

Analogy just looks like high level perception: why a domain-general approach to analogical mapping is right

KENNETH D. FORBUS*, DEDRE GENTNER*,
ARTHUR B. MARKMAN‡, and RONALD W. FERGUSON*

* *Institute for the Learning Sciences, 1890 Maple Avenue,
Evanston, IL 60201, USA*

tel.: +1 847 491 7699

fax: +1 847 491 5258

email: forbus@ils.nwu.edu

‡ Columbia University

Abstract. Hofstadter and his colleagues have criticized current accounts of analogy, claiming that such accounts do not accurately capture interactions between processes of representation construction and processes of mapping. They suggest instead that analogy should be viewed as a form of *high level perception* that encompasses both representation building and mapping as indivisible operations within a single model. They argue specifically against SME, our model of analogical matching, on the grounds that it is modular, and offer instead programs such as Mitchell and Hofstadter's Copycat as examples of the high level perception approach. In this paper we argue against this position on two grounds. First, we demonstrate that most of their specific arguments involving SME and Copycat are incorrect. Second, we argue that the claim that analogy is high-level perception, while in some ways an attractive metaphor, is too vague to be useful as a technical proposal. We focus on five issues: (1) how perception relates to analogy, (2) how flexibility arises in analogical processing, (3) whether analogy is a domain-general process, (4) how micro-worlds should be used in the study of analogy, and (5) how best to assess the psychological plausibility of a model of analogy. We illustrate our discussion with examples taken from computer models embodying both views.

Keywords: analogy, similarity, structure-mapping, cognitive simulation, qualitative physics

Received 14 February 1997; revision accepted 30 March 1997

1. Introduction

The field of analogy is widely viewed as a cognitive science success story. In few other research domains has the connection between computational and psychological work been as close and as fruitful as in this one. This collaboration, along with significant influences from philosophy, linguistics, and history of science, has led to a substantial degree of theoretical and empirical convergence among researchers in the field (e.g. Falkenhainer et al. 1986, 1989, Holyoak and Thagard 1989, Halford 1992, Keane et al.

1994). There has been progress both in accounting for the basic phenomena of analogy and in extending analogy theory to related areas, such as metaphor and mundane similarity, and to more distant areas such as categorization and decision making (see Holyoak and Thagard 1995, Gentner and Holyoak 1997, Gentner and Markman 1997). Though there are still many debated issues, there is a fair degree of consensus on certain fundamental theoretical assumptions. These include the usefulness of decomposing analogical processing into constituent sub-processes, such as *retrieving* representations of the analogs, *mapping* (*aligning* the representations and *projecting inferences* from one to the other), *abstracting* the common system, and so on, and the fact that the mapping process is a domain-general process that is the core defining phenomenon of analogy (Gentner 1989).

Hofstadter and his colleagues express a dissenting view. They argue for a ‘high-level perception’ approach to analogy (Chalmers *et al.* 1992, Mitchell 1993, French 1995, Hofstadter, 1995a) and are sharply critical of the structure-mapping research program and related approaches. Indeed, Hofstadter (1995a, pp. 155–165) even castigates Waldrop (1987) and Boden (1991) for praising models such as SME and ACME. This paper is a response to these criticisms.

Hofstadter and his colleagues argue against most current approaches to modelling analogical reasoning. One of their major disagreements is with the assumption that mapping between two analogues can be separated from the process of initially perceiving both analogues. As Chalmers *et al.* put it: ‘We argue that perceptual processes cannot be separated from other cognitive processes even in principle, and therefore that traditional artificial-intelligence models cannot be defended by supposing the existence of a “representation module” that supplies representations ready-made’ (Chalmers *et al.* 1992, p. 185).

Hofstadter (1995a, pp. 284–285) is even more critical: ‘SME is an algorithmic but psychologically implausible way of finding what the structure-mapping theory would consider to be the best mapping between two given representations, and of rating various mappings according to the structure-mapping theory, allowing such ratings then to be compared with those given by people’. Hofstadter (1995b, p. 78) further charges analogy researchers with ‘trying to develop a theory of analogy making while bypassing both gist extraction and the nature of concepts...’ an approach ‘as utterly misguided as trying to develop a theory of musical aesthetics while omitting all mention of both melody and harmony’. Writing of Holyoak and Thagard’s (1995) approach to analogy, he states that it is ‘to hand shrink each real-world situation into a tiny, frozen caricature of itself, containing precisely its core and little else’.

Hofstadter and colleagues are particularly critical of the assumption that analogical mapping can operate over pre-derived representations and of the associated practice of testing the simulations using representations designed to capture what are believed to be human construals. ‘We believe that the use of hand-coded, rigid representations will in the long run prove to be a dead end, and that flexible, content-dependent, easily adaptable representations will be recognized as an essential part of any accurate model of cognition’ (Chalmers *et al.* 1992, p. 201). Rather, they propose the metaphor of ‘high-level perception’ in which perception is holistically integrated with higher forms of cognition. They cite Mitchell and Hofstadter’s Copycat model (Mitchell 1993) as a model of high-level perception. Chalmers *et al.* (1992) claim that the flexibility of human cognition cannot be explained by any more modular account.

We disagree with many of the theoretical and empirical points made by Hofstadter and his colleagues. In this paper we present evidence that the structure-mapping

algorithm embodied in the SME approach can capture significant aspects of the psychological processing of analogy. We consider and reply to the criticisms made against SME and correct some of Hofstadter's (1995a) and Chalmers *et al.*'s (1992) claims that are simply untrue factually. We begin in Section 2 by summarizing Chalmers *et al.*'s notion of high-level perception and outlining general agreements and disagreements. Section 3 describes the simulations of analogical processing involved in the specific arguments: SME (and systems that use it) and Copycat. This section both clears up some of the specific claims Chalmers *et al.* make regarding both systems, and provides the background needed for the discussion in Section 4. There we outline five key issues in analogical processing and compare our approach with that of Chalmers *et al.* (1992) with regard to them. Section 5 summarizes the discussion.

2. Chalmers *et al.*'s notion of high level perception

Chalmers *et al.* (1992) observe that human cognition is extraordinarily flexible, far more so than is allowed for in today's cognitive simulations. They postulate that this flexibility arises because, contrary to most models of human cognition, there is no separation between the process of creating representations from perceptual information and the use of these representations. That is, for Chalmers *et al.* there *is* no principled decomposition of cognitive processes into 'perceptual processes' and 'cognitive processes'. While conceding that it may be possible informally to identify aspects of our cognition as either perception or cognition, Chalmers *et al.* claim that building a computational model that separates the two cannot succeed. Specifically, they identify analogy with 'high-level perception', and argue that this holistic notion cannot productively be decomposed.

One implication of this view is that cognitive simulations of analogical processing must always involve a 'vertical' slice of cognition (see (Morrison and Dietrich 1995) for a similar discussion). That is, a simulation must automatically construct its internal representations from some other kind of input, rather than being provided with them directly by the experimenters. In Copycat, for instance, much of the information used to create a match in a specific problem is automatically generated by rules operating over a fairly sparse initial representation. Chalmers *et al.* point out that Copycat's eventual representation of a particular letter-string is a function of not just the structure of the letter-string itself, but also of the other letter-strings it is being matched against.




Figure 1. Three abstract views of perception and cognition.

2.1. Overall points of agreement and disagreement

Chalmers *et al.*'s view of analogy as high-level perception has its attractive features. For instance, it aptly captures a common intuition that analogy is 'seeing as'. For example, when Rutherford thought of modelling the atom as if it were the solar system, he might be said to have been 'perceiving' the atom as a solar system. It further highlights the fact that analogical processing often occurs outside of purely verbal situations. Yet while we find this view is in some respects an attractive metaphor, we are less enthusiastic about its merits as a technical proposal, especially the claim of the inseparability of the processes.

We agree with Chalmers *et al.* that understanding how analogical processing interacts with perception and other processes of building representations is important but we disagree that such interactions necessitate a holistic account. Figure 1 illustrates three extremely coarse-grained views of how perception and cognition interact. Part (a) depicts a classic stage model, in which separate processes occur in sequence. This is the straw man that Chalmers *et al.* argue against. Part (b) depicts Chalmers *et al.*'s account. The internal structure either is not identifiable in principle (the literal reading of Chalmers *et al.*'s claims) or the parts interact so strongly that they cannot be studied in isolation (how Chalmers *et al.* actually conduct their research). Part (c) depicts what we suggest is a more plausible account. The processes that build representations are interleaved with the processes that use them. With this view, there is value in studying the processes in isolation, as well as in identifying their connections with the rest of the system. We will return to this point in Section 3.

3. A comparison of some analogical processing simulations

Hofstadter's claims concerning how to simulate analogical processing can best be evaluated in the context of the models. We now turn to the specific simulations under discussion, SME and Copycat.

3.1. Simulations using structure-mapping theory

Gentner's (1983, 1989) *structure-mapping* theory of analogy and similarity decomposes analogy and similarity processing into several processes (not all of which occur for every instance of comparison), including representation, access, mapping (alignment and inference), evaluation, adaptation, verification, and schema-abstraction. For instance, the *mapping* process operates on two input representations, a *base* and a *target*. It results in one or a few *mappings*, or interpretations, each consisting of a set of correspondences between items in the representations and a set of *candidate inferences*, which are surmises about the target made on the basis of the base representation plus the correspondences. The set of constraints on correspondences include *structural consistency*, i.e. that each item in the base maps to at most one item in the target and vice versa (the *1:1 constraint*) and that if a correspondence between two statements is included in an interpretation, then so must correspondences between its arguments (the *parallel connectivity* constraint). Which interpretation is chosen is governed by the *systematicity* constraint: preference is given to interpretations that match systems of relations in the base and target.

