Level 2

Computational Models of Analogy-Making

Boicho Kokinov, New Bulgarian University, Sofia, Bulgaria
Robert M French, University of Liège, Belgium

CONTENTS

Analogy-making is the process of finding or constructing a common relational structure in the descriptions of two situations or domains and making inferences by transferring knowledge from the familiar domain (the 'base' or 'source') to the unfamiliar domain (the 'target'), thus enriching our knowledge about the latter.

INTRODUCTION

0098.001 Analogy-making is crucial for human cognition. Many cognitive processes involve analogy-making in one way or another: perceiving a stone as a human face, solving a problem in a way similar to another problem previously solved, arguing in court for a case based on its common structure with another case, understanding metaphors, communicating emotions, learning, or translating poetry from one language to another (Gentner *et al*., 2001). All these applications require an abstract mapping to be established between two cases or domains based on their common structure (common systems of relations). This may require re-representation of one (or both) of the domains in terms of the other one (or in terms of a third domain). The first domain is called the base, or source, and the second is called the target.

0098.002 Analogy-making is a basic cognitive ability. It appears to be present in humans from a very early age, and develops over time. It starts with the simple ability of babies to imitate adults and to recognize when adults are imitating them, progresses to children's being able to recognize an analogy between a picture and the corresponding real object, and culminates in the adult ability to make complex analogies between various situations. This seems to suggest that analogy-making serves as the basis for numerous other kinds of human thinking; hence the importance of developing computational models of analogy-making.

Analogy-making involves at least the following sub-processes: building a representation, retrieving a 'base' for the analogy, mapping this base to the 'target', transferring unmapped elements from the base to the target, thereby making inferences, evaluating the validity and applicability of these inferences, and learning from the experience – which includes generalizing from specific cases and, possibly, developing general mental schemata. There are, at present, no models that incorporate all of these sub-processes. Individual models focus on one or more of them. 0098.003

Representation-Building

The process of representation-building is absent from most models of analogy-making. Typically, representations are fed into the model. However, there are some models (e.g., ANALOGY, Copycat, Tabletop, Metacat) that do produce their own high-level representations based on essentially unprocessed input. These models (Mitchell, 1993; Hofstadter *et al*., 1995; French, 1995) attempt to build flexible, context-sensitive representations during the course of the mapping phase. Other models, such as AMBR (Kokinov and Petrov, 2001), perform re-representation of old episodes. 0098.004

Retrieval

The retrieval process has been extensively studied experimentally. Superficial similarity is the most important factor in analogical retrieval: the retrieval of a base for analogy is easier if it shares similar objects, similar properties, and similar general themes with the target. Structural similarity, the familiarity of the domain from which the analogy is drawn, the richness of its representations, and the presence of generalized schemata, also facilitate retrieval. Most models of retrieval are based 0098.005

on exhaustive search of long-term memory (LTM) and on the assumption that old episodes have context-independent, encapsulated representations. There are, however, exceptions (e.g., AMBR) that rely on context-sensitive reconstruction of old episodes performed in interaction with the mapping process.

Mapping

0098.006 Mapping is the core of analogy-making. All computational models of analogy-making include mapping mechanisms, i.e., means of discovering which elements of the base correspond to which elements of the target. The difficulty is that one situation can be mapped to a second situation in many different ways. We might, for example, make a mapping based on the color of the objects in both the base and target (the red-shirted person in the base domain would be mapped to the red-shirted person in the target domain). This would, in general, be a very superficial mapping (but might, nonetheless, be appropriate on occasion). We could also map the objects in the two domains based on their relational structures. For example, we could decide that it was important to map the giver–receiver relationship in the first domain to the same relationship in the target domain, ignoring the fact that in the base domain the giver had a red shirt and in the target domain the giver was wearing a blue shirt.

0098.007 Experimental work has demonstrated that finding this type of structural isomorphism between base and target domains is crucial for mapping (Gentner, 1983). Object similarity also plays a role in mapping, although generally a secondary one. A third factor is the pragmatic importance of various elements in the target: we want to find mappings that involve the most important elements in the target. Searching for the appropriate correspondences between the base and the target is a computationally complex task that can become infeasibly time-consuming if the search is unconstrained.

Transfer

0098.008 New knowledge then has to be inserted into the target domain based on the mapping. For example, suppose a new brand of car appears on the market, and that this car maps well onto another brand of car that is small, fast, and handles well on tight curves. But you also know that this latter brand of car is frequently in need of repair. You then wonder whether the new brand of car will also be in frequent need of repair.

