



Computational models of analogy

Dedre Gentner^{1*} and Kenneth D. Forbus²

Analogy is a core process in human cognition. A number of valuable computational models of analogy have been created, capturing aspects of how people compare representations, retrieve potential analogs from memory, and learn from the results. In the past 25 years, this area has progressed rapidly, fueled by strong collaboration between psychologists and Artificial Intelligence (AI) scientists, with contributions from linguists and philosophers as well. There is now considerable consensus regarding the constraints governing the mapping process. However, computational models still differ in their focus, with some aimed at capturing the range of analogical phenomena at the cognitive level and others aimed at modeling how analogical processes might be implemented in neural systems. Some recent work has focused on modeling interactions between analogy and other processes, and on modeling analogy as a part of larger cognitive systems.

© 2010 John Wiley & Sons, Ltd. *WIREs Cogn Sci*

ANALOGICAL COMPARISON AND ITS ROLE IN COGNITION

Analogy has been studied from a variety of perspectives. The computational modeling of analogy, conducted in collaboration between psychologists and AI scientists, has provided a valuable source of insights which have led to a deeper theoretical understanding of analogy and the roles it plays in human cognition. This article summarizes some of these insights, as well as some specific computational models of analogy.

Analogy involves the comparison of two structured representations. That is, the representations being compared typically include labeled relationships between entities and between other relations. Such representations contrast sharply with representations lacking internal structure, such as those based on independent features or multidimensional vectors (see Ref 1). This representational choice is dictated by a large set of findings indicating that people are sensitive to relational structure in processing analogy,²⁻⁴ and even in visual comparisons.⁵⁻⁸ Computational models provide insights as to why this must be. One well-known characteristic of analogy is that it can

suggest new inferences; indeed, the most familiar type of analogy is one in which a familiar *base* (or *source*) domain is mapped to a less familiar (or more abstract) *target* domain, with the result that a new prediction or explanation is mapped from the base to the target. Importantly, this kind of inference is selective: not everything known about the base domain is mapped to the target. Computationally, this kind of selective inference can be captured if we assume (a) that people have structured representations in which higher order relations constrain the lower-order relations and (b) that analogical mapping operates to prefer matching systems of relations governed by higher order constraining relations such as *cause* or *implies* rather than isolated matches (Gentner's *Systematicity* principle⁹). Psychological studies bear out this assumption. For example, when people were given analogous scenarios designed so that two pertinent facts were present in the base but not the target, and asked to make a new inference about the target, they inferred whichever fact was connected (in the base) via a higher order causal relation to another matching fact.² In other words, they did not simply bring across *any* fact present in the base but not the target; their inferences were implicitly geared toward finding a larger matching system.

As this example suggests, causal relations often serve as higher order relations in analogical processing. When the antecedents of a causal relation are matched, the consequent is projected to hold in the new (target) situation (prediction); and if instead the consequents are matched, the antecedents are projected to hold in the new situation (explanation or

*Correspondence to: gentner@northwestern.edu

¹Department of Psychology, Weinberg College of Arts and Sciences, Swift Hall 102, 2029 Sheridan Road, Evanston, IL 60208-2710, USA

²EECS, Northwestern University, Ford 3-320, 2133 Sheridan Road, Evanston, IL 60208, USA

DOI: 10.1002/wcs.105

abduction). But other kinds of higher order relations can also serve to constrain analogical inference, including logical and mathematical relations, such as implication, and perceptual regularities, such as symmetry or monotonicity. To capture the processing of causal theories, explanations, logical proofs, and other such inferential systems requires structured representations; they cannot be effectively represented via mental distance models or feature set models (for further discussion, see Ref 1).

In the early days of analogical research, analogical processing was viewed as a rarified mental operation, occurring only infrequently. Today the emerging consensus among analogy researchers is quite different. While spontaneous analogies between dramatically different domains are indeed rare, the same kinds of comparison processes that make distant analogies possible also appear to underlie the more mundane, everyday similarity comparisons we make, including perceptual comparisons.^{5–8,10} These within-domain analogical comparisons might even explain behavior commonly attributed to rule-following.^{11–13} The psychological evidence pointing to the same comparison processes underlying a wide range of cognitive phenomena has motivated explorations of larger-scale models of the roles in analogy in cognitive processing, a new frontier for analogy research.

COMPUTATIONAL MODELS OF ANALOGICAL PROCESSES

Analogy is generally decomposed into multiple subprocesses, as follows:

1. *Retrieval*: Given a situation, find an analog that is similar to it.
2. *Mapping*: Given two situations, align them structurally to produce a set of *correspondences* that indicate ‘what goes with what,’ *candidate inferences* that follow from the analogy, and a *structural evaluation score* which provides a numerical measure of how well the base and target align.
3. *Abstraction*: The results of comparison may be stored as an abstraction, producing a schema or other rule-like structure.
4. *Rerepresentation*: Given a partial match, people may alter one or both analogs to improve the match.

Finally, we note some other processes that, although not specific to analogy, are nonetheless important to it. First, psychological evidence suggests

that *encoding* has a large effect on analogical processing. How two situations are encoded strongly influences whether one will retrieve the other from LTM, as well as whether they will yield a good alignment when they are compared. Thus, how situations are encoded is of great importance to analogical processing. Second, in addition to the structural evaluation that is specific to analogy, the results of an analogy often receive a more general evaluation: e.g., are the inferences factually true (or at least plausible) and are they relevant to the current context.