Structure-mapping theory incorporates computational-level or information-level assumptions about analogical processing, in the sense discussed by Marr (1982). Each of the theoretical constraints is motivated by the role analogy plays in cognitive processing. The 1:1 and parallel connectivity constraints ensure that the candidate

inferences of an interpretation are well-defined. The systematicity constraint reflects a (tacit) preference for inferential power in analogical arguments. Structure-mapping theory provides an account of analogy that is independent of any specific computer implementation. It has broad application to a variety of cognitive tasks involving analogy, as well as to tasks involving ordinary similarity comparisons, including perceptual similarity comparisons (cf. Medin *et al.* 1993, Gentner and Markman 1995, 1997).

In addition to mapping, structure-mapping theory makes claims concerning other processes involved in analogical processing, including retrieval and learning. The relationships between these processes are often surprisingly subtle. Retrieval, for instance, appears to be governed by overall similarity, because this is an ecologically sound strategy for organisms in a world where things that look alike tend to act alike. On the other hand, in learning conceptual material, a high premium is placed on structural consistency and systematicity, since relational overlap provides a better estimate of validity for analogical inferences than the existence of otherwise disconnected correspondences.

As Marr pointed out, eventually a full model of a cognitive process should extend to the algorithm and mechanism levels of description as well. We now describe systems that use structure-mapping theory to model cognitive processes, beginning with SME.

3.1.1. *SME*. The SME simulation takes as input two descriptions, each consisting of a set of propositions. The only assumption we make about statements in these descriptions is that (1) each statement must have an identifiable predicate and (2) there is some means of identifying the roles particular arguments play in a statement. Predicates can be relations, attributes,¹ functions, logical connectives, or modal operators. Representations that have been used with SME include descriptions of stories, fables, plays, qualitative and quantitative descriptions of physical phenomena, mathematical equations, geometric descriptions, visual descriptions, and solutions of problems.

Representation is a crucial issue in our theory, for our assumption is that the results of a comparison process depend crucially on the representations used. We further assume that human perceptual and memorial representations are typically far richer than required for any one task.² Thus we do not assume that the representations given to SME contain all logically possible (or even relevant) information about a situation. Rather, the input descriptions are intended as particular psychological construals—collections of knowledge that someone might bring to bear on a topic in a particular context. The content and form of representations can vary across individuals and contexts. Thus, the colour of a red ball may be encoded as *colour (ball) = red* on some occasions, and as *red (ball)* on others. Each of these construals has different implications about the way this situation will be processed (see Gentner *et al.* 1995) for a more detailed treatment of this issue)

This issue of the size of the construals is important. Chalmers *et al.* (1992, p. 200) argue that the mapping processes used in SME ‘all use very small representations that have the relevant information selected and ready for immediate use’. The issues of the richness and psychological adequacy of the representations, and of the degree to which they are (consciously or unconsciously) pre-tailored to create the desired mapping results, are important issues. But although we agree that more complex representations should be explored than those typically used by ourselves and other researchers—including Hofstadter and his colleagues—we also note three points relevant to this

criticism: (1) SME's representations typically contain irrelevant as well as relevant information, and misleading as well as appropriate matches, so that the winning interpretation is selected from a much larger set of potential matches; (2) in some cases, as described below, SME has been used with very large representations, certainly by comparison with those of Copycat; and (3) on the issue of hand-coding, SME has been used with representations built by other systems for independent purposes. In some experiments the base and target descriptions inputted into the SME program are written by human experimenters. In other experiments and simulations (e.g. PHINEAS, MAGI, MARS) many of the representations are computed by other programs. SME's operation on these descriptions is the same in either case.

Given the base and target descriptions, SME finds globally consistent interpretations via a local-to-global match process. SME begins by proposing correspondences, referred to as *match hypotheses*, in parallel between statements in the base and target. Not every pair of statements can match; structure-mapping theory postulates the *tiered identity* constraint to describe when statements may be aligned. Initially, two statements can be aligned if either (1) their predicates are identical or (2) their predicates are functions, and aligning them would allow a larger relational structure to match. Then, SME filters out match hypotheses that are structurally inconsistent, using the 1:1 and parallel connectivity constraints of structure-mapping theory described in the previous section. Depending on context (including the system's current goals (cf. (Falkenhainer 1990b)), more powerful re-representation techniques may be applied to see if two statements can be aligned in order to achieve a larger match (or a match with potentially relevant candidate inferences).

Mutually consistent collections of match hypotheses are gathered into a small number of global interpretations of the comparison referred to as *mappings*³ or *interpretations*. For each interpretation, *candidate inferences* about the target—that is, statements about the base that are connected to the interpretation but are not yet present in the target—are imported into the target. An evaluation procedure based on Gentner's (1983) systematicity principle is used to compute an evaluation for each interpretation, leading to a preference for deep connected common systems (Forbus and Gentner 1989).

The SME algorithm is very efficient. Even on serial machines, the operations involved in building networks of match hypotheses and filtering can be carried out in polynomial time, and the greedy merge algorithm used for constructing interpretations is linear in the worst case, and generally fares far better empirically. How does SME do at capturing significant aspects of analogical processing? It models the local-to-global nature of the alignment process (see (Goldstone and Medin 1994) for psychological evidence). Its evaluations ordinarily match human soundness judgments. It models the drawing of inferences, an important form of analogical learning. However, the real power of modelling analogical mapping as a separable process can best be seen in the larger simulations that use SME as a component. One of the first of these, and the one that best shows the use of analogy in building representations, is Falkenhainer's PHINEAS.

3.1.2. *PHINEAS: a simulation of analogical learning in physical domains.* The PHINEAS program (Falkenhainer 1987, 1988, 1990a) learns physical theories by analogy with previously understood examples. Its design exploits several modules that have themselves been used in other projects, including SME, QPE (Forbus 1990), an

implementation of qualitative process theory (Forbus 1984), and DATMI (Decoste 1990),⁴ a measurement interpretation system. The architecture of PHINEAS is illustrated in Figure 2.

The best way to illustrate how PHINEAS works is by example. The program starts with the description of the behaviour of a physical system, described in qualitative terms. In one example, PHINEAS is given the description of the temperature changes that occur when a hot brick is immersed in cold water. The program first attempts to understand the described behaviour in terms of its current physical theories, by using QPE to apply these theories to the new situation and qualitatively simulate the kinds of behaviour that can occur, and then uses DATMI to construct explanations of the observations in terms of the simulated possibilities. For the example given PHINEAS did not have a model of heat or heat flow, so it could not find any physical processes to explain the observed changes. In such circumstances PHINEAS turns to analogy to seek an explanation.




Figure 2. The architecture of PHINEAS. In PHINEAS, SME was used as a module in a system that learns qualitative models of physical phenomena via analogy. PHINEAS' map/analyse cycle is a good example of how SME can be used in systems that interleave representation construction with other operations.

To derive an explanation, PHINEAS attempts to find an analogous behaviour in its database of previously-explained examples. These examples are indexed in an abstraction hierarchy by their observed behaviours.⁵ Based on global properties of the new instance's behaviour, PHINEAS selects a potentially analogous example from this hierarchy. When evaluating a potential analogue, PHINEAS uses SME to compare the behaviours, which generates a set of correspondences between different physical aspects of the situations. These correspondences are then used with SME to analogically infer an explanation for the new situation, based on the explanation for the previously understood situation. Returning to the immersed brick example, the most promising candidate explanation is a situation where liquid flow causes two pressures to equilibrate. To adapt this explanation for the original behaviour PHINEAS creates a new process, PROCESS-1 (which we will call heat flow for simplicity after this), which is analogous to the liquid flow process, using the correspondences between aspects of the two behaviours. In this new physical process, the relationships that held for pressure in the liquid flow situation are hypothesized to hold for the corresponding temperature parameters in the new situation.

Generating the initial physical process hypothesis via analogical inference is only the first step. Next PHINEAS must ensure that the hypothesis is specified in enough detail to actually reason with it. For instance, in this case it is not obvious what the analogue to liquid is, nor what constitutes a flow path, in the new heat flow situation. It resolves these questions by a combination of reasoning with background knowledge about the physical world (e.g. that fluid paths are a form of connection, and that immersion in a liquid implies that the immersed object is in contact with the liquid) and by additional analogies. Falkenhainer calls this the *map/analyse cycle*. Candidate inferences are examined to see if they can be justified in terms of background knowledge, which may in turn lead to further matching to see if the newly applied background knowledge can be used to extend the analogy further. Eventually, PHINEAS extends its candidate theory into a form that can be tested, and proceeds to do so using the combination of QPE and DATMI to see if the newly extended theory can explain the original observation.

We believe that PHINEAS provides a model for the use of analogy in learning, and indeed for the role of analogy in abduction tasks more generally. The least psychologically plausible part of PHINEAS's operation is the retrieval component, in which a domain-specific indexing vocabulary is used to filter candidate experience (although it might be a reasonable model of expert retrieval). On the other hand PHINEAS's map/analyse cycle and its method of using analogy in explanation and learning are we believe plausible in their broad features as a psychological model.

The omission of PHINEAS from Chalmers *et al.*'s (1992) discussion of analogy (and from Hofstadter's (1995a) discussions) is striking, since it provides strong evidence against their position.⁶ PHINEAS performs a significant learning task, bringing to bear substantial amounts of domain knowledge in the process. It can extend its knowledge of the physical world, deriving new explanations by analogy, which can be applied beyond the current situation. Thus, PHINEAS provides a solid refutation of the claim of Chalmers *et al.* that systems that interleave a general mapping engine with other independently developed modules cannot be used to flexibly construct their own representations.