Transfer is present in some form in most models of analogy-making, and is typically integrated with mapping. Transfer is considered by some researchers as an extension of the mapping already established, adding new elements to the target. 0098.009

Evaluation

Evaluation is the process of establishing the likelihood that the transferred knowledge will be applicable to the target domain. In the example above, the evaluation process would have to assign the degree of confidence we would have that the new car would also be in frequent need of repair. Evaluation is often implicit in the mechanisms of mapping and transfer. 0098.010

Learning

Only a few models of analogy-making have incorporated learning mechanisms. This is somewhat surprising since analogy-making is clearly a driving force behind much learning. However, some models are capable of generalization from the base and target, or from multiple exemplars, to form an abstract schema, as in LISA (Hummel and Holyoak, 1997) and the SEQL model based on SME (Falkenhainer *et al*., 1989). 0098.011

Below we will review a number of important computational models of analogy-making belonging to different classes and following different approaches. First the 'symbolic' models will be presented. These models employ separate local representations of objects, relations, propositions and episodes (e.g., 'John', 'chair', 'run', 'greater than', 'John ate fish', 'my birthday party last year'). Then, 'connectionist' models will be presented. Here the objects, relations, and episodes are represented as overlapping patterns of activation in a neural network. Finally, a third, hybrid class of models will be presented. These models combine symbolic representations with connectionist activations. They are based on the idea that cognition is an emergent property of the collective behavior of many simple agents. 0098.012

SYMBOLIC MODELS

ANALOGY

The earliest computational model of analogy-making, ANALOGY, was developed by Thomas Evans (1964). This program solves multiple-choice geometric analogy problems of the form '*A* is to *B* 0098.013

as *C* is to what?' taken from intelligence tests and college entrance examinations.

0098.014 An important feature of this program is that the input is not a high-level description of the problem, but a low-level description of each component of the figure – dots, simple closed curves or polygons, and sets of closed curves or polygons. The program builds its own high-level representation describing the figures in *A*, *B*, *C*, and all given alternatives for the answer, with their properties and relationships – for example, `((P1 P2 P3) . ((INSIDE P1 P2) (LEFT P1 P3) (LEFT P2 P3)))`. Then the program represents the relationship between *A* and *B* as a set of possible rules describing how figure *A* is transformed into figure *B* – for example, `((MATCH P2 P4) (MATCH P1 P5) (REMOVE P3))` which means that figure P_2 from *A* corresponds to figure P_4 from *B*, P_1 corresponds to P_5, and the figure P_3 does not have a corresponding figure and is therefore deleted. Then each such rule is applied to *C* in order to get one of the alternative answers. In fact, each such rule would be generalized in such a way as to allow *C* to be applied to *D*. Finally, the most specific successful rule would be selected as an outcome. Arguably, one of the most significant features of the program is its ability to represent the target problem on its own – a feature that has been dropped in most recent models.

Structure Mapping Theory

0098.015 The most influential family of computational models of analogy-making have been those based on Dedre Gentner's (1983) 'structure mapping theory' (SMT). This theory was the first to explicitly emphasize the importance of structural similarity between base and target domains, defined by common systems of relations between objects in the respective domains. Numerous psychological experiments have confirmed the crucial role of relational mappings in producing convincing and sound analogies. There are several important assumptions underlying the computational implementation of SMT called SME (Falkenhainer *et al.*, 1989): (1) mapping is largely isolated from other analogy-making sub-processes (such as representation, retrieval and evaluation) and is based on independent mechanisms; (2) relational matches are preferred over property matches; (3) only relations that are identical in the two domains can be put into correspondence; (4) relations that are arguments of higher-order relations that can also be mapped have priority, following the 'systematicity principle' that favors systems of relations over isolated relations; and (5) two or three interpretations are constructed by a 'greedy merge' algorithm that generally finds the 'best' structurally coherent mapping. Early versions of SME mapped only identical relations and relied solely on relational structure. This purely structural approach was intended to ensure the domain independence of the mapping process. Recent versions of SME have made some limited use of pragmatic aspects of the situation, as well as re-representation techniques that allow initially non-matching predicates to match.

The MAC/FAC model of analogical retrieval 0098.016
(Forbus *et al.*, 1995) was intended to be coupled with SME. This model assumes that episodes are encapsulated representations of past events; they have a dual encoding in LTM: a detailed predicate-calculus representation of all the properties and relations of the objects in an episode and a shorter summary (a vector representation indicating the relative frequencies of the predicates that are used in the detailed representation). The retrieval process has two stages. The first stage uses the vector representations to perform a superficial search for episodes that share predicates with the target problem. The episode vectors in LTM that are close to the target vector are selected for processing by the second stage. The second stage uses the detailed predicate-calculus representations of the episodes to select the one that best matches the target. These two stages reflect the dominance of superficial similarity as well as the influence of structural similarity.