This functional decomposition fits with psychological evidence that different subprocesses have different characteristics. For example, mapping is known to be sensitive to structural overlap, while retrieval is dominated by surface overlap.^{14–16} Further, not all comparisons lead to the formation of schema, so generalization may or may not occur from a specific analogy. Finally, some models integrate two or more of these operations into a single process, while others use separate process models for each functional process.

We now discuss these subprocesses and how different models try to capture them. (See Table 1 for a list of the models we discuss along with their chief characteristic features.) We begin with mapping—the one subprocess that all the models aim to capture. Mapping is the core defining process for analogy. One might be given both analogs, thereby eliminating retrieval; one might or might not need to rerepresent, or to draw an abstraction from the analogy; but without mapping, there is no analogy. Indeed, most of the models we discuss do not attempt to capture retrieval from memory, nor abstraction from multiple exemplars. Since all the models include some version of analogical mapping, this is our logical starting point. After discussing approaches to mapping, we go on to examine the other subprocesses in sequence.

Mapping

The mapping process takes as input two structured representations, the *base* (sometimes called *source*) and *target* and computes one or more *mappings*. Each mapping consists of a set of *correspondences*, each linking a particular item (entity or statement) in the base with a particular item (entity or statement) in the target. It can also contain *candidate inferences*, which are surmises about what is true in one description based on projecting structure from the other, as discussed above. Typically a mapping also includes a numerical score, indicating its structural quality.

Structure-mapping theory^{9,17–20} proposes that the following constraints govern the mapping process:

TABLE 1 | Computational Models of Analogy and Their Key Characteristics

Name	Processes	Type	General?	Key feature
ACME	Mapping	Connectionist	Yes	Network used for multiple constraint satisfaction
AMBR	Mapping	Hybrid	Yes	Based on distributed micro-agent framework
ARCS	Retrieval, mapping	Connectionist	Yes	Parallel first-stage matches potential analogs; ACME used as second-stage matcher
CAB	Mapping	Connectionist	Yes	Uses middle-out algorithm plus parallel constraint satisfaction
CARL	Mapping	Symbolic	No	Understanding analogies for programming, first incremental matcher
Copycat	Encoding, mapping	Hybrid	No	Letter-string analogies, using rules governed by simulated annealing for encoding
DORA	Retrieval, Mapping	Connectionist		Models early relation-learning as combining of role-relations
DUAL	Encoding, Retrieval, Mapping	Hybrid	Yes	Uses AMBR for mapping, same distributed agent framework for retrieval and encoding
EMMA	Retrieval, Mapping	Hybrid	No	Used Latent Semantic Analysis to model predicate similarity
HDTP	Mapping	Symbolic	Yes	Uses antiunification to construct generalization
IAM	Mapping	Symbolic	Yes	First general-purpose incremental matcher
LISA	Retrieval, Mapping	Structured connectionist	Yes	Uses microfeatures and projection-based algorithm neurally inspired
MAC/FAC	Retrieval	Symbolic	Yes	Parallel first-stage vector match to filter candidates; SME used as stage 2 matcher
NLAG	Mapping	Symbolic	Yes	Top-down algorithm
SQL	Generalization	Symbolic	Yes	Uses SME to compare exemplars, produces probabilistic generalizations
SME	Mapping	Symbolic	Yes	Middle-out: parallel initial stage followed by structurally consistent kernels & greedy merge algorithm
Tabletop	Encoding, mapping	Hybrid	No	Place settings, using rules governed by simulated annealing for encoding
Winston	Mapping	Symbolic	Yes	Early bottom-up algorithm; later, importance-dominated matching

- *Structural consistency*: Structural consistency is defined by two constraints:
 - *1:1 constraint*: Each item in the base maps to at most one item in the target, and vice-versa.
 - *Parallel connectivity*: If a correspondence between two statements is included in a mapping, then so must correspondences between its arguments.
- *Systematicity*: Mappings that place systems of relations—especially those governed by higher order constraining relations—into correspondence are preferred.
- *Tiered identicality*: Identical matches between predicates and functions are preferred. By default, relations must match identically, but

non-identical functions can be aligned if such alignments would support a larger overlapping structure. Depending on task demands, this can be relaxed further to allow non-identical relations to correspond, if they are suggested by a larger structure and satisfy additional criteria.

Viewed computationally, structural consistency ensures that candidate inferences can be projected consistently: Without these constraints, it is unclear what substitutions should be made when projecting inferences.^{21,22} Although there are cases in which people appear to compute correspondences that violate the 1:1 constraint,²³ evidence from inference patterns indicates that people are shifting between

different mappings for the analogy. Within each mapping, the 1:1 constraint is respected, and inferences are made only on the basis of a structurally consistent mapping.^{19,24,25} Most current models of analogical processing incorporate the structural consistency constraint.

The systematicity constraint increases the likelihood that the winning interpretation of an analogy will be an interconnected system of relations, rather than a large set of coincidental matches. Systematicity pushes the mapping process toward producing candidate inferences, because interconnected systems often contain further inferences that can be projected from the base to the target. Among current models (as described below), structure-mapping engine (SME), Analogical Constraint Mapping Engine (ACME) and connectionist analogy builder (CAB) incorporate the systematicity constraint. In other models (e.g., Incremental Analogical Mapper (IAM), Learning and Inference with Schemas and Analogies (LISA)) a similar constraint of preferring highly connected base structure to import may be used.