3.1.3. *Other simulations using SME.* SME has been used in a variety of other cognitive simulations. These include the following.

- **SQL**: a simulation of abstraction processes in concept learning (Skorstad *et al.* 1988). Here SME was used to explore whether abstraction-based or exemplar-based accounts best accounted for sequence effects in concept learning. The input stimuli were representations of geometric figures.
- **MAC/FAC**: a simulation of similarity-based retrieval (Gentner and Forbus 1991, Law *et al.* 1994, Forbus *et al.* 1995). In MAC/FAC, SME is used in the second stage of retrieval to model the human preference for structural reminders. The first stage is a simple matcher whose output estimates what SME will produce on two structured representations and can be implemented in first-generation connectionist hardware in parallel, and thus has the potential to scale to human-sized memories. MAC/FAC has been tested with simple metaphors, stories, fables, Shakespeare plays,⁷ and descriptions of physical phenomena.
- **MAGI**: a simulation of symmetry detection (Ferguson 1994). MAGI uses SME to map a representation against itself, to uncover symmetries and regularities within a representation. It has been tested with examples from the visual perception literature, conceptual materials,⁸ and combined perceptual/functional representations (i.e. diagrams and functional descriptions of digital logic circuits).
- **MARS**: a simulation of analogical problem solving (Forbus *et al.* 1994). MARS uses SME to import equations from a previously-worked thermodynamics problems⁹ to help it solve new problems. This simulation is the first in a series of systems that we are building to model the range of expert and novice behaviours in problem solving and learning.

The last two systems use a new version of SME, ISME (Forbus *et al.* 1994), which allows incremental extension of the descriptions used as base and target (see (Burstein 1988) and (Keane 1990)).¹⁰ This process greatly extends SME's representation-building capabilities.

3.2. *Psychological research using SME*

SME has been used to simulate and predict the results of psychological experiments on analogical processing. For example, we have used SME to model the developmental shift from focusing on object matches to focusing on relational matches in analogical processing. The results of this simulation indicate that it is possible to explain this shift in terms of a change of knowledge rather than as a change in the basic mapping process itself (Kotovsky and Gentner 1990). Another issue is that of competing mappings, as noted above. SME's operation suggests that when two attractive mappings are possible, the competition between mappings may lead to confusion. This effect has been shown for children (Rattermann and Gentner 1990, Gentner, *et al.* 1995) and to some extent for adults (Markman and Gentner 1993a). A third issue is that SME's structural alignment process for similarity has led to the possibility of a new understanding of *dissimilarity*, based on alignable differences between representations (Markman and Gentner 1993b, Gentner and Markman 1994, Markman and Gentner 1996). In all these cases, SME has been used to verify the representational and processing assumptions underlying the psychological results. These studies suggest many different ways in which analogy may interact with other reasoning processes, including, but not limited to, representation construction.

3.3. Copycat: a model of high-level perception

Copycat operates in a domain of alphabetic strings (see (Chalmers *et al.* 1992, Mitchell 1993, Hofstadter 1995a) for descriptions of Copycat, and (French 1995, Hofstadter 1995a) for descriptions of related programs in different domains). It takes as input problems of the form ‘If the string *abc* is transformed into *abd*, what is the string *aabbcc* transformed into?’ From this input and its built-in rules, Copycat derives a representation of the strings, finds a rule that links the first two strings, and applies that rule to the third string to produce an answer (such as *abbdd*). Copycat’s architecture is a blackboard system (cf Erman *et al.* 1980, Englemore and Morgan 1988), with domain-specific rules¹¹ that perform three tasks: (1) adding to the initial representation, by detecting groups and sequences, (2) suggesting correspondences between different aspects of the representations, and (3) proposing transformation rules to serve as solutions to the problem, based on the outputs of the other rules. As with other blackboard architectures, Copycat’s rules operate (conceptually) in parallel, and probabilistic information is used to control which rules are allowed to fire. Each of these functions is carried out within the same architecture by the same mechanism and their operation is interleaved. Chalmers *et al.* (1992) claim that they are ‘inseparable’.

Concepts in this domain consist of: letters, e.g. *a*, *b*, and *c*; groups, e.g. *aa*, *bb*, and *cc*; and relationships involving ordering—e.g. *successor*, as in *b* is the successor of *a*. A property that both Mitchell (1993) and Chalmers *et al.* (1992) emphasise is that mappings in Copycat can occur between non-identical relationships. Consider for example two strings, *abc* versus *cba*. Copycat can recognize that the first group is a sequence of successors, while the second is a sequence of predecessors. When matching these two strings, Copycat would allow the concepts *successor* and *predecessor* to match, or, in their terminology, to ‘slip’ into each other. Copycat has a pre-determined list of concepts that are allowed to match, called the *Slipnet*. In Copycat, all possible similarities between concepts are determined a priori. The likelihood that a concept will slip in any particular situation is also governed by a parameter called *conceptual depth*. Deep concepts are less likely to slip than shallow ones. The conceptual depth for each concept is, like the links in the Slipnet, hand-selected a priori by the designers of the system.

The control strategy used in Copycat’s blackboard is a form of simulated annealing. The likelihood that concepts will slip into one another is influenced by a global parameter called *computational temperature*, which is initially high but is gradually reduced, creating a gradual settling. This use of temperature differs from simulated annealing in that the current temperature is in part a function of the system’s ‘happiness’ with the current solution. Reaching an impasse may cause the temperature to be reset to a high value, activating rules that remove parts of the old representation, and thus allow new representations to be built.

4. Dimensions of analogy

We see five issues as central to the evaluation of Chalmers *et al.*’s claims with regard to analogical processing as follows.

1. How does perception relate to analogy?
2. How does flexibility arise in analogical processing?
3. Is analogy a domain-general process?
4. How should micro-worlds be used in the study of analogy?
5. How should the psychological plausibility of a model of analogy be assessed?

This section examines these questions, based both on the comparison of SME, PHINEAS, and Copycat above, as well as drawing on the broader computational and psychological literature on analogy

4.1. *How does perception relate to analogy?*

Chalmers *et al.* (1992) argue that, because perception and comparison interact and are mutually dependent, they are inseparable and cannot be productively studied in isolation. But as discussed in Section 2.1, dependencies can arise through interleaving of processes; they need not imply 'in principle' nonseparability. (After all, the respiratory system and the circulatory system are highly mutually dependent, yet studying them as separate but interacting systems has proven extremely useful.) Contrary to Chalmers *et al.*'s claims, even Copycat can be analysed in terms of modules that build representations and other modules that compare representations. Mitchell (1993) provides just such an analysis, cleanly separating those aspects of Copycat that create new representations from those responsible for comparing representations, and showing how these parts interact.

Hofstadter's call for more perception in analogical modelling might lead one to think that he intends to deal with real-world recognition problems. But the high-level perception notion embodied in Copycat is quite abstract. The program does not take as input a visual image, nor line segments, nor even a geometric representation of letters. Rather, like most computational models of analogy, it takes propositional descriptions of the input, which in the case of Copycat consists of three strings of characters, e.g. $abc \rightarrow abd$; $rst \rightarrow ?$ Copycat's domain of operation places additional limits on the length and content of the letter-strings. The perception embodied in Copycat consists of taking this initial sparse propositional description and executing rules that install additional assertions about sequence properties of the English language alphabet. This procedure is clearly a form of representation generation, but (as Chalmers *et al.* (1992) note) falls far short of the complexity of perception.

So far we have considered what the high-level perception approach bundles in with analogical mapping. Let us now consider two things it leaves out. The first is retrieval of analogues from memory. Since Copycat's mapping process is inextricably mixed with its (high-level) perceptual representation-building processes, there is no way to model being reminded and pulling a representation from memory. Yet work on case-based reasoning in artificial intelligence (e.g. Schank 1982, Hammond 1990, Kolodner 1994) and in psychology (e.g. Kahneman and Miller 1986, Holyoak and Koh 1987, Ross 1987, Gentner *et al.* 1993) suggests that previous examples play a central role in the representation and understanding of new situations and in the solution of new problems. To capture the power of analogy in thought, a theory of analogical processing must go beyond analogies between situations that are perceptually present. It must address how people make analogies between a current situation and stored representations of past situations, or even between two or prior situations.

Investigations of analogical retrieval have produced surprising and illuminating results. It has become clear that the kinds of similarity that govern memory access are quite different from the kinds that govern mapping once two cases are present. The pattern of results suggests the fascinating generalization that similarity-based memory access is a stupider, more surface driven, less structurally sensitive process than analogical mapping (Holyoak and Koh 1987, Keane 1988, Gentner *et al.* 1993). In our research we explicitly model the analogical reminding process by adding retrieval processes to SME in a system called MAC/FAC (many are called / but few are chosen)

(Forbus *et al.* 1995). The ARCS model of Thagard *et al.* (1990) represents the corresponding extension to ACME. Thus by decomposing analogical processing into modules, we gain the ability to create accounts which capture both perceptual and conceptual phenomena.

The second omission is learning. Copycat has no way to store an analogical inference, nor to derive an abstract schema that represents the common system (in SME's terms, the interpretation of the analogy, or mapping). For those interested in capturing analogy's central role in learning, such a modelling decision is infelicitous to say the least, although Hofstadter's approach can be defended as a complementary take on the uses of analogy. A central goal in our research with SME is to capture long-term learning via analogy. Three specific mechanisms have been proposed by which domain representations are changed as a result of carrying out an analogy: schema abstraction, inference projection, and re-representation (Gentner *et al.* 1997). The fluid and incremental view of representation embodied in Copycat cannot capture analogy's role in learning.

