.) The extracted relational abstractions can then influence encoding. With sufficient domain knowledge, the set of known abstractions – such as *flow rate* or *positive feedback situation* – becomes firm enough to allow relational encoding and retrieval.

The post-access processes can be influenced both by individual training and by local strategies. I suspect that this is the area in which training in thinking skills can be of most benefit. For example, people may learn better skills for checking potential matches and rejecting bad matches, and perhaps also skills for tinkering with potential matches to make them more useful (Clement, 1983, 1986). However, I suspect that some parts of the system will always remain outside direct volitional control. To return to the fishing analogy, we can learn to bait the hook better, and once the fish bites we can learn better skills for landing it, identifying it, and deciding whether to keep it or throw it back. But no matter how accurate the preaccess and post-access processes, there is always uncertainty in the access itself. When we throw the hook into the current we cannot determine exactly which fish will bite. A strategically managed interplay between discretionary and automatic processes may be the most productive technique for analogical reasoning.

Conclusion

In this chapter I have suggested that different kinds of similarity participate differently in transfer. In particular, I have proposed de-

composing similarity into subclasses of *analogy*, *mere-appearance*, and *literal similarity* and transfer into subprocesses of *access*, *mapping*, *storing inferences*, and *extracting commonalities*. Although many issues remain to be worked out, it seems clear that this finer-grained set of distinctions will allow a more fruitful discussion of similarity based learning.

NOTES

This research was supported by the Office of Naval Research under Contract N00014–5–K–0559, NR667–551. The developmental studies were supported by the National Institute of Education under Contract 400–80–0031 awarded to the Center for the Study of Reading. I thank Ann Brown, Allan Collins, Judy DeLoache, Ken Forbus, Doug Medin, and Brian Ross for many insightful discussions of these issues, and Cathy Clement, Rogers Hall, Mike Jeziorski, Doug Medin, Andrew Ortony, Mary Jo Rattermann, Bob Schumacher, and Janice Skorstad for their helpful comments on a draft of this paper.

1 For now, I will use the term *analogical learning* to refer to both learning by analogy and learning by literal similarity. Later in the chapter I will distinguish between analogy and similarity.
2 This account has benefited from the comments and suggestions of my colleagues since my first proposal in 1980. Here and there I will indicate some ways in which the theory has changed.
3 The notation in Figure 7.2 is equivalent to a predicate calculus representation; I use it because it emphasizes structural relationships (see Norman & Rumelhart, 1975; Palmer, 1978).
4 In this analogy, the function PRESSURE in the water domain must be mapped onto TEMPERATURE in the heat domain. Like objects, functions on objects in the base can be put in correspondence with different functions in the target in order to permit mapping a larger systematic chain.
5 An ongoing question in our research is whether mere-appearance matches should be viewed as including first-order relations as well as object attributes.
6 The reason that attributes are ignored, rather than being placed in correspondence with other attributes, is to permit analogical matches between rich objects and sparse objects.
7 Adding functions to the representation is a change from my former position, which distinguished only between object attributes (one-place predicates) and relations (two-or-more-place predicates). I thank Ken Forbus, Brian Falkenhainer, and Janice Skorstad for discussions on this issue.
8 Currently, the global evaluation is extremely simple: The match-hypothesis evidence scores are simply summed for each Gmap. Although we have developed more elaborate schemes for computing the goodness of the Gmaps, this simple summation has proved extremely effective. We have tried SME on over 40 analogies, and in every case its highest-ranked Gmap is the one humans prefer.

9 The term *module* here should not be taken in the Fodorian sense. I assume that analogical processing is not innate or hard-wired but, at least in part, learned; nor do I assume that the analogy processor is impenetrable, although its workings may be opaque.
10 As in all top-down expectation situations, comprehension should be easier with a supporting context and harder when context leads to the wrong expectations; but the basic analogy processes do not *require* a context.
11 It should be noted that since this chapter was written Holyoak has revised his position. His recent work incorporates many of the structural constraints discussed here while still postulating a central role for contextual goals (Thagard & Holyoak, 1988).
12 Much of the developmental literature has been couched in terms of *metaphor* rather than *analogy*. Often, the items called *metaphors* are figurative comparisons that adults would treat as analogies.
13 It is not clear whether this shift is due to a developmental change in analytical reasoning skills or simply to an increase in domain knowledge, especially relational knowledge (Brown, this volume; Brown & Campione, 1984; Carey, 1984; Chi, 1978; Crisafi & Brown, 1986; Gentner, 1977a,b, 1988; Larkin, McDermott, Simon, & Simon, 1980; Reynolds & Ortony, 1980; Siegler, 1988; Vosniadou & Ortony, 1986).
14 These results, especially in problem-solving contexts, are problematic for the plan-based indexing view held by many researchers in artificial intelligence. See Gentner (1987) for a discussion.
15 This echoes the relational shift in the development of analogy from an early focus on surface commonalities to the adult focus on relational commonalities. How much we should make anything of this parallel is unclear.
16 We may perhaps learn to guide access by the indirect route of changing the contents of working memory so that a different set of matches arises. However, this is not a very fine-tuned method. I thank Brian Ross for discussions of this issue.