The tiered identity constraint addresses a key problem that faces any model of mapping: What possible correspondences between items in base and target should be considered? The answer chosen to this question is one of the two factors that determine the computational complexity of the analogical mapping model. If semantic constraints on the matches are not considered, then the matching problem becomes an example of the general graph isomorphism problem, which is known to be NP-complete—that is, it is unlikely that any algorithm can solve it exactly in polynomial time. An algorithm operates in polynomial time if its consumption of a resource (like time or number of processing units) rises in a way that is bounded by a polynomial in the size of some property of its inputs. An NP-complete method such as a pure graph-matching model seems implausible, given the ubiquity and fluency of analogical comparison in human cognition; it would be implausibly slow, if done serially, or require too much hardware, if done in parallel. Further, psychological studies that have pitted syntactic matches (pure graph matches) against semantic matches have found that people attend almost exclusively to the semantic matches in processing analogies.¹⁹ Therefore, many analogy models impose semantic restrictions on which items in the base and target can match, thereby reducing the complexity of matching. Requiring that predicates be identical, or at least very similar, is a strong semantic constraint that rules out the vast majority of possible matches. Such a restriction brings the number of potential correspondences down to $O(N^2)$, where

N is the size of the base and target descriptions—a more psychologically plausible degree of complexity.^a Even so, whatever test is used to determine the closeness of each pair of items must be cheap; allowing arbitrary inference for each decision would be prohibitive as the default operation within a core cognitive process. Structure-mapping's tiered identity constraint basically starts by allowing only statements with identical predicates to match, and allowing other matches only when they would create a larger structure. Consider, for example

```
B: (implies
    (connectedAtContact A B)
    (movesWith A B))
T: (implies
    (rotationallyConnectedTo C D)
    (movesWith C D))
```

Falkenhainer's minimal ascension criterion²⁶ allows non-identical substitutions for predicates playing corresponding roles in a larger structure *if* they have a close common superordinate in the predicate hierarchy. Here, `connectedAtContact` and `rotationallyConnectedTo` have a common superordinate relationship, `connectedTo`, and allowing those statements to match would allow the implication to match (otherwise, parallel connectivity would be violated).

A number of other solutions have been proposed for testing the closeness of relations. For example, CAB²⁷ uses a scheme that is similar to minimal ascension, allowing matches between relations that are close, with weightings inversely proportional to relational distance. In CAB, this is done for all item matches (instead of being restricted to just those pairs that stand to create a larger structure match, as in SME), so the number of matches considered will in general be larger. In CopyCat^{28,29} and TableTop³⁰, which relationships can be matched is hard-wired in a table associated with each program (the *slipnet*). ACME³¹ used a similarity table to decide which predicates could match, deriving the similarity table when possible from WordNet lexical constraints. In Environmental Model of Analogy (EMMA),³² co-occurrence information computed via Latent Semantic Analysis (LSA) was used for semantic filtering in mapping and retrieval. This method failed to match human mapping preferences, which are governed by relational matches. However, it provided a reasonable match to retrieval patterns which, as described below, are more sensitive to surface matches.

The second major choice in models of mapping is how the mapping is constructed. There are three

basic plans: bottom-up, top-down or middle-out. Bottom-up models (e.g., Ref 33)^b generate sets of correspondences between entities and see which relations can match as a consequence. Top-down models (e.g., CARL,³⁴ Greiner's NLAG model,³⁵ IAM,³⁶ and LISA³⁷ start from key statements in the base and attempt to find matches for them in the target. Neither of these approaches scales particularly well. For bottom-up strategies, if there are N entities in the base and target, then the search space of entity correspondences is of size $N!$.³³ For top-down strategies, there is the additional problem of selecting which aspects of the base to project, followed by finding ways to fit it to the target. Consequently, such models tend to use a variety of heuristics to minimize search. The third class of algorithm is the *middle-out*—or more accurately, *local-global*—matching process introduced in the SME,^{39,40} and also used by ACME and CAB. Local-global algorithms begin by finding all possible identity matches between the potential analogs, in parallel—both low-level information, such as object attributes, and high-level information, such as causal relations. This creates an initial set of correspondences based on identity. Arguments of these candidate statements are then aligned, often using weaker identity constraints, since they are already known to be part of a larger structure. Potential correspondences between entities are hypothesized only when there are statements that place them into alignment.

A related distinction is whether the mapping process is alignment-first or projection-first. Alignment-first models, such as SME or ACME, begin by aligning the base and target, and on the basis of the alignment, further inferences are projected from base to target. Projection-first models, such as LISA, begin by projecting information from the base (or *driver*) to the target (or *recipient*). In practice, alignment-first models are generally use either a bottom-up or a middle-out order, and projection-first models use top-down matching order.

Note that the surface properties of the entities themselves are assumed to be represented explicitly in the base and target (e.g., Gray(Fido)), and such statements also lead to correspondences. This means that entities that are similar (i.e., that have similar attributes) will be suggested as possible correspondences. Because the local match stage is assumed to happen in parallel, attribute matches can either support or conflict with relational matches in the subsequent merge algorithm. Thus, literal similarity matches, in which the entity matches are consistent with the maximal relational alignment, are very fast to compute, consistent with the human pattern.³ In

contrast, cross-mapped matches, in which the entity matches are inconsistent with the maximal relational alignment, are more challenging to compute; If the entity matches are sufficiently rich relative to the potential relational alignment, SME may settle on an entity match, missing the relational alignment.⁴¹

Once the local correspondences have been found, they must be combined into mappings—structurally consistent systems of correspondences that constitute the output of the match process. SME first groups the initial correspondences into internally consistent groups (*kernels*), assigning each an evaluation score by combining the numerical scores computed for each node. Then it uses a greedy merge algorithm to combine the kernels into globally consistent mapping(s). More than one mapping can be produced if they are sufficiently close in size. SME's incremental greedy merge algorithm yields polynomial-time performance.⁴² ACME and CAB use a similar local-global strategy to construct local correspondences, but use parallel constraint satisfaction implemented via a connectionist network to create a pattern of activation corresponding to a mapping.