The holistic view of processing taken by Hofstadter's group obscures the multiplicity of processes that must be modelled to capture analogy in action. This can lead to misunderstandings. In their description of SME, Chalmers *et al.* state that '... the SME program is said to discover an analogy between an atom and the solar system' (Chalmers *et al.*, 1992, p. 196). We do not know who "said" this but it certainly was not said by us. By our account, *discovering* an analogy requires spontaneously retrieving one of the analogues as well as carrying out the mapping.¹² But this attack is instructive, for it underscores Hofstadter's failure to take seriously the distinction between a model of analogical mapping and a model of the full discovery process.

It is worth considering how Falkenhainer's map/analyse cycle (described in Section 3.2.2) could be applied to perceptual tasks. An initial representation of a situation would be constructed, using bottom-up operations on, say, an image. (There is evidence for bottom-up as well as top-down processes in visual perception, e.g. (Marr 1982, Kosslyn 1994). Comparing two objects based on the bottom-up input descriptions leads to the formation of an initial set of correspondences. The candidate inferences drawn from this initial mapping would then provide questions that can be used to drive visual search and the further elaboration of the initial representations. The newly-added information in turn would lead to additional comparisons, continuing the cycle.

Consider the two comparisons in Figure 3 (drawn from Medin *et al.* 1993) as an example. In the comparison between parts (a) and (b) in Figure 3, people who were asked to list the commonalities of these figures said that both have three prongs. In contrast, people who listed the commonalities of the comparison of parts (b) and (c)




Figure 3. An example of how comparison can be used to reduce visual ambiguity. Subjects asked to list the commonalities between A and B said that each has three prongs, while subjects asked to list the commonalities between B and C said that each has four prongs. Since the ambiguous figure is identical in both cases, this demonstrates that similarity processing can be used to resolve visual ambiguities (Medin *et al.* 1993).

in Figure 3 said that both items have four prongs. Thus, the same item was interpreted as having either three or four prongs depending on the object it was compared with. The initial visual processing of the scene would derive information about the contours of the figures, but the detection of the regularities in the portions of the contours that comprise the 'hands' would be conservative, identifying them as bumps, but nothing more. When compared with the three-pronged creature, the hypothesis that the creature with the fourth bump has only three prongs might lead to the clustering of the three bumps of roughly the same size as prongs. When compared with the four-pronged creature, the hypothesis that the creature has four prongs might lead to the dismissal of the size difference as irrelevant. The map analyse cycle allows representation and mapping to interact while maintaining some separation. Recently Ferguson has simulated this kind of processing for reference frame detection with MAGI (Ferguson, 1994). This example suggests that perceptual processing can, in principle, be decomposed into modular sub-tasks. A major advantage of decomposition is identifying what aspects of a task are general-purpose modules, shared across many tasks. The conjectured ability of candidate inferences to make suggestions that can drive a visual search is we believe, a fruitful avenue for future investigation.

4.2. *How does flexibility arise in analogical processing?*

A primary motivation for Hofstadter's casting of analogy as high-level perception is to capture the creativity and flexibility of human cognition. Chalmers *et al.* (1992, p. 201) suggest that this flexibility entails cognitive processes in which 'representations can gradually be built up as the various pressures evoked by a given context manifest themselves'. This is clearly an important issue, worthy of serious consideration. We now examine the sources of flexibility and stability in both Copycat and SME.

We start by noting that comparisons are not infinitely flexible. As described in Section 4.1, people are easily able to view the ambiguous item (Figure 3*b*) as having three prongs when comparing it to Figure 3(*a*) and four prongs when comparing it to Figure 3(*c*). However, people cannot view the item in Figure 3(*b*) as having six prongs, because it has an underlying structure incompatible with that interpretation. There are limits to flexibility.

Another example of flexibility comes from the pair of pictures in Figure 4. In these pictures the robots are *cross-mapped*, that is, they are similar at the object level yet play different roles in the two pictures. People deal flexibly with such cross-mappings. They can match the two pictures either on the basis of like objects, by placing the two robots in correspondence, or on the basis of like relational roles, in which case the robot in the top picture is placed in correspondence with the repairman in the bottom picture. Interestingly, people do not mix these types of similarity (Goldstone *et al.* 1991). Rather, they notice that, in this case, the attribute similarity and the relational similarity are in opposition. SME's way of capturing this flexibility is to allow the creation of more than one interpretation of an analogy. Like human subjects, it will produce both an object-matching interpretation and a relation-matching interpretation. As with human judges, the relational interpretation will usually win out, but may lose to the object interpretation if the object matches are sufficiently rich (Gentner and Rattermann 1991, Markman and Gentner 1993a).

How does Copycat model the flexibility of analogy and the more general principle that cognitive processes are themselves 'fluid'? In Copycat (and in Tabletop) (French 1995)), a major source of flexibility is held to be the ability of concepts to 'slip' into each other, so that non-identical concepts can be seen as similar if that helps make a

good match. Chalmers *et al.* (1992) contrast this property with SME's rule that relational predicates (though not functions and entities) must be identical to match, claiming that Copycat is thus more flexible. Let us compare how Copycat and SME work, to see which scheme really is more flexible.

Like SME, Copycat relies on local rules to hypothesize correspondences between individual statements as part of its mapping operations (Any matcher must constrain the possible correspondences; otherwise everything would match with everything else.) Recall from Section 3.4 that Copycat's constraints come from two sources: a *Slipnet* and a notion of *conceptual depth*. A Slipnet contains links between predicates. For two statements to match, either their predicates must be identical, or there must be a link connecting them in the Slipnet. Each such link has a numerical *weight*, which influences the likelihood that predicates so linked will be placed in correspondence. (Metaphorically, the weight suggests how easy it is for one concept to 'slip into




Figure 4. An example of flexibility in comparison (Markman and Gentner 1990, 1993a).

another'.) These weights are pre-associated with *pairs* of concepts. In addition, each predicate has associated with it a *conceptual depth*, a numerical property indicating how likely it is to be involved in non-identical matches. Predicates with high conceptual depth are less likely to match non-identically than predicates with low conceptual depth.

Both the weights on predicate pairs (the Slipnet) and the conceptual depths of individual predicates are hand-coded and pre-set. Because these representations do not have any other independent motivation for their existence, there are no particular constraints on them, aside from selecting values which make Copycat work in an appealing way. This is not flexibility: it is hand-tailoring of inputs to achieve particular results, in exactly the fashion that Chalmers *et al.* decry. Because of this design, Copycat is unable to make correspondences between classes of statements that are not explicitly foreseen by its designers. Copycat cannot learn, because it cannot modify or extend these hand-coded representations that are essential to its operation. More fundamentally, it cannot capture what is perhaps the most important creative aspect of analogy: the ability to align and map systems of knowledge from different domains.

SME, despite its seeming rigidity, is in important ways more flexible than Copycat. At first glance this may seem wildly implausible. How can a system that requires identity in order to make matches between relational statements qualify as flexible? The relational identity requirement provides a strong, domain-independent, *semantic* constraint. Further, the requirement is not as absolute as it seems, for matches between non-identical functions are allowed, when sanctioned by higher-order structure. Thus SME can place different aspects of complex situations in correspondence when they are represented as functional dimensions. This is a source of *bounded* flexibility. For example, SME would fail to match two scenes represented as louder (Fred, Gina) and bigger (Bruno, Peewee). But if the situations were represented in terms of the same relations over different dimensions—as in greater (loudness(F), loudness(G)) and greater (size(B), size(P))—then the representations can be aligned. Moreover in doing so SME aligns the dimensions of *loudness* and *size*. If we were to extend the comparison—for example, by noting that a megaphone for Gina would correspond to stilts for Peewee—this dimensional alignment would facilitate understanding of the point that both devices would act to equalize their respective dimensions. We have found that online comprehension of metaphorical language is facilitated by consistent dimensional alignments (Gentner and Boronat 1991, Gentner and Imai 1992).

The contrast between SME and Copycat can be illustrated by considering what would happen if both systems were given the following problem with two choices:

If $abc \rightarrow abd$ then *Mercury, Venus, Earth* \rightarrow ??

(1) *Mercury, Venus, Mars* or (2) *Mercury, Venus, Jupiter*

In order to choose the correct answer (1), SME would need representational information about the two domains, e.g. the greater-than relations along the dimension of *closeness to sun* for the planets and the dimension of *precedence in alphabet* for the letters. It could then choose the best relational match, placing the two unlike dimensions in correspondence. But no amount of prior knowledge about the two domains taken separately would equip Copycat to solve this analogy. It would have to have advance knowledge of the cross-dimensional links, e.g. that *closer to sun* could slip into *preceding in alphabet*. The ability of SME to place non-identical

functions in correspondence allows it to capture human ability to see deep analogies between well-understood domains even when they are juxtaposed for the first time.

Despite the above arguments, we agree that there may be times when identity should be relaxed. This consideration has led to our *tiered identity constraint*, which allows non-identical predicates to match (1) if doing so would lead to substantially better or more useful matches, and (2) if there is some principled reason to justify placing those particular predicates in correspondence. One method for justifying non-identical predicate matches is Falkenhainer's *minimal ascension* technique, which was used in PHINEAS (Falkenhainer 1987, 1988, 1990a). Minimal ascension allows statements involving non-identical predicates to match if the predicates share a close common ancestor in a taxonomic hierarchy, when doing so would lead to a better match, especially one that could provide relevant inferences. This is a robust solution for two reasons. First, the need for matching non-identical predicates is determined by the program itself, rather than a priori. Second, taxonomic hierarchies have multiple uses, so that there are sources of external constraint on building them.