REFERENCES

Asch, S. E., & Nerlove, H. (1960). The development of double function terms in children: An exploratory investigation. In B. Kaplan & S. Wapner (Eds.), *Perspectives in psychological theory*. New York: International Universities Press.

Billow, R. M. (1975). A cognitive development study of metaphor comprehension. *Developmental Psychology, 11*, 415–423.

Brooks, L. (1978). Nonanalytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

Brown, A. L., & Campione, J. C. (1984). Three faces of transfer: Implications for early competence, individual differences, and instruction. In M. Lamb, A. Brown, & B. Rogoff (Eds.), *Advances in developmental psychology* (Vol. 3, pp. 143–192). Hillsdale, NJ: Erlbaum.

Buckley, S. (1979). *Sun up to sun down*. New York: McGraw-Hill.

Burstein, M. H. (1986). Concept formation by incremental analogical rea-

soning and debugging. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 351–370). Los Altos, CA: Kaufmann.

Carbonell, J. G. (1981, August). Invariance hierarchies in metaphor interpretation. *Proceedings of the Third Annual Conference of the Cognitive Science Society* (pp. 292–295), Berkeley, CA.

Carbonell, J. G. (1983). Learning by analogy: Formulating and generalizing plans from past experience. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 1, pp. 137–161). Palo Alto, CA: Tioga.

Carey, S. (1984). Are children fundamentally different kinds of thinkers and learners than adults? In S. F. Chipman, J. W. Segal, & R. Glaser (Eds.), *Thinking and learning skills: Current research and open questions* (Vol. 2, pp. 485–517). Hillsdale, NJ: Erlbaum.

Chi, M. T. H. (1978). Knowledge structures and memory development. In R. Siegler (Ed.), *Children's thinking: What develops?* Hillsdale, NJ: Erlbaum.

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

Clement, J. (1981). Analogy generation in scientific problem solving. *Proceedings of the Third Annual Conference of the Cognitive Science Society* (pp. 137–140), Berkeley, CA.

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

Clement, J. (1986, August). Methods for evaluating the validity of hypothesized analogies. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, Amherst, MA.

Clement, C. A., & Gentner, D. (1988). Systematicity as a selection constraint in analogical mapping. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal.

Crisafi, M. A., & Brown, A. L. (1986). Analogical transfer in very young children: Combining two separately learned solutions to reach a goal. *Child Development, 57*, 953–968.

DeLoache, J. S. (in press). The development of representation in young children. In H. W. Reese (Ed.), *Advances in child development and behavior* (Vol. 21). New York: Academic Press.

Dent, C. H. (1984). The developmental importance of motion information in perceiving and describing metaphoric similarity. *Child Development, 55*, 1607–1613.

Elio, R., & Anderson, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 397–417.

Elio, R., & Anderson, J. R. (1984). The effects of information and learning mode on schema abstraction. *Memory & Cognition, 12*, 20–30.

Falkenhainer, B. (1987a). *An examination of the third stage in the analogy process: Verification-based analogical reasoning* (Tech. Rep. UIUCDCS-R86-1302). Urbana-Champaign: University of Illinois, Department of Computer Science.

Falkenhainer, B. (1987b). Scientific theory formation through analogical inference. *Proceedings of the Fourth International Machine Learning Workshop* (pp. 218–229). Los Altos, CA: Kaufmann.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of the American Association for Artificial Intelligence* (pp. 272–277), Philadelphia.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (in press). The structure-mapping engine. *Artificial Intelligence*.

Forbus, K. D., & Gentner, D. (1983, June). Learning physical domains: Towards a theoretical framework. *Proceedings of the 1983 International Machine Learning Workshop* (pp. 198–202), Monticello, IL. Also in R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2). Los Altos, CA: Kaufmann, 1986.

Gardner, H., Kircher, M., Winner, E., & Perkins, D. (1975). Children's metaphoric productions and preferences. *Journal of Child Language, 2*, 1–17.

Gentner, D. (1977a). Children's performance on a spatial analogies task. *Child Development, 48*, 1034–1039.