In CopyCat^{28,29} and TableTop,³⁰ representation and mapping are interleaved; the mapping can influence the representations. They use a *parallel terraced scan* for both representation and mapping. Rules are used to elaborate the input representations and to suggest correspondences between them, based on a table of allowable correspondences. Heuristic estimates of interestingness are used to control how much processing power each rule gets. This inspired key features of Associative Memory-Based Reasoning (AMBR),⁴³ which stores knowledge in units that are smaller than cases. These are viewed as active agents, whose processing is governed by a combination of spreading activation, constraint satisfaction and marker passing. Its structural correspondence mechanism uses a parallel local-global process similar to SME's.

A different approach to mapping is to search for a way that the base could be transformed into the target, thus making them identical. For example, Hahn and Chater's^{44,45} *Representational Distortion* account defines similarity according to the complexity of the transformations needed to make one identical to the other. RD models to date have focused on modeling perceptual comparisons.

As described below, it is useful to consider the kinds of overlap that can occur between two descriptions. In Gentner's typology,⁹ *literal similarity* matches involve overlap in both relations and attributes: e.g., one Prius is typically quite like another Prius. *Analogy* matches involve mostly overlap in relations, with little surface overlap. Cross-domain

analogies, such as solar system/atom, are examples of this type. Literal similarity matches express within-domain comparisons: e.g., starting up one Prius is like starting up another. *Mere-appearance* or *surface* matches involve mostly overlap in attributes and perhaps a few first-order relations: e.g., A toy car looks like a real car. Importantly, structural alignment algorithms do not need *a priori* 'modes' to look for different types of mappings. In SME, the same alignment process is used for all these match types; it will produce a different outcome—analogy, literal similarity or mere-appearance—based on the kind of match between base and target.

An important aspect of analogical reasoning is computing new inferences from the analogy. In symbolic projection-based models, like IAM and heuristic-driven theory projection (HDTP) (see below), the non-overlapping projected structure is the inference. In alignment-first models like SME, inferences are computed after the common structure is identified, by finding propositions connected to the common system in one analog (the base), but not yet present in the other. Candidate inferences are essentially adding structure, which is somewhat more difficult in connectionist models. CAB does not model inferences at all, and ACME required the insertion of a special unit representing the form of a desired inference in the target. LISA can recruit new units to represent projected relations, but given that it is limited to working with around three relations at a time, it does not appear to be able to do complex nested inferences.

The process of structural alignment appears to be used psychologically to compute differences as well as similarities.^{5,46} An *alignable difference* is a difference that is related to the commonalities represented by the mapping. For example, both a regular car and a Prius have keys with which they are started, but the key of a Prius is a block of plastic while the key of a regular car is a flat piece of metal. Such differences can be detected by pairs of conflicting candidate inferences: e.g., what the respective keys are made of. Alignable differences are more salient in human comparison than are differences that are not related to the mapping (*non-alignable differences*)⁴⁶: e.g., differences in party affiliation between Prius drivers and Hummer drivers.

Retrieval

One of the surprises in the psychology of analogy is that retrieval is governed by different constraints than mapping. Given two potential analogs, people prefer mappings involving relational structure, with more systematic structure being preferred, over surface matches. In contrast, retrieval from long-term

memory is dominated by surface matches.^{14–16} Specifically, if literally similar memories are available for a given probe item, those are most likely to be retrieved, followed by mere-appearance matches, with purely relational analogies being the least frequent. Although the dominance of surface over relational matches might seem like a design flaw in human memory, it has been proposed that this is a reasonable strategy ecologically, since (a) things that look alike tend to be alike in causal properties as well and (b) mental representations are skewed toward concrete surface properties, since those are what are delivered by perception and hence highly likely to be encoded.¹⁷ This disconnect has led some researchers to propose that retrieval should be viewed as a separate process. ARCS⁴⁷ for example, used a two-stage connectionist network, which first filtered candidate memory items in parallel and then used ACME to match the best. Likewise, when MAC/FAC²⁰ has a case in working memory, it computes a simple feature vector from the structural representation and uses that in a parallel search of long-term memory. It generates up to three candidate memory items. This generates a mix of literally similar and surface-similar matches, with an occasional analogy. The corresponding structural representations are then compared in parallel via SME to produce one or more reminders.

Another alternative is to consider retrieval and mapping as an integrated process. For example, LISA³⁷ represents propositions using a symbolic connectionist scheme, where the roles arguments play in relations are reified and connected to semantic features, as entities are. These shared semantic features help prime both retrieval of descriptions and mapping connections between entities. Similarly, AMBR⁴³ uses hybrid symbolic-connectionist architecture to encode LTM contents, with retrieval operations interleaved with the mapping process.

While most analogy research uses structured representations, a recent proposal views relations as transformations between points in a continuous similarity space. For example, in the special case of single-relation comparisons (i.e., *A is to B as C is to ?*), Leech et al.⁴⁸ propose that relational priming can be used to explain the sequence of phenomena found in development. Such analogies are solved, they argue, by retrieving a relevant transformation between A and B, which then primes both the retrieval of a relationship between C and what it is transformed into, using a recurrent network.