Gentner's ideas – in particular, their emphasis on 0098.017
the structural aspects of analogical mappings – have been very influential in the area of computational analogy-making and have been applied in contexts ranging from child development to folk physics. Various improvements and variants of SME have been developed, and it has been included as a module in various practical applications.

Other Symbolic Models

A number of other symbolic models have helped to 0098.018
advance our understanding of analogy-making. Jaime Carbonell proposed the concept of derivational analogy, where the analogy is drawn not with the final solution of the old problem, but with its derivation, i.e., with the way of reaching the solution, an approach developed further by Manuela Veloso. Smadar Kedar-Cabelli developed a model of purpose-directed analogy-making in concept learning. Mark Burstein developed a model

called CARL which learned from multiple analogies combining several bases. Mark Keane and his colleagues developed an incremental model of mapping, IAM, which helps explain the effects of order of presentation of the material.

CONNECTIONIST MODELS

0098.019 Research in the field of analogy-making has, until recently, been largely dominated by the symbolic approach, for an obvious reason: symbolic models are well equipped to process and compare the complex structures required for analogy-making. In the early years of the new connectionist paradigm, these structures were very difficult to represent in a connectionist network. However, advances in connectionist representation techniques have allowed distributed connectionist models of analogy to be developed. Most importantly, distributed representations provide a natural internal measure of similarity, thereby allowing the system to handle the problem of similar but not identical relations in a relatively straightforward manner. This ability is essential to analogy-making and has proved hard for symbolic models to implement.

Multiple Constraints Theory

0098.020 The earliest attempt to design an architecture in which analogy-making was an emergent process of activation states of neuron-like objects was proposed by Keith Holyoak and Paul Thagard (1989) and implemented in a model called ACME. In this model, structural similarity, semantic similarity, and pragmatic importance determine a set of constraints to be simultaneously satisfied. The model is supplied with representations of the target and of the base, and proceeds to build a localist constraint-satisfaction connectionist network where each node corresponds to a possible pairing hypothesis for an element of the base and an element of the target. For example, if the base is 'train' and the target is 'car' then all elements of trains will be mapped to all elements of cars; there will therefore be hypothesis nodes created not only for 'locomotive → motor' but also for 'locomotive → license plate', 'locomotive → seat-belt buckle', etc. The excitatory and inhibitory links between these nodes implement the structural constraints. In this way, contradictory hypothesis nodes compete and do not become simultaneously active, while consistent ones mutually support each other. The network gradually moves towards an equilibrium state, and the best set of consistent mapping hypotheses (e.g., 'locomotive → motor', 'rails → road', etc.) wins. The relaxation of the network provides a parallel evaluation of all possible mappings and finds the best one, which is represented by the set of most active hypothesis nodes.

ARCS is another related model of retrieval. It is coupled with ACME and operates in a similar fashion. However, while mapping is dominated by structural similarity, retrieval is dominated by semantic similarity. 0098.021

STAR

STAR-1 was the first distributed connectionist model of analogy-making (Halford *et al.*, 1994). It is based on the tensor product connectionist models developed by Smolensky. A proposition like `MOTHER-OF (CAT, KITTEN)` is represented by the tensor product of the three vectors corresponding to `MOTHER-OF`, `CAT`, and `KITTEN`: $\mathtt{MOTHER\text{-}OF} \otimes \mathtt{CAT} \otimes \mathtt{KITTEN}$. The tensor product in this case is a three-dimensional array of numbers where the number in each cell is the product of the three corresponding coordinates. This representation allows any of the arguments, or the relational symbol, to be extracted by a generalized dot product operation: $(\mathtt{MOTHER\text{-}OF} \otimes \mathtt{CAT}) \bullet (\mathtt{MOTHER\text{-}OF} \otimes \mathtt{CAT} \otimes \mathtt{KITTEN}) = \mathtt{KITTEN}$. The LTM of the system is represented by a tensor that is the sum of all tensor products representing the individual statements (the main restriction being that the propositions are simple and have the same number of arguments). Using this type of representation, STAR-1 solves proportional analogy problems like 'cat is to kitten as mare is to what?' 0098.022

STAR-2 (Wilson *et al.*, 2001) maps complex analogies by sequentially focusing on various parts of the domains – simple propositions with no more than four dimensions – and finding the best map for the arguments of these propositions by parallel processing in the constraint satisfaction network (similarly to ACME). The fact that the number of units required for a tensor product representation increases exponentially with the number of arguments of a predicate implies processing constraints in the model. Wilson *et al.* claim that humans are subject to similar processing constraints: specifically, they can, in general, handle a maximum of four dimensions of a situation concurrently. The primary interest of the modelers is in exploring and explaining capacity limitations of human beings and achieving a better understanding of the development of analogy-making capabilities in children. 0098.023