Gentner, D. (1977b). If a tree had a knee, where would it be? Children's performance on simple spatial metaphors. *Papers and Reports on Child Language Development, 13*, 157–164.

Gentner, D. (1980). *The structure of analogical models in science* (BBN Tech. Rep. 4451). Cambridge, MA: Bolt, Beranek & Newman.

Gentner, D. (1982). Are scientific analogies metaphors? In D. Miall (Ed.), *Metaphor: Problems and perspectives*. Brighton: Harvester Press.

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

Gentner, D. (1987). Analogical inference and analogical access. In A. Prieditis (Ed.), *Analogica: Proceedings of the First Workshop in Analogical Reasoning*. London: Pitman.

Gentner, D. (1988). Metaphor as structure-mapping: The relational shift. *Child Development, 59*, 47–59.

Gentner, D., & Clement, C. A. (in press). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation*. New York: Academic Press.

Gentner, D., Falkenhainer, B., & Skorstad, J. (1987, January). Metaphor: The good, the bad, and the ugly. *Proceedings of the Third Conference on Theoretical Issues in Natural Language Processing* (pp. 155–159), Las Cruces, NM.

Gentner, D., & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (pp. 99–129). Hillsdale, NJ: Erlbaum.

Gentner, D., & Jeziorski, M. (in press). Historical shifts in the use of analogy in science. In B. Gholson, W. R. Shadish, & A. Graesser (Eds). *Psychology of Science*. Cambridge: Cambridge University Press.

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

Gentner, D., & Rattermann, M. J. (in preparation). Analogical access: A good match is hard to find.

Gentner, D., & Schumacher, R. M. (1986, October). Use of structure-mapping theory for complex systems. *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (pp. 252–258). Atlanta, GA.

Gentner, D., & Stuart, P. (1983). *Metaphor as structure-mapping: What develops* (Tech. Rep. 5479). Cambridge, MA: Bolt, Beranek & Newman.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10*, 277–300.

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

Gick, M. L., & 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.

Hall, R. (in press). Computational approaches to analogical reasoning: A comparative analysis. *Artificial Intelligence.*

Hesse, M. B. (1966). *Models and analogies in science.* Notre Dame, IN: University of Notre Dame Press.

Hofstadter, D. R. (1984). The Copycat project: An experiment in nondeterministic and creative analogies (Memo 755). Cambridge, MA: MIT, Artificial Intelligence Laboratory.

Hofstadter, D. R., Mitchell, T. M., & French, R. M. (1987). *Fluid concepts and creative analogies: A theory and its computer implementation* (Tech. Rep. 87–1). Ann Arbor: University of Michigan, Fluid Analogies Research Group.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 19, pp. 59–87). New York: Academic Press.

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

Indurkhya, B. (1986). Constrained semantic transference: A formal theory of metaphors. *Synthese, 68*, 515–551.

Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception and concept learning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 18, pp. 1–47). New York: Academic Press.

Kedar-Cabelli, S. (1985, August). Purpose-directed analogy. *Proceedings of the Seventh Annual Conference of the Cognitive Science Society* (pp. 150–159), Irvine, CA.

Kedar-Cabelli, S. (1988). Analogy: From a unified perspective. In D. H. Helman (Ed.), *Analogical reasoning: Perspectives of artificial intelligence, cognitive science, and philosophy.* Boston: Reidel.

Keil, F., & Batterman, N. A. (1984). A characteristic-to-defining shift in the development of word meaning. *Journal of Verbal Learning and Verbal Behavior, 23*, 221–236.

Kemler, D. G. (1983). Holistic and analytical modes in perceptual and cognitive development. In T. J. Tighe & B. E. Shepp (Eds.), *Perception, cognition, and development: Interactional analysis* (pp. 77–102). Hillsdale, NJ: Erlbaum.

Kogan, N. (1975, April). *Metaphoric thinking in children: Developmental and individual-difference aspects.* Paper presented at the meeting of the Society for Research in Child Development, Denver, CO.

Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science, 208*, 1335–1342.

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

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207–238.

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

Miller, G. A., Gallanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior.* New York: Holt, Rinehart & Winston.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289–316.

Newport, E. L. (1984). Constraints on learning: Studies in the acquisition of American Sign Language. *Papers and Reports on Child Language Development, 23*, 1–22.

Norman, D. A., Rumelhart, D. E., & LNR Research Group (1975). *Explorations in cognition.* San Francisco: Freeman.

Novick, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*, 510–520.

Ortony, A. (1979). Beyond literal similarity. *Psychological Review, 86*, 161–180.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259–302). Hillsdale, NJ: Erlbaum.