However, our preferred technique for achieving flexibility while preserving the identity constraint is to *re-represent* the non-matching predicates into sub-predicates, permitting a partial match. Copycat is doing a simple, domain-specific form of re-representation when alternate descriptions for the same letter-string are computed. However, the idea of re-representation goes far beyond this. If identity is the dominant constraint in matching, then analogizers who have regularized their internal representations (in part through prior re-representation processes) will be able to use analogy better than those who have not. There is some psychological evidence for this *gentrification of knowledge*. Kotovsky and Gentner (1990) found that four-year-olds initially only choose cross-dimensional perceptual matches by chance (e.g. in deciding whether *black-grey-black* should be matched with *big-little-big* or with a foil such as *big-big-little*). But children could come to perceive these matches if they were given intensive within-domain experience or, interestingly, if they were taught words for higher-order perceptual patterns such as symmetry. We speculate that initially children may represent their experience using idiosyncratic internal descriptions (Gentner and Rattermann 1991). With acculturation and language-learning, children come to represent domains in terms of a canonical set of dimensions. This facilitates cross-domain comparisons, which invite further re-representation, further acting to canonicalize the child's knowledge base. Subsequent cross-domain comparisons will then be easier. Gentner *et al.* (1995) discuss some mechanisms of re-representation that may be used by children. Basically, re-representation allows relational identity to arise out of an analogical alignment, rather than acting as a strict constraint on the input descriptions.

A second source of flexibility in SME, again seemingly paradoxically, is its rigid reliance on structural consistency. The reason is that structural consistency allows the generation of candidate inferences. Remember that a candidate inference is a surmise about the target, motivated by the correspondences between the base and the target. To calculate the form of such an inference requires knowing unambiguously what goes with what (provided by satisfying the 1:1 constraint) and that every part of the statements that correspond can be mapped (provided by satisfying the parallel connectivity constraint). This reliance on one-to-one mapping in inference is consistent with the performance of human subjects (Markman, in press). The fact that structural consistency is a domain-general constraint means that SME can (and does)

generate candidate inferences in domains not foreseen by its designers. Copycat, on the other hand, must rely on domain-specific techniques to propose new transformation rules.

A third feature that contributes to flexibility is SME's initially blind local-to-global processing algorithm. Because it begins by blindly matching pairs of statements with identical predicates, and allowing connected systems to emerge from these local identities, it does not need to know the goal of an analogy in advance. Further, it is capable of working simultaneously on two or three different interpretations for the same pair of analogues.

Is SME sufficiently flexible to fully capture human processing? Certainly not yet. But the routes towards increasing its flexibility are open, and are consistent with its basic operation. One route is to increase its set of re-representation techniques, a current research goal. Flexibility to us entails the capability of operating across a wide variety of domains. This ability has been demonstrated by SME. It has been applied to entire domains not foreseen by its designers (as described above), as well as sometimes surprising its designers even in domains they work in. Flexibility also entails the ability to produce different interpretations of the same analogy where appropriate. Consider again the example in Figure 4, which illustrates a typical cross-mapping. As discussed earlier, human subjects entertain two interpretations, one based on object-matching and one based on relational-role matching. SME shows the same pattern, and like people it prefers the interpretation based on like relational roles, so that the robot doing the repairing is placed in correspondence with the person repairing the other robot (see (Markman and Gentner 1993a) for a more detailed description of these simulations). It should be noted that few computational models of analogy are able to handle cross-mappings successfully. Many programs, such as ACME (Holyoak and Thagard 1989), will generate only a single interpretation that is a mixture of the relational similarity match and the object similarity match. The problem cannot even be posed to Copycat, however, because its operation is entirely domain-specific. This, to us, is the ultimate inflexibility.

4.3. *Is analogy a domain-general process?*

A consequence of Chalmers *et al.*'s argument that perception cannot be split from comparison is that one should not be able to make domain-independent theories of analogical processing. However, there is ample evidence to the contrary in the literature. In the genre of theories that are closest to SME, we find a number of simulations that have made fruitful predictions concerning human phenomena, including ACME (Holyoak and Thagard 1989), IAM (Keane 1990, Keane *et al.* 1994), SIAM (Goldstone and Medin 1994), REMIND (Lange and Wharton 1993), and LISA (Hummel and Holyoak 1997).

Even in accounts that are fundamentally different from the present accounts, e.g. bottom-up approaches such as one of Winston's (1975) early models, or top-down approaches (Kedar-Cabelli 1985, Greiner 1988), there are no serious domain-specific models. This is partly because of the problems that seem natural to analogy. The most dramatic and visible role of analogy is as a mechanism for conceptual change, where it allows people to import a set of ideas worked out in one domain into another. Obviously, domain-specific models of analogy cannot capture this signature phenomenon.

There are grave dangers with domain-specific models. The first danger is that the model can be hostage to irrelevant constraints. One way to test the validity of the

inevitable simplifications made in modelling is to triangulate, testing the model with a wide variety of inputs. Limiting a model to a specific domain dramatically reduces the range over which it can be tested. Another way to test the validity of simplifications is to see if they correspond to natural constraints. Surprisingly little effort has been made to examine the psychological plausibility of the simplifying assumptions that go into Copycat. Mitchell (1993) has described an initial experiment designed to see if human subjects perform similarly to Copycat in its domain. This study produced mixed results; more efforts of this kind would be exceedingly valuable. Likewise, French (1995) has presented the results of some studies examining human performance in his Tabletop domain, in which people make correspondences between tableware on a table. Again, this effort is to be applauded. But in addition to carrying out more direct comparisons, the further question needs to be addressed of whether and how these domains generalize to other domains of human experience. At present we have no basis for assuming that the domain specific principles embodied in Copycat are useful beyond a narrow set of circumstances.

The second danger of domain-specific models is that it is harder to analyse the model, to see *why* it works. For example, Mitchell (1993) notes that in Copycat, only one type of relationship may be used to describe a created group. Thus, in grouping the *ttt* in the letter-string *rssttt*, Copycat sometimes describes it as a group of three things, and other times as a group of the letter *t* (to choose, it probabilistically picks one or the other, with shorter strings being more likely to be described by their length than by their common letter). This is partly due to a limitation in the mapping rules for Copycat, which can only create a single matching bond between two objects. For example, it could create either a letter-group bond or a triad group bond between *ttt* and *uuu*, but not both. Why should this be? (Note that this is quite different from the situation with humans. People consider a match between two things better the more structurally consistent relations they have in common.) As far as we can tell the ban on having more than a single mapping bond between any two objects is a simple form of the one-to-one matching criterion found in SME. This prevents one letter from being matched to more than one other, which in most aspects of Copycat's operation is essential, but it backfires in not being able to create matches along multiple dimensions. Human beings, on the other hand, have no problem matching along multiple dimensions. In building domain-specific models the temptation to tweak is harder to resist, because the standard for performance is less difficult than for domain-independent models.

4.4. *Micro-worlds and real worlds: bootstrapping in Lilliput*

A common criticism of Copycat is that its domain of letter-strings is a 'toy' domain, and that nothing useful will come from studying this sliver of reality. Hofstadter and his colleagues counter that the charge of using toy domains is more accurately levelled at other models of analogy (like SME), which leave many aspects of their domains unrepresented. Our purpose here is not to cudgel Copycat with the toy domain label. We agree with Hofstadter that a detailed model of a small domain can be very illuminating. But it is worth examining Hofstadter's two arguments for why SME is more toylike than Copycat.

First, Hofstadter, with some justice, takes SME and ACME to task because of the rather thin domain semantics in some of their representations. For example, he notes that even though SME's representations contain labels such as 'heat' and 'water', 'The only knowledge the program has of the two situations consists of their syntactic

structures... it has no knowledge of any of the concepts involved in the two situations' (Hofstadter 1995a, p. 278). This is a fair complaint for some examples.¹³ However, the same can be said of Copycat's representations. Copycat explicitly factors out every perceptual property of letters, leaving only their identity and sequencing information (i.e. where a letter occurs in a string and where it is in an alphabet). There is no representation of the geometry of letters: Copycat would not notice that 'b' and 'p' are similar under a flip, for instance, or that 'a' looks more like 'a' than 'A' does.

The second argument raised by Hofstadter and his colleagues concerns the size and tailoring of the representations. Although they acknowledge that SME's representations often include information irrelevant to the mapping, (Chalmers *et al.* 1992, p. 201) state:

The mapping processes used in most current computer models of analogy-making, such as SME, all use very small representations that have the relevant information selected and ready for immediate use. For these programs to take as input large representations that include all available information would require a radical change in their design.

Compare the letter-string domain of Copycat with the qualitative physics domain of PHINEAS. There are several ways one might measure the complexity of a domain or problem:

- Domain size: how many facts and rules does it take to express the domain?
- Problem size: how many facts does it take to express the particular situation or problem?
- Elaboration size: how many facts are created when the system understands a particular problem?

In Copycat the domain size is easy to estimate, because we can simply count the number of rules, the number of links in the Slipnet, and the number of predicates. In PHINEAS it is somewhat harder, because much of its inferential power comes from the use of QPE, a qualitative reasoning system that was developed independently and has been used in a variety of other projects and systems. In order to be as fair as possible, we exclude from our count the contents of QPE and the domain-independent laws of QP theory (even though these are part of PHINEAS's domain knowledge), instead, we will count only the number of statements in its particular physical theories. We also ignored the size of PHINEAS's initial knowledge base of explained examples, even though this would again weigh in favour of our claim. Table 1 shows the relative counts on various dimensions.

The number of expressions is only a rough estimate of the complexity of a domain, for several reasons. First, higher-order relations may add more complexity than lower-order relations. Copycat has no higher-order relations, while PHINEAS does. Further, PHINEAS does not have a Slipnet to handle predicate matches. Instead it uses higher-order relational matches to promote matching non-identical predicates. Second, ISA links and partonomy links are not represented in the same way in both systems. Finally, the representation changes significantly enough in Copycat that it is not clear whether to include all relations constructed over the entire representation-building period, or simply to take the maximum size of the representation that Copycat constructs at any one time.