Generalization

Psychologically, comparing specific descriptions can lead to a generalization.^{49–51} How and when this happens is still very much an open question. The simplest model is to simply replace the entities in the common structure with variables, to produce a simple schema, with the candidate inferences becoming consequences associated with that schema. Winston's³⁸ system produced rules based on comparisons, but also stored the original cases along with the rule so that the precedents could be re-examined when applying the rule. LISA³⁷ can be used to produce schemas by a form of self-supervised learning, adding new units to represent the commonalities found during its mapping process, although it does not compute probabilities for aspects of its generalizations.

Another model of generalization is *antiunification*, which involves finding the least general unifier of two expressions. A unifier of two expressions is a statement with variables that, with the appropriate substitutions of values for variables, will be identical to the two expressions. A unifier of an expression thus captures in some sense what is common between two expressions. A least general unifier is a unifier with the fewest variables, thereby preserving more of the shared common structure between the two expressions. Antiunification has been used in analogy models such as HDTP^{52,53}, which has been used to model proportional analogies⁵⁴ and geometric analogies.⁵⁵

SME computes a generalization by preserving the common structure from a pair of analogs, and has been used to model the psychological phenomenon that such generalizations are more abstract and more likely to be transferred to future analogs than are the initial items.⁵⁶ It is also sometimes desirable to generalize over a set of examples—for example, when learning a new category. SEQL⁵⁷ constructs generalizations over multiple examples. For any concept being learned, SEQL stores the first exemplar and then compares the next exemplar to the first. If they are sufficiently similar, their generalization is stored. SEQL maintains two lists, a list of generalizations and a list of unassimilated exemplars. Each subsequent exemplar is first compared against the generalization(s), using SME. If it is sufficiently similar, it is assimilated into the generalization (which may be altered by the alignment with the exemplar). In the assimilation process, overlapping statements are merged, with corresponding entities that are not identical being replaced by placeholders. (Unlike most generalization algorithms, variables are not introduced.) A probability is associated with each statement and is updated during the

assimilation process.⁵⁸ For example, a generalization about swans might include the information

```
probabilityOf(White(Swan81), 0.99)
probabilityOf(Black(Swan81), 0.01)
```

based on the observed frequency of individuals that have been assimilated into that generalization. Accidental properties lead to low-probability statements, which eventually are filtered out (or become inaccessible) if their probability remains low for a long time. If no generalization is close enough, the new example is then compared against the list of exemplars, and if it is similar enough to one of them, a new generalization is constructed by the same assimilation process. This model supports disjunctive concepts (since there can be more than one generalization) and exceptional cases associated with a concept (through the list of exemplars). Note that the assimilation process does not introduce variables: the abstracted entities are still concrete, albeit now more prototypical of the concept. SEQL has been used in models of grammar learning in infants,¹³ learning spatial prepositions across multiple languages,⁵⁹ and hypothesizing perpetrators in terrorist attacks.⁶⁰

Encoding and Rerepresentation

All models of mapping and retrieval are sensitive to the particular representations used to encode the base, target, and memory items. Computationally, this can be understood as a consequence of the expense of arbitrary inference. Since similarity computations appear to be 'inner loop' core operations of cognition, i.e., they are used throughout cognitive processes, they cannot rely on exponential (or worse) processes to find matches. This sensitivity has often resulted in the use of hand-coded representations, tailored to the needs of the particular model. However, in some cases, independent models of other psychological processes have been used to create descriptions that serve as inputs to analogical processing. PHINEAS used descriptions of physical behavior produced by qualitative simulation. Natural language input has been used with analogy-based models of learning intuitive physics⁶¹ and moral decision-making.⁶² A sketch understanding system has been used to produce inputs for modeling the learning of spatial prepositions⁵⁹ and geometric analogies.⁶³

There is psychological evidence that comparison can affect the final representations of the analogs.^{50,64,65} How to best model this interdependency between initial encoding, comparison process and comparison and final encoding is still an open question. Hofstadter's group has argued that mapping

cannot be separated from encoding, and as noted above, their CopyCat and TableTop models run rules that do encoding and rules that do matching at the same time. One disadvantage of these models is that they are domain-specific; neither could operate in the other's domain. AMBR⁴³ takes a similar approach, embedding knowledge in *micro-agents* whose speed of processing is a function of their estimated relevance to the current problem. AMBR uses the same mechanism to implement both semantic and episodic memory, thus providing a functional equivalent of long-term memory, which is missing from CopyCat and TableTop.

The alternative to complete integration is to interleave encoding and analogical processing. PHINEAS⁶⁶ introduced a *map-analyze* cycle, where the results of the first round of similarity computation were used to influence subsequent processing. Salvucci and Anderson⁶⁷ provide evidence that mapping and other kinds of problem solving can be tightly interleaved, using a story mapping task. By assuming that attributes are computed before relations, psychological findings on response times in visual similarity tasks have been modeled with SME, which can incrementally update its mappings in response to the ongoing encoding of the base and target.⁶⁸

One way that analogy influences encoding is via rerepresentation. That is, people seem to reconstrue the contents of the base and/or the target in order to improve the match.³ In the HDTP approach, rerepresentation is implemented by logical inference rules which operate as part of the antiunification process.⁶⁹ The constraints of structure-mapping theory have been used to derive a theory of rerepresentation,⁷⁰ identifying opportunities for rerepresentation based on local changes that could improve the overall match.