LISA

0098.024 John Hummel and Keith Holyoak (1997) proposed an alternative computational model of analogy-making using distributed representations of structure relying on dynamic binding. The idea is to introduce an explicit time axis so that patterns of activation can oscillate over time (thus the timing of activation becomes an additional parameter independent of the level of activation). Patterns of activation oscillating in synchrony are considered to be bound together, while those oscillating out of synchrony are not. For example, 'John hired Mary' requires synchronous oscillation of the patterns for 'John' and 'Employer' alternating with synchronous oscillation of the patterns for 'Mary' and 'Employee'. Periodic alternation of the activation of the two pairs represents the whole statement. However, if the statement is too complex there will be too many pairs that need to fire in synchrony. Based on research on single-cell recordings, Hummel and Holyoak believe that the number of such pairs of synchronously firing concepts cannot exceed six. Representations in LISA's working memory are distributed over the network of semantic primitives, but representations in long-term memory are localist – there are separate units representing the episode, the propositions, their components, and the predicates, arguments, and bindings. Retrieval is performed by spreading activation, while mapping is performed by learning new connections between the most active nodes. LISA successfully integrates retrieval of a base with the mapping of the base and target, even though retrieval and mapping are still performed sequentially (mapping starts only after one episode is retrieved).

HYBRID MODELS

0098.025 Two groups of researchers have independently produced similar models of analogy-making based on the idea that high-level cognition emerges as a result of the continual interaction of relatively simple, low-level processing units, capable of doing only local computations. These models are a combination of the symbolic and connectionist approaches. Semantic knowledge is incorporated in order to compute the similarity between elements of the two domains in a context-sensitive way.

Copycat and Related Architectures

0098.026 The family of Copycat and Tabletop architectures (Mitchell, 1993; Hofstadter *et al*., 1995; French, 1995) was explicitly designed to integrate top-down semantic information with bottom-up emergent processing. Copycat solves letter-string analogies of the form '*ABC* is to *ABD* as *KLM* is to what?' and gives plausible answers such as *KLN* or *KLD*. The architecture of Copycat involves a working memory, a semantic network (simulating long-term memory) defining the concepts used in the system and their relationships, and the 'coderack' – the procedural memory of the system – a store for small, nondeterministic computational agents ('codelets') working on the structures in the working memory and continually interacting with the semantic network. Codelets can build new structures or destroy old structures in working memory. The system gradually settles towards a consistent set of structures that will determine the mapping between the base and the target.

0098.027 The most important feature of these models of analogy-making is their ability to build up their own representations of the problem, in contrast with most other models which receive the representations of the base and target as input. Thus these models abandon traditional sequential processing and allow representation-building and mapping to run in parallel and continually influence each other. The partial mapping can influence further representation-building, thus allowing the gradual construction by the program of context-sensitive representations. In this way, the mapping may force us to see a situation from an unusual perspective in terms of another situation, and this is an essential aspect of creative analogy-making.

AMBR

0098.028 AMBR (Kokinov, 1994) solves problems by analogy. For example, 'how can you heat some water in a wooden vessel, being in the forest?' The solution, heating a knife in a fire and immersing it in the water, is found by analogy with boiling water in a glass using an immersion heater.

0098.029 The AMBR model is based on DUAL, a general cognitive architecture. The LTM of DUAL consist of many micro-agents, each of which represents a small piece of knowledge. Thus concepts, instances and episodes are represented by (possibly overlapping) coalitions of micro-agents. Each micro-agent is hybrid: its symbolic part encodes the declarative or procedural knowledge it is representing, while its connectionist part computes the agent's activation level, which represents the relevance of this knowledge to the current context. The symbolic processors run at a speed proportional to their computed relevance, so the behavior of the system is highly context-sensitive. The AMBR model implements the interactive parallel work of recollection, mapping and transfer that emerge from the collective behavior of the agents and which produces the analogy. Recollection in AMBR-2 (Kokinov and Petrov, 2001) is reconstruction of the base episode in working memory by activating relevant aspects of event information, of general knowledge, and of other episodes, and forming a coherent representation which will correspond to the target problem. The model exhibits illusory memories, including insertions from general knowledge and blending with other episodes, and context and priming effects. Some of these phenomena have been experimentally confirmed in humans.