So, in order to estimate the complexity fairly, we use the following heuristics. First, for domain complexity, we count the number of entities, the number of entity categories, the number of rules the domain follows, and the number of relational predicates used. Then, for problem complexity, we simply count the number of entities

Table 1. Relative complexity of Copycat and PHINEAS domain theories

	Copycat	PHINEAS
Entities	26 letters and 5 numbers	10 predefined entities plus arbitrary number of instantiated entities
Entity types	2	13 in type hierarchy
Relational predicates	26	174 (including 50 quantity relations)
Rules	24 rules (codelet types) and 41 slippages between predicates	64 rules. Also 10 views, and 9 physical processes (approximately 135 axioms when expanded into clause form)

and the number of relations. For Copycat, we count the total number of relational expressions created, even when those expressions are later thrown away in favour of other representations.

For the domain comparison (Table 1), the results clearly show the relative complexity of PHINEAS compared with Copycat. Copycat has a set of 31 entities (26 letters and 5 numbers), which are described using a set of 24 codelet rules and 41 slippages,¹⁴ represented in a description language containing only 26 predicates. PHINEAS, on the other hand, has a domain which contains 10 predefined entities (such as alcohol and air) as well as an arbitrary number of instantiations of 13 predefined entity types. There are 64 general rules in the domain theory, as well as multiple rules defined in each of 9 process descriptions and 10 view descriptions, for a total of approximately 112–160 rules (assuming that each process or view description contains an average of 3–5 rules (again, not counting the rules in the QPE rule-engine itself)). The relational language of PHINEAS is much richer than Copycat's, with 174 different predicates defined in its relational language (including 50 quantity types).

The problem complexity of PHINEAS is similarly much higher than Copycat's. For example, take the first examples given for PHINEAS in (Falkenhainer 1988) and for Copycat in (Mitchell 1993). For the IJK problem in Copycat, there are 9 entities that are described via 15 relational expressions¹⁵ (21 if the predicate matches created in the Slipnet are counted). On the other hand, PHINEAS's caloric heat example contains 11 entities (split between base and target) that are described via 88 relational expressions (see Table 2). Similar results may be obtained by comparing other examples from PHINEAS and Copycat.

Despite Chalmers *et al.*'s claims that Copycat excels in representation-building, it seems clear that PHINEAS actually constructs larger and more complex representations.

4.4.1. *The dangers of micro-worlds.* Micro-worlds can have many advantages. But they work best when they allow researchers to focus on a small set of general issues. If chosen poorly, research in micro-worlds can yield results that only apply to a small set of issues specific to that micro-world. The use of Blocks World in 1970s' artificial

Table 2. Relative complexity of Copycat and PHINEAS demonstration problems.

	Copycat (IJK example)	PHINEAS (Caloric heat example)
Entities	9 entities	11 entities (7 in base, 4 in target)
Relations between entities	15 relations ¹⁵	88 relations (55 in base, 33 in target)

intelligence vision research provides an instructive example of the dangers of micro-worlds. First, carving off 'scene analysis' as an independent module that took as input perfect line drawings was, in retrospect, unrealistic: visual perception has top-down as well as bottom-up processing capabilities (cf. recent work in animate vision, e.g. (Ballard 1991)). Second, vision systems that built the presumptions of the micro-world into their very fabric (e.g. all lines will be straight and terminate in well-defined vertices) often could not operate outside their tightly constrained niche. The moral is that the choice of simplifying assumptions is crucial.

Like these 1970s' vision systems, Copycat ignores the possibility of memory influencing current processing and ignores learning. Yet these issues are central to why analogy is interesting as a cognitive phenomenon. Copycat is also highly selective in its use of the properties of its string-rule domain. This extensive use of domain-specific information is also true of siblings of Copycat such as French's Tabletop (French 1995).

If we are correct that the analogy mechanism is a domain-independent cognitive mechanism, then it is important to carry out research in multiple domains to ensure that the results are not hostage to the peculiarities of a particular micro-world.

5. How should the psychological plausibility of a model of analogy be assessed?

Both Hofstadter's group and our own group have the goal of modelling human cognition, but we have taken very different approaches. Our group, and other analogy researchers such as Holyoak, Keane, and Halford, follow a more-or-less standard cognitive science paradigm in which the computational model is developed hand-in-hand with psychological theory and experimentation. The predictions of computational models are tested on people, and the results are used to modify or extend the computational model, or in the case of competing models, to support one model or the other.¹⁶ Further, because we are interested in the *processes* of analogical thinking as well as in the output of the process, we have needed to 'creep up' on the phenomena from several different directions. We have carried out several scores of studies, using a range of methods—free interpretation, reaction time, ratings, protocol analysis, and so on. We are still a long way from a full account.

This research strategy contrasts with that of Hofstadter (1995a, p. 359), who states:

What would make a computer model of analogy-making in a given domain a good model? Most cognitive psychologists have been so well trained that even in their sleep they would come up with the following answer: *Do experiments on a large number of human subjects, collect statistics, and make your program imitate those statistics as closely as possible.* In other words, a good model should act very much like Average Ann and Typical Tom (or even better, like an average of the two of them). Cognitive psychologists tend to be so convinced of this principle as essentially the

only way to validate a computer model that it is almost impossible to talk them out of it. But that is the job to be attempted here.

We note in passing that most cognitive psychologists would be startled to see this characterization. The central goal of most cognitive psychologists is to model the *processes* by which humans think. The job would be many times easier if matching output statistics were all that mattered.

Hofstadter (1995a, p. 354) goes on to propose specific ways in which Copycat and Tabletop might be compared with human processing. For example, answers that seem obvious to people should appear frequently in the program's output, and answers that seem far-fetched to people should appear infrequently in the output; answers that seem elegant but subtle should appear infrequently but with a high quality rating in the program's behaviour. Further, if people's preferred solutions shift as a result of a given order of prior problems,¹⁷ then so should the program's solution frequencies and quality judgments. Also, the program's most frequent pathways to solutions 'should seem plausible from a human point of view'. These criteria seem eminently reasonable from a psychological point of view. But Hofstadter (1995a, p. 364) rejects the psychologist's traditional methods:

Note that these criteria ... can all be assessed informally in discussions with a few people, without any need for extensive psychological experimentation. None of them involves calculating averages or figuring out rank-orderings from questionnaires filled out by large numbers of people.

... such judgments [as the last two above] do not need to be discovered by conducting large studies; once again, they can easily be gotten from casual discussions with a handful of friends.

The trouble with this method of assessment is that it is hard to find out when one is wrong. One salubrious effect of doing experiments on people who do not care about one's hopes and dreams is that one is more or less guaranteed a supply of humbling and sometimes enlightening experiences. Another problem with Hofstadter's method is that no matter how willing the subject, people simply do not have introspective access to all their processes.

In explaining why he rejects traditional psychology methods, Hofstadter (1995a, p. 359) states:

Who would want to spend their time perfecting a model of the performance of *lackluster* intellects when they could be trying to simulate *sparkling* minds? Why not strive to emulate, say, the witty columnist Ellen Goodman or the sharp-as-a-tack theoretical physicist Richard Feynman?

... In domains where there is a vast gulf between the taste of sophisticates and that of novices, it makes no sense to take a bunch of novices, average their various tastes together, and then use the result as a basis for judging the behavior of a computer program meant to simulate a sophisticate.

He notes later that traditional methods are appropriate when one single cognitive mechanism, or perhaps the interaction of a few mechanisms, is probed, because these might reasonably be expected to be roughly universal across minds.

This suggests that some of these differences in method and in modelling style stem from a difference in goals. Whereas psychologists seek to model general mechanisms—and we in particular have made the bet that analogical mapping and comparison in general is one such mechanism—Hofstadter is interested in capturing an extraordinary thinker. We have, of course, taken a keen interest in whether our mechanisms apply to extraordinary individual thinkers. There has been considerable work applying structure-mapping and other general process models to cases of scientific discovery. For example, Nersessian (1992) has examined the use of analogies by Maxwell and Faraday; Gentner *et al.* (1997) have analysed Kepler's writings, and have run SME

simulations to highlight key features of the analogies Kepler used in developing his model of the solar system,¹⁸ and Dunbar (1995) has made detailed observations of the use of analogy in microbiology labs. These analyses of analogy in discovery suggest that many of the processes found in ordinary college students may also occur in great thinkers. But a further difference is that Hofstadter is not concerned with analogy exclusively, but also with its interaction with the other processes of ‘high-level perception’. His aim appears to be to capture the detailed performance of one or a few extraordinary individuals engaged in a particular complex task—one with a strong aesthetic component. This is a unique and highly interesting project. But it is not one that can serve as a general model for the field.

6. Summary and conclusions

We consider the process of arriving at answer *wyz* to be very similar, on an abstract level, to the process whereby a full-scale conceptual revolution takes place in science (Hofstadter 1995, p. 261).

Hofstadter and his colleagues make many strong claims about the nature of analogy, as well as about their research program (as embodied in Copycat), and our own. Our goals here have been to correct mis-statements about our research program and to respond to these claims about the nature of analogy, many of which are not supported or are even countermanded by data. Chalmers *et al.* (1992) argued that analogy should be viewed as ‘high-level perception’. We believe this metaphor obscures more than it clarifies. While it appropriately highlights the importance of building representations in cognition, it undervalues the importance of long-term memory, learning, and even perception, in the usual sense of the word. Finally, we reject Hofstadter’s claim that analogy is inseparable from other processes. On the contrary, the study of analogy as a domain-independent cognitive process that can interact with other processes has led to rapid progress.