Over the long term, gains in expertise appear to lead to changes in encoding,⁷¹ as well as to greater likelihood of relational retrieval.⁷² Forbus et al.²⁰ proposed that uniform relational encoding is a characteristic of expert encoding, and that this promotes relational retrieval. Finlayson and Winston⁷³ argue that encodings that lead to intermediate-sized representations provide more expert-like retrieval.

LARGER SCALE SIMULATIONS

The hypothesis that analogy plays a central role in human cognition suggests using analogy within larger-scale models that capture broader swaths of human cognition, in which models of individual processes used in analogy are used in their proposed functional roles. The integrated encoding approaches

(e.g., Hofstadter's group, AMBR) provide examples of this. A path-mapping model of analogical mapping has been developed for Adaptive Control of Thought-Rational (ACT-R),⁶⁷ supporting the integrated modeling of mapping with other kinds of ACT-R modeling. Recently several cognitive architectures have been proposed with analogy at their core. DUAL⁷⁴ uses AMBR for mapping and retrieval, using a hybrid symbolic/connectionist representation scheme. DUAL has been used to model priming and context effects on problem-solving. The Companions cognitive architecture⁷⁵ is based on structure-mapping models. SME is used for mapping, MAC/FAC for retrieval, and SEQL for generalization. To date the companions architecture is the only one that has been tested in experiments in which the inputs were produced by groups other than the researchers, and where the results were independently evaluated by other organizations (e.g., learning Advanced Placement (AP) Physics, evaluated by the Educational Testing Service, and learning simple games, evaluated by the US Naval Research Laboratory).

One challenge for current models of analogy is what these investigations are revealing about the scale of representations used in analogical processing. For example, the number of relationships needed to represent problems and solutions in technical domains (e.g., physics, engineering) is on the order of 10 or more, with visual descriptions being substantially larger. Today's connectionist models cannot handle such representations: For example, the particular synchronous binding scheme used by LISA means that it can match only three relations at a time; whether it can handle complex analogies by shifting the focus of attention around different parts of the representations remains an open question. On the other hand, SME does not currently model working memory limitations at all, which is also unrealistic. Coming up with unified models, that are both capable of human-like performance on realistic tasks and have a clear, biologically plausible implementation, remains an open problem. A second challenge to current models is the question of hand-coding. In most current models, the representations are created by the experimenters, leading to the tailorability concern: that is, that (whether knowingly or not) the researchers have encoded the items in such a way as to give them the desired results. One way to avoid hand-coding is to use pre-existing databases and automatic (or semi-automatic) parsing and semantic representation of the input text.⁶² Another route, at least for visual materials, is to use automatic spatial encoding of sketched materials.^{59,68}

CONCLUSION

By combining constraints and insights from cognitive psychology and artificial intelligence, substantial progress has been made in modeling a variety of phenomena involved in analogical processing. Existing models of analogical mapping, retrieval, and generalization have been used to model a wide variety of psychological phenomena, and have been used to make predictions that shed new light on cognition. One promising future direction is the use of analogical models to capture learning processes in cognitive development (e.g., Refs 13,61,76,77). Another potentially fruitful direction is the use of analogical simulations in intelligent tutoring systems and learning environments, to improve education and training.

Much research remains, of course, before we have a complete account of analogical processing. For example, work on larger-scale simulations involving

analogical processing, to explore the roles analogy plays in perception, reasoning, categorization, and learning represents an exciting new frontier which is only beginning to be explored. Another exciting direction is seeking biologically plausible ways of implementing analogical processing, which will need to evolve as we gain better understanding of neural systems.

NOTES

^a For example, if a parallel implementation is assumed, and a fixed upper bound on the size of description to be processed, only N^2 processing units would have to be set aside to represent correspondences.

^b Winston later added importance-dominated matching to the simulation,³⁸ a top-down feature.

REFERENCES

1. Markman AB. *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates; 1999.
2. Clement CA, Gentner D. Systematicity as a selection constraint in analogical mapping. *Cognitive Science* 1991, 15:89–132.
3. Gentner D, Kurtz K. Relations, objects, and the composition of analogies. *Cognitive Science* 2006, 30:609.
4. Spellman BA, Holyoak KJ. If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology* 1992, 62:913–933.
5. Gentner D, Sagi E. Does “different” imply a difference? A comparison of two tasks. In: Sun R, Miyake N, eds. *Proceedings of the Twenty-eighth Annual Meeting of the Cognitive Science Society*, 2006.
6. Goldstone RL, Medin DL, Gentner D. Similarity involving attributes and relations: Judgments of similarity and difference are not inverses. *Psychological Science* 1990, 1:64–69.
7. Love BC, Rouder JN, Wisniewski EJ. A structural account of global and local processing. *Cognitive Psychology* 1999, 38:291–316.
8. Markman AB, Gentner D. Structural alignment during similarity comparisons. *Cognitive Psychology* 1993, 25:431–467.
9. Gentner D. Structure-mapping: a theoretical framework for analogy. *Cognitive Science* 1983, 7:155–170.
10. Goldstone RL, Medin DL. Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, & Cognition* 1994, 20:29–50.
11. Cheng PW, Holyoak KJ. Pragmatic reasoning schemas. *Cognitive Psychology* 1985, 17:391–416.
12. Gentner D, Medina J. Similarity and the development of rules. *Cognition* 1998, 65:263–297.
13. Kuehne SE, Gentner D, Forbus KD. Modeling infant learning via symbolic structural alignment. In: Gleitman L, Joshi AK, eds. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, 2000, 286–291.
14. Gentner D, Rattermann MJ, Forbus KD. The roles of similarity in transfer: Separating retrievability from inferential soundness. *Cognitive Psychology* 1993, 25:524–575.
15. Holyoak KJ, Koh K. Surface and structural similarity in analogical transfer. *Memory & Cognition* 1987, 15:332–340.
16. Ross BH. Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory and Cognition* 1989, 15:456–468.
17. Gentner D. The mechanisms of analogical learning. In: Vosniadou S, Ortony A, eds. *Similarity and Analogical Reasoning* London: Cambridge University Press; 1989, 199–241.
18. Gentner D. Why we’re so smart. In: Gentner D., Goldin-Meadow S., eds. *Language in Mind: Advances in the Study of Language and Cognition*. Cambridge, MA: MIT Press; 2003, 195–236.
19. Gentner D, Markman AB. Defining structural similarity. *The Journal of Cognitive Science* 2006, 6:1–20.
20. Forbus KD, Gentner D, Law K. MAC/FAC: A model of similarity-based retrieval. *Cognitive Science* 1995, 19:141–205.