CONCLUSIONS

0098.030 The field of computational modeling of analogy-making has moved from the early models, which were intended mainly to demonstrate that computers could, in fact, be programmed to do analogy-making, to complex models that make nontrivial predictions of human behavior. Researchers have come to appreciate the need for structural mapping of the base and target domains, for integration of and interaction between representation-building, retrieval, mapping and learning, and for systems that can potentially scale up to the real world. Computational models of analogy-making have now been applied to a large number of cognitive domains (Gentner *et al.*, 2001). However, many issues remain to be explored in the endeavor to model the human capacity for analogy-making, one of our most important cognitive abilities.

References

Evans TG (1964) A heuristic program to solve geometric-analogy problems. In: *Proceedings of the Spring Joint Computer Conference* **25**: 327–338 [Reprinted in: Fischler M and Firschein O (eds) (1987) *Readings in Computer Vision*. Los Altos, CA: Morgan Kaufmann.]

Falkenhainer B, Forbus KD and Gentner D (1989) The structure-mapping engine: algorithm and examples. *Artificial Intelligence* **41**: 1–63.

Forbus K, Gentner D and Law K (1995) MAC/FAC: a model of similarity-based retrieval. *Cognitive Science* **19** (2): 141–205.

French R (1995) *The Subtlety of Sameness: A Theory and Computer Model of Analogy-Making*. Cambridge, MA: MIT Press.

Gentner D (1983) Structure-mapping: a theoretical framework for analogy. *Cognitive Science* **7**(2): 155–170.

Gentner D, Holyoak K and Kokinov B (eds) (2001) *The Analogical Mind: Perspectives From Cognitive Science*. Cambridge, MA: MIT Press.

Halford G, Wilson W, Guo J *et al.* (1994) Connectionist implications for processing capacity limitations in analogies. In: Holyoak K and Barnden J (eds) *Advances in Connectionist and Neural Computation Theory*, **vol. II** 'Analogical Connections', pp. 363–415. Norwood, NJ: Ablex.

Hofstadter D and the Fluid Analogies Research Group (1995) *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. New York, NY: Basic Books.

Holyoak K and Thagard P (1989) Analogical mapping by constraint satisfaction. *Cognitive Science* **13**: 295–355.

Hummel J and Holyoak K (1997) Distributed representations of structure: a theory of analogical access and mapping. *Psychological Review* **104**: 427–466.

Kokinov B (1994) A hybrid model of analogical reasoning. In: Holyoak K and Barnden J (eds) *Advances in Connectionist and Neural Computation Theory*, **vol. II** 'Analogical Connections', pp. 247–318. Norwood, NJ: Ablex.

Kokinov B and Petrov A (2001) Integration of memory and reasoning in analogy-making: the AMBR model. In: (Gentner *et al.*, 2001), pp. 59–124.

Mitchell M (1993) *Analogy-Making as Perception: A Computer Model*. Cambridge, MA: MIT Press.

Wilson W, Halford G, Gray B and Phillips S (2001) The STAR-2 model for mapping hierarchically structured analogs. In: (Gentner *et al.*, 2001), pp. 125–159.

Further Reading

Barnden J and Holyoak K (eds) (1994) *Advances in Connectionist and Neural Computation Theory*, **vol. III** 'Analogy, Metaphor, and Reminding'. Norwood, NJ: Ablex.

Gentner D, Holyoak K and Kokinov B (eds) (2001) *The Analogical Mind: Perspectives From Cognitive Science*. Cambridge, MA: MIT Press.

Hall R (1989) Computational approaches to analogical reasoning: a comparative analysis. *Artificial Intelligence* **39**: 39–120.

Holyoak K and Barnden J (eds) (1994) *Advances in Connectionist and Neural Computation Theory*, **vol. II** 'Analogical Connections'. Norwood, NJ: Ablex.

Holyoak K, Gentner D and Kokinov B (eds) (1998) *Advances in Analogy Research: Integration of Theory and Data from the Cognitive, Computational, and Neural Sciences*. Sofia: New Bulgarian University Press.

Holyoak K and Thagard P (1995) *Mental Leaps*. Cambridge, MA: MIT Press.

Vosniadou S and Ortony A (eds) (1989) *Similarity and Analogical Reasoning*. New York, NY: Cambridge University Press.

Keywords: (Check)

computational model; psychology of analogical reasoning; case-based reasoning; connectionism; symbolic modeling; constraint satisfaction; emergent cognition; long-term memory; episode; tensor product

1. Please check coauthor details.
2. ‘Other Symbolic Models’, last sentence, ‘… which helps explain the effects of order of presentation of the material’: Do you mean it helps explain effects observed in humans? As these effects have not been mentioned, perhaps