Palmer, S. E., & Kimchi, R. (1985). The information processing approach to cognition. In T. Knapp & L. C. Robertson (Eds.), *Approaches to cognition: Contrasts and controversies* (pp. 37–77). Hillsdale, NJ: Erlbaum.

Polya, G. (1954). *Induction and analogy in mathematics: Vol. 1. Of mathematics and plausible reasoning.* Princeton, NJ: Princeton University Press.

Rattermann, M. J., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. *Proceedings of the Ninth Annual Meeting of the Cognitive Science Society* (pp. 23–34), Seattle.

Reddy, D. R., Erman, L. D., Fennell, R. D., & Neely, R. B. (1973, August). HEARSAY Speech understanding system: An example of the recognition process. *Proceedings of the Third International Joint Conference on Artificial Intelligence* (pp. 185–193), Stanford University.

Reed, S. K. (1987). A structure-mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 124–139.

Reed, S. K., & Ackinclose, C. C. (1986). *Selecting analogous solutions: Similarity vs. inclusiveness.* Unpublished manuscript, Florida Atlantic University, Boca Raton.

Reed, S. K., Ernst, G. W., & Banerji, R. (1974). The role of analogy in transfer between similar problem states. *Cognitive Psychology, 6*, 436–450.

Reynolds, R. E., & Ortony, A. (1980). Some issues in the measurement of children's comprehension of metaphorical language. *Child Development, 51*, 1110–1119.

Ross, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology, 16*, 371–416.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 629–639.

Ross, B. H., & Sofka, M. D. (1986). *Remindings: Noticing, remembering, and using specific knowledge of earlier problems*. Unpublished manuscript, University of Illinois, Department of Psychology, Urbana-Champain.

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

Rumelhart, D. E., & Norman, D. A. (1981). Analogical processes in learning. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 335–339). Hillsdale, NJ: Erlbaum.

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

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.

Schumacher, R. M. (1987). *Similarity-based reminding: Effects of distance and encoding on retrieval*. Unpublished master's thesis, University of Illinois, Urbana-Champaign.

Schumacher, R., & Gentner, D. (in press). Transfer of training as analogical mapping. *Communications of the IEEE*.

Shepp, B. E. (1978). From perceived similarity to dimensional structure: A new hypothesis about perceptual development. In E. Rosch & B. B. Lloyd (Eds.) *Cognition and categorization* (pp. 135–167). Hillsdale, NJ: Erlbaum.

Seigler, R. S. (1988). Mechanisms of cognitive growth: Variation and selection. In R. J. Sternberg (Ed.), *Mechanisms of cogitive development*. Prospect Heights, IL: Waveland Press.

Skorstad, J., Falkenhainer, B., & Gentner, D. (1987). Analogical processing: A simulation and empirical corroboration. *Proceedings of the American Association for Artificial Intelligence*. Seattle.

Skorstad, J., Gentner, D., & Medin, D. (1988). Abstraction processes during concept learning: A structural view. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal.

Smith, L. B., & Kemler, D. G. (1977). Developmental trends in free classification: Evidence for a new conceptualization of perceptual development. *Journal of Experimental Child Psychology, 24*, 279–298.

Sternberg, R. J. (1977). Component processes in analogical reasoning. *Psychological Review, 84*, 353–378.

Thagard, P., & Holyoak, K. (1988, August). Analogical problem solving: A constraint satisfaction approach. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal.

Thorndike, E. L. (1903). *Educational psychology*. New York: Lemcke & Buechner.

Tourangeau, R., & Sternberg, R. J. (1981). Aptness in metaphor. *Cognitive Psychology, 13*, 27–55.

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

Van Lehn, K., & Brown, J. S. (1980). Planning nets: A representation for formalizing analogies and semantic models of procedural skills. In R. E. Snow, P. A. Federico, & W. E. Montague (Eds.), *Aptitude, learning, and instruction: Cognitive process analyses* (Vol. 2, pp. 95–137). Hillsdale, NJ: Erlbaum.

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

Vosniadou, S. & Ortony, A. (1986). Testing the metaphoric competence of the young child: Paraphrase versus enactment. *Human Development, 29*, 226–230.

Waldrop, M. (1987). Causality, structure, and common sense. *Science, 237*, 1297–1299.

Wattenmaker, W. D., Nakamura, G. V., & Medin, D. L. (1986). *Relationships between similarity-based and explanation-based categorization*. Unpublished manuscript, University of Illinois, Department of Psychology, Urbana-Champaign.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the Association for Computing Machinery, 23*, 689–703.

Winston, P. H. (1982). Learning new principles from precedents and exercises. *Artificial Intelligence, 19*, 321–350.