21. Gentner D. Are scientific analogies metaphors?. In: Miall DS, ed. *Metaphor: Problems and Perspectives*. Brighton, England: Harvester Press; 1982, 106–132.
22. Gentner D, Clement C. Evidence for relational selectivity in the interpretation of analogy and metaphor. In: Bower GH, ed. *The Psychology of Learning and Motivation, Advances in Research and Theory*. New York: Academic Press; 1988, 307–358.
23. Spellman BA, Holyoak KJ. Pragmatics in analogical mapping. *Cognitive Psychology* 1996, 31:307–346.
24. Markman AB. Constraints on analogical inference. *Cognitive Science* 1997, 21:373–418.
25. Krawczyk DC, Holyoak KJ, Hummel JE. The one-to-one constraint in analogical mapping and inference. *Cognitive Science* 2005, 29:797–806.
26. Falkenhainer B. *Learning from physical analogies*, Technical report no. UIUCDCS-R-88-1479. Ph.D. thesis. Urbana-Champaign: University of Illinois; 1988.
27. Larkey LB, Love BC. CAB: Connectionist Analogy Builder. *Cognitive Science* 2003, 27:781–794.
28. Mitchell M. *Analogy-making as perception: A computer model*. Cambridge, MA: The MIT Press; 1993.
29. Hofstadter DH. *Fluid concepts and creative analogies*. New York: Basic Books; 1995.
30. French RM. *The subtlety of similarity*. Cambridge, MA: The MIT Press; 1995.
31. Holyoak KJ, Thagard PR. A computational model of analogical problem solving. In: Vosniadou S, Ortony A, eds. *Similarity and Analogical Reasoning*. New York: Cambridge University Press; 1989, 242–266.
32. Ramscar MJA, Pain HG. Can a real distinction be made between cognitive theories of analogy and categorization?. *Proceedings of the 18th Annual Conference of the Cognitive Science Society*. San Diego: University of California; 1996.
33. Winston PH. Learning and reasoning by analogy. *Communications of the ACM* 1980, 23:689–703.
34. Burstein MH. A model of learning by incremental analogical reasoning and debugging. *Proceedings of the National Conference on Artificial Intelligence*; 1983, 45–48.
35. Greiner R. Learning by understanding analogies. In: Prieditis A, ed. *Analogica*. Los Altos, CA: Kaufmann; 1988, 1–36.
36. Keane MT. On order effects in analogical mapping: predicting human error using IAM. In: Moore JD, Lehmann JF, eds. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum; 1995.
37. Hummel JE, Holyoak KJ. LISA: A computational model of analogical inference and schema induction. *Psychological Review* 1997.
38. Winston PH. Learning by augmenting rules and accumulating censors. In: Michalski RS, Carbonell JG, Mitchell TM, eds. *Machine Learning: An Artificial Intelligence Approach*. Los Altos, CA: MorganKaufmann; 1986, 45–61.
39. Falkenhainer B, Forbus K, Gentner D. The Structure-Mapping Engine. *Proceedings of AAAI-86*, Philadelphia, PA, August 1986.
40. Falkenhainer B, Forbus KD, Gentner D. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 1989, 41:1–63.
41. Loewenstein J, Gentner D. Relational language and the development of relational mapping. *Cognitive Psychology* 2005, 50:315–353.
42. Forbus KD, Ferguson RW, Gentner D. Incremental structure mapping In: *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, 1994.
43. Kokinov BN, Petrov AA. In: Integrating memory and reasoning in analogy-making: the AMBR model. Gentner D, Holyoak KJ, Kokinov BN, eds. *The Analogical Mind: Perspectives From Cognitive Science*. Cambridge, MA: MIT Press; 2001, 161–196.
44. Hahn U, Chater N, Richardson LB. Similarity as transformation *Cognition* 2003, 87:1–32
45. Hodgetts C.J. Hahn U., Chater N. Transformation and alignment in similarity. *Cognition* 2009, 113:62–79.
46. Markman AB, Gentner D. Splitting the differences: a structural alignment view of similarity. *Journal of Memory and Language* 1993, 32:517–535.
47. Thagard P, Holyoak KJ, Nelson G, Gochfeld D. Analog retrieval by constraint satisfaction. *Artificial Intelligence* 1990, 46:259–310.
48. Leech R, Mareschal D, Cooper R. Analogy as relational priming: a developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences* 2008, 31:357–414.
49. Gick ML, Holyoak KJ. Schema induction and analogical transfer. *Cognitive Psychology* 1983, 15:1–38.
50. Catrambone R, Holyoak KJ. Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory and Cognition* 1989, 15:1147–1156.
51. Loewenstein J, Thompson L, Gentner D. Analogical learning in negotiation teams: Comparing cases promotes learning and transfer. *Academy of Management Learning and Education* 2003, 2:119–127.
52. Gust H, Kuhnberger KU, Schmid U. metaphors and heuristic-driven theory projection (HDTP). *Theoretical Computer Science* 2006, 354:98–117.
53. Schwering A, Krumnack U, Kuhnberger KU, Gust H. Syntactic principles of heuristic-driven theory projection. *Cognitive Systems Research* 2009, 10:251–269.
54. Schmid U, Gust H, Kuhnberger K-U, Burghardt J. In: An algebraic framework for solving proportional and predictive analogies. Schmalhofer F, Young R, Katz G, eds. *Proceedings of the European Conference on Cognitive Science (EuroCogSci 2003)*. Osnabrück, Germany: Lawrence Erlbaum; September 10–13, 2003, 295–300.