There are things to admire about Copycat. It is an interesting model of how representation construction and comparison can be interwoven in a simple, highly familiar domain, in which allowable correspondences might be known in advance. Copycat’s search technique, with gradually lowering temperature, is an intriguing way of capturing the sense of settling on a scene interpretation. Moreover there are some points of agreement: both groups agree on the importance of dimensions such as the clarity of the mapping, and that comparison between two things can alter the way in which one or both are conceived. But Copycat’s limitations must also be acknowledged. The most striking of these is that every potential non-identical correspondence—and its evaluation score—is domain-specific and hand-coded by its designers, forever barring the creative use of analogy for cross-domain mappings or for transferring knowledge from a familiar domain to a new one. In contrast, SME’s domain-general alignment and mapping mechanism can operate on representations from different domains and find whatever common relational structure they share. It has been used with a variety of representations (some built by hand, some built by others, some built by other programs) and has run on dozens if not hundreds of analogies whose juxtaposition was not foreseen by its designers. (True, its success depends on having at least some common representational elements, but this we argue is true of human analogizers as well.) Further, Copycat itself contradicts Chalmers *et al.*’s (1992) claims concerning the holistic nature of high-level perception and analogy, for Mitchell’s (1993) analysis of Copycat demonstrates that it can be analysed into modules.

Debates between research groups have been a motivating force in the advances made in the study of analogy. For example, the roles of structural and pragmatic factors in analogy are better understood as a result of debates in the literature (see Holyoak 1985, Gentner and Clement 1988, Gentner 1991, Keane *et al.* 1994, Spellman and Holyoak 1997, Markman in press). However, these debates first require accurate characterizations of the positions and results on both sides of the debate. It is in this spirit that we sought to correct systematic errors in the descriptions of our work that appear in (Chalmers *et al.* 1992) and again in (Hofstadter 1995a), e.g. the claim that SME is limited to small representations that contain only the relevant information. As Section 3 points out, SME has been used with hand-generated representations, with representations generated for other analogy systems, and with representations generated by other kinds of models altogether (such as qualitative reasoners). SME has been used in combination with other modules in a variety of cognitive simulations and performance programs. In other words, SME is an existence proof that modelling alignment and mapping as domain-general processes can succeed, and can drive the success of other models. Although Chalmers *et al.* never mention our psychological work (which shares an equal role with the simulation side of our research), we believe that it too says a great deal about analogy and its interactions with analogy with other cognitive processes. In our view the evidence is overwhelmingly in favour of SME and its associated simulations over Copycat as a model of human analogical processing.

Acknowledgements

This work was supported by the Cognitive Science Division of the office of Naval Research Grant N00014-89-J1272. We thank Brian Bowdle, Jon Handler, Laura Kotovsky, Mary Jo Rattermann, Phil Wolff, and Eric Dietrich, as well as the Similarity and Analogy group and the SME group for helpful discussions on this topic. Special thanks are due to Kendall Gill for taking over at small forward.

Notes

1. Attributes are unary predicates representing properties of their argument which in the current description are not further decomposed. Examples include Red (ball32) and Heavy (sun).
2. See, for example, the discussion of the *specificity conjecture* in (Forbus and Gentner 1989).
3. Using a greedy merge algorithm, as described in (Forbus and Oblinger 1990), and extended in (Forbus *et al.* 1994). Hofstadter appears to be unaware of the use of this algorithm, '... certainly, the exhaustive search SME performs through all consistent mappings is psychologically implausible' (Hofstadter 1995a, p. 283).
4. Another system, TPLAN (Hogge 1987), a temporal planner, was used in some PHINEAS simulations for designing experiments.
5. Examples of behavioural classifications include *dual-approach* (e.g. two parameters approaching each other) and *cyclic* (e.g. parameters that cycle through a set of values). The abstraction hierarchy is a plausible model of expert memory, but we believe our more recent MAC/FAC model would provide a more psychologically plausible model for most situations.
6. In this connection, we must correct an inaccuracy. In Hofstadter's (1995) reprint of (Chalmers *et al.* 1992), a disclaimer is added on page 185: 'Since this article was written, Ken Forbus, one of the authors of SME, has worked on modules that build representations in "qualitative physics." Some work has also been done on using these representations as input to SME.' However, the use of these representations, and PHINEAS, was discussed in the (Falkenhainer *et al.* 1989) paper cited by Chalmers *et al.* (1992).
7. The representation of fables and plays were supplied by Paul Thagard.
8. This includes its namesake example, a representation of O. Henry's '*The Gift of the Magi*'.
9. Representations for the previously-worked problems are automatically generated by CyclePad (Forbus and Whalley 1994), an intelligent learning environment designed to help students learn engineering thermodynamics. CyclePad is currently being used in education experiments by students at Northwestern University, Emiston, IL and the US Naval Academy.

10. MAGI and MARS appeared after the Chalmers *et al.* (1992) paper, so while they constitute evidence for the utility of modular accounts of analogy, we cannot fault Chalmers *et al.* for not citing them (although this does not apply to SEQL, MAC/FAC, and PHINEAS). However, many of the main claims in the paper by Chalmers *et al.* are repeated in later books by French (1995) and by Hofstadter (1995a) despite the availability of counter-evidence.
11. These rules are called ‘codelets’ in papers describing Copycat.
12. A similar comment occurs in Hofstadter’s (1995) discussion of the ‘Socrates is the midwife of ideas’ analogy analysed by Kittay (1987) as simulated in Holyoak and Thagard’s ACME: ‘At this point, the tiny, inert predicate calculus cores are conflated with the original full-blown situations, subtly leading many intelligent people to such happy conclusions as that the program has insightfully leaped to a cross-domain analogy ...’. Here too, the simulation was presented only as a model of mapping, not the full process of discovery.
13. However, SME escapes this charge for the representations it has borrowed from qualitative physics programs, which have a richly interconnected domain structure. (There is still, of course, no true external reference, but this is equally true for all the models under discussion.) See also Ferguson (1994), which uses visual representations computed automatically from a drawing program.
14. Some of the codelets and most of the slipnodes are really used for mapping, rather than representation-building, so we are actually overcounting the number of relevant rules here.
15. The 15 relations for the IJK example include three each of the *leftmost*, *rightmost*, and *middle* relations, two grouping relations, and four letter-successor relations.
16. Examples are the comparison of MAC/FAC and ARCS as models of similarity-based retrieval (Forbus *et al.* 1995), the comparison of SME and ACME as accounts of analogical inference (Clement and Gentner, 1991, Spellman and Holyoak 1993, Markman in press), and comparisons of ACME, SME, and IAM (Keane *et al.* 1994).
17. Burns (1996) has shown that such order effects do occur: people’s preferred solutions on letter-string analogies shift as a result of prior letter-string analogies.
18. We hasten to state that we do not consider ourselves to have captured Kepler’s discovery process.

References

- Ballard, D. H. (1991) Animate vision. *Artificial Intelligence*, **48**: 57–86.
- Boden, M. A. (1991) *The Creative Mind: Myths and Mechanisms* (New York: Basic Books).
- Burns, B. B. (1996) Meta-analogical transfer: transfer between episodes of analogical reasoning. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **22**(4): 1032–1048.
- Burstein, M. H. (1988) Incremental learning from multiple analogies. In A. Prieditis (eds.), *Analogica* (Los Altos, CA: Morgan Kaufmann) pp. 37–62.
- Chalmers, D. J., French, R. M. and Hofstadter, D. R. (1992) High-level perception, representation and analogy: a critique of artificial intelligence methodology. *Journal of Experimental & Theoretical Artificial Intelligence*, **4**: 185–211.
- Clement, C. A. and Gentner, D. (1991) Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, **15**: 89–132.
- Decoste, D. (1990) Dynamic across-time measurement interpretation. In *Proceedings of the Ninth National Conference on Artificial Intelligence*.
- Dunbar, K. (1995) How scientists really reason: scientific reasoning in real-world laboratories. In R. J. Sternberg and J. E. Davidson (eds.) *The nature of insight* (Cambridge, MA: MIT Press), pp. 365–396.
- Englemore, B. and Morgan, T. (1988) *Blackboard Systems* (Cambridge, MA: MIT Press).
- Erman, L. D., Hayes-Roth, F., Lesser, V. R. and Reddy, D. R. (1980) The Hearsay II speech understanding system: integrating knowledge to resolve uncertainty. *Computing Surveys*, **12**(2) 213–253.
- Falkenhainer, B. (1987) An examination of the third stage in the analogy process: verification-based analogical learning. In *Proceedings of IJCAI-87*, pp. 260–263.
- Falkenhainer, B. (1988) Learning from physical analogies: a study in analogy and the explanation process. PhD thesis, University of Illinois at Urbana-Champaign.
- Falkenhainer, B. (1990a) A unified approach to explanation and theory formation. In Shrager and Langley (eds.) *Computational Models of Scientific Discovery and Theory Formation* (San Mateo, CA: Morgan Kaufmann). Also in Shavlik and Dietterich (eds.) (1990) *Readings in Machine Learning* (San Mateo, CA: Morgan Kaufmann).
- Falkenhainer, B. (1990b) Analogical interpretation in context. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (Cambridge, MA: Erlbaum).
- Falkenhainer, B., Forbus, K. D. and Gentner, D. (1986) The structure-mapping engine. In *Proceedings of the Fifth National Conference on Artificial Intelligence*.
- Falkenhainer, B., Forbus, K. D. and Gentner, D. (1989) The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, **41**(1): 1–63.

- Ferguson, R. W. (1994) MAGI: Analogy-based encoding using symmetry and regularity. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum).
- Forbus, K. D. (1984) Qualitative process theory. *Artificial Intelligence*, **24**(1): 85–168.
- Forbus, K. D. (1990). The qualitative process engine. In D. S. Weld and J. de Kleer (eds.) *Readings in Qualitative Reasoning about Physical Systems* (San Mateo, CA: Morgan Kaufmann) pp. 220–235.
- Forbus, K. D., Ferguson, R. W., and Gentner, D. (1994) Incremental structure-mapping. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum), pp. 313–318.
- Forbus, K. D. and Gentner, D. (1989) Structural evaluation of analogies: what counts? In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (Ann Arbor, MI: Erlbaum) pp. 341–348.
- Forbus, K. D., Gentner, D. and Law, D. (1995) MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, **19**(2): 141–205.
- Forbus, K. and Whalley, P. (1994) Using qualitative physics to build articulate software for thermodynamics education. In *Proceedings of AAAI-94*, Seattle, Washington, USA.
- French, R. M. (1995) *The Subtlety of Similarity* (Cambridge, MA: MIT Press).
- Gentner, D. (1983) Structure-mapping: a theoretical framework for analogy. *Cognitive Science*, **23**: 155–170.
- Gentner, D. (1989) The mechanisms of analogical learning. In S. Vosniadou and A. Ortony (eds.) *Similarity and Analogical Reasoning* (New York: Cambridge University Press) 199–241.
- Gentner, D. and Boronot, C. B. (1991) Metaphors are (sometimes) processed as generative domain-mappings. Paper presented at the symposium on Metaphor and Conceptual Change, Meeting of the Cognitive Science Society, Chicago, IL, USA.
- Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P. and Forbus, K. D. (1997) Conceptual change via analogical reasoning: a case study of Johannes Kepler. *Journal of the Learning Sciences*, **6**(1): 3–40.
- Gentner, D. and Forbus, K. D. (1991) MAC/FAC: A model of similarity-based retrieval. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum).
- Gentner, D. and Holyoak, K. J. (1997) Reasoning and learning by analogy. *American Psychologist*, **52**: 32–34.
- Gentner, D. and Imai, M. (1992) Is the future always ahead? Evidence for system-mappings in understanding space-time metaphors. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (Bloomington, IN: Erlbaum).
- Gentner, D. and Markman, A. B. (1994) Structural alignment in comparison: no difference without similarity. *Psychological Science*, **5**(3): 152–158.
- Gentner, D. and Markman, A. B. (1995) Similarity is like analogy. In C. Caciari (ed.) *Similarity*. (Brussels: BREPOLs) pp. 111–148.
- Gentner, D. and Markman, A. B. (1997) Structural alignment in analogy and similarity. *American Psychologist*, **52**(1): 45–56.
- Gentner, D. and Rattermann, M. J. (1991) Language and the career of similarity. In S. A. Gelman and J. P. Byrnes (eds.) *Perspectives on Language and Thought: Interrelations in Development* (Cambridge: Cambridge University Press), pp. 225–277.
- Gentner, D., Rattermann, M. J. and Forbus, K. D. (1993) The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology*, **25**(4), 524–575.
- Gentner, D., Rattermann, M. J., Markman, A. B. and Kotovsky, L. (1995) Two forces in the development of relational structure. In T. Simon and G. Halford (eds.) *Developing Cognitive Competence: New Approaches to Process Modeling* (Hillsdale, NJ: Erlbaum).
- Goldstone, R. L. and Medin, D. L. (1994) Similarity, interactive-activation and mapping. In K. J. Holyoak and J. A. Barnden (eds.) *Advances in Connectionist and Neural Computation Theory*, Vol. 2, *Analogical Connections* (Norwood, NJ: Ablex).
- Goldstone, R. L., Medin, D. L. and Gentner, D. (1991) Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology*, **23**: 222–262.
- Greiner, R. (1988) Learning by understanding analogies. *Artificial Intelligence*, **35**: 81–125.
- Halford, G. S. (1992) Analogical reasoning and conceptual complexity in cognitive development. *Human Development* **35**(4), 193–217.
- Hammond, K. J. (1990) Explaining and repairing plans that fail. *Artificial Intelligence*, **45**: 173–228.
- Hofstadter, D. H. (1995a) *Fluid Concepts and Creative Analogies* (New York: Basic Books).
- Hofstadter, D. H. (1995b) A review of mental leaps: analogy in creative thought. *A. I. Magazine*, **16**: 75–80.
- Holyoak, K. J. and Koh, K. (1987) Surface and structural similarity in analogical transfer. *Memory and Cognition*, **15**(4): 332–340.
- Holyoak, K. J. and Thagard, P. (1989) Analogical mapping by constraint satisfaction. *Cognitive Science*, **13**(3): 295–355.
- Holyoak, K. J. and Thagard, P. (1995) *Mental Leaps: Analogy in Creative Thought* (Cambridge, MA: MIT Press).

- Hummel, J. E. and Holyoak, K. J. (1997) Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, **104**(3), 427–466.
- Kahneman, D. and Miller, D. T. (1986) Norm theory: comparing reality to its alternatives. *Psychological Review*, **93**(2): 136–153.
- Keane, M. T. G. (1990) Incremental analogizing: theory and model. In K. J. Gilhooly, M. T. G. Deane, R. H. Logie and G. Erdos (eds.) *Lines of Thinking* (London: Wiley) pp. 221–235.
- Keane, M. T., Ledgeway, T. and Duff, S. (1994) Constraints on analogical mapping: a comparison of three models. *Cognitive Science*, **18**, 387–438.
- Kedar-Cabelli, S. (1985) Toward a computational model of purpose-directed analogy. In A. Prieditis (eds.) *Analogica*. (San Mateo, CA: Morgan Kaufmann) pp. 89–108.
- Kittay, E. (1987) *Metaphor: Its Cognitive Force and Linguistic Structure* (Oxford: Clarendon).
- Kolodner, J. L. (1994) *Case-based Reasoning* (San Mateo, CA: Morgan Kaufmann).
- Kosslyn, S. (1994) *Image and Brain* (Cambridge, MA: MIT Press).
- Kotovsky, L. and Gentner, D. (1990). Pack light: you will go farther. In J. Dinsmore and T. Koschmann (eds.) *Proceedings of the Second Midwest Artificial Intelligence and Cognitive Science Society Conference*, Carbondale, IL, pp. 60–67.
- Lange, T. E. and Wharton, C. M. (1993) Dynamic memories: analysis of an integrated comprehension and episodic memory retrieval model. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum).
- Law, B. K., Forbus, K. D. and Gentner, D. (1994) Simulating similarity-based retrieval: a comparison of ARCS and MAC/FAC. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum).
- Markman, A. B. (in press) Constraints on analogical inference. *Cognitive Science*.
- Markman, A. B. and Gentner, D. (1990) Analogical mapping during similarity judgements. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Erlbaum).
- Markman, A. B. and Gentner, D. (1993a) Structural alignment during similarity comparisons. *Cognitive Psychology*, **25**(4): 431–467.
- Markman, A. B. and Gentner, D. (1993b) Splitting the differences: a structural alignment view of similarity. *Journal of Memory and Language*, **32**(4): 517–535.
- Markman, A. B. and Gentner, D. (1996) Commonalities and differences in similarity comparisons. *Memory and Cognition*, **24**(2): 235–249.
- Marr, D. (1982) *Vision* (New York: W. H. Freeman and Company).
- Medin, D. L., Goldstone, R. L. and Gentner, D. (1993) Respects for similarity. *Psychological Review*, **100**(2): 254–278.
- Mitchell, M. (1993) *Analogy-making as Perception: A Computer Model* (Cambridge, MA: MIT Press).
- Morrison, C. and Dietrich, E. (1995) Structure-mapping vs. high-level perception: the mistaken fight over the explanation of analogy. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (Pittsburgh, PA: Erlbaum), pp. 678–682.
- Nersessian, N. J. (1992) How do scientists think? Capturing the dynamics of conceptual change in science. In R. Giere (ed.) *Cognitive Models of Science* (Minneapolis, MI: University of Minnesota Press) pp. 3–44.
- Rattermann, M. J. and Gentner, D. (1990) The development of similarity use: it's what you know, not how you know it. In *Proceedings of the Second Midwest Artificial Intelligence and Cognitive Science Society Conference* (Carbondale, IL) pp. 54–59.
- Rattermann, M. J. and Gentner, D. (1991) Language and the career of similarity. In S. A. Gelman and J. P. Byrnes (eds.) *Perspectives on Language and Thought: Interrelations in Development* (London: Cambridge University Press).
- 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**(4): 629–639.
- Schank, R. C. (1982) *Dynamic Memory: A Theory of Learning in Computers and People* (Cambridge: Cambridge University Press).
- Skorstad, J., Gentner, D. and Medin, D. (1988) Abstraction processes during concept learning: a structural view. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (Montreal: Erlbaum).
- Spellman, B. A. and Holyoak, K. J. (1993) An inhibitory mechanism for goal-directed analogical mapping. In *The Fifteenth Annual Meeting of the Cognitive Science Society*, (pp. 947–952). Boulder, CO: Lawrence Erlbaum Associates.
- Spellman, B. A. and Holyoak, K. J. (1996) Pragmatics in analogical mapping. *Cognitive Psychology*, **31**: 307–346.
- Thagard, P., Holyoak, K. J., Nelson, G. and Gochfeld, D. (1990) Analog retrieval by constraint satisfaction. *Artificial Intelligence*, **46**: 259–310.
- Waldrop, M. (1987) Causality, structure, and common sense. *Science*, **237**: 1297–1299.
- Winston, P. H. (1975) Learning structural descriptions from examples. In P. H. Winston (ed.) *The Psychology of Computer Vision* (New York, NY: McGraw-Hill).