55. Schwering A, Gust H, Kuhnberger K-U, Krumnack U. Solving geometric proportional analogies with the analogy model HDTP. *Proceedings of CogSci09*; 2009.
56. Gentner D, Loewenstein J, Thompson L, Forbus K. Reviving inert knowledge: Analogical abstraction supports relational retrieval of past events. *Cognitive Science* 2009, 33:1343–1382.
57. Kuehne SE, Forbus KD, Gentner D, Quinn B. SEQL: Category learning as progressive abstraction using structure mapping. In: Gleitman LR, Joshi AK, eds., *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, Philadelphia, PA; 2000, 770–775.
58. Halstead D, Forbus KD. Transforming between propositions and features: bridging the gap. *Proceedings of AAAI-05*. Pittsburgh, PA; 2005.
59. Lockwood K, Lovett A, Forbus K. Automatic Classification of Containment and Support Spatial Relations in English and Dutch. In: *Proceedings of Spatial Cognition*; 2008.
60. Halstead D, Forbus K. Some Effects of a Reduced Relational Vocabulary on the Whodunit Problem. *Proceedings of IJCAI-2007*, Hyderabad, India; 2007.
61. Friedman SE, Forbus KD. Learning naive physics models and misconceptions. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Amsterdam, Netherlands; 2009.
62. Dehghani M, Tomai E, Forbus K, Iliev R, Klenk M. MoralDM: a computational modal of moral decision-making. *Proceedings of the 30th Annual Conference of the Cognitive Science Society (CogSci)*, Washington, DC; 2008.
63. Lovett A, Tomai E, Forbus K, Usher J. Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science* 2009.
64. Loewenstein J, Thompson L, Gentner D. Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review* 1999, 6:586–597.
65. Gentner D, Namy L. Comparison in the development of categories. *Cognitive Development* 1999, 14:487–513.
66. Falkenhainer B. A unified approach to explanation and theory formation. In: Shrager J, Langley P, eds. *Computational Models of Scientific Discovery and Theory Formation*. Morgan Kaufmann Publishers; 1990, 157–196.
67. Salvucci DD, Anderson JR. Integrating analogical mapping and general problem solving: the path-mapping theory. *Cognitive Science* 2001, 25:67–110
68. Lovett A, Gentner D, Forbus K, Sagi E. Using analogical mapping to simulate time-course phenomena in perceptual similarity *Cognitive Systems Research* 2009, 10:216–228.
69. Krumnack U, Gust H, Kuhnberger K-U, Schwering A. Re-representation in a logic-based model for analogy-making. *Proceedings of the 21st Australasian Joint Conference on Artificial Intelligence*, Auckland, New Zealand; 2008.
70. Yan J, Forbus KD, Gentner D. A theory of rerepresentation in analogical matching. *Proceedings of the 25th Annual Conference of the Cognitive Science Society*; 2003.
71. Chi M, Feltovich P, Glaser R. Categorization and representation of physics problems by experts and novices. *Cognitive Science* 1981, 5:121–152.
72. Novick LR. Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1988, 14:510–520.
73. Finlayson M, Winston PH. Intermediate features and information-level constraint on analogical retrieval. *Proceedings of CogSci05*; 2005.
74. Kokinov B. A hybrid model of reasoning by analogy. In: Holyoak K, Barnden J, eds. *Advances in Connectionist and Neural Computation Theory. Volume 2: Analogical Connections*. Norwood, NJ: Ablex; 1994, 247–320.
75. Forbus K, Klenk M, Hinrichs T. Companion cognitive systems: Design goals and lessons learned so far. *IEEE Intelligent Systems* vol. 2009, 24:36–46.
76. Doumas. LAA, Hummel JE, Sandhofer CM. A theory of the discovery and predication of relational concepts. *Psychological Review* 2008, 115:1–43.
77. Friedman S, Taylor J, Forbus K. Learning naive physics models by analogical generalization. *Proceedings of the 2nd International Analogy Conference*, Sofia, Bulgaria; 2009.

FURTHER READING

- French RM. The computational modeling of analogy-making. *Trends in Cognitive Science* 2002, 6:200–205.
- Gentner D, Holyoak KJ, Kokinov BN. *The analogical mind: Perspectives from cognitive science*. Cambridge, MA: MIT Press; 2001.
- Gentner D, Markman AB. Structure-mapping in analogy and similarity. *American Psychologist* 1997, 52:45–56.
- Kokinov B, French RM. Computational models of analogy making. In: Nadel L., ed. *Encyclopedia of Cognitive Science*. London: MacMillan; 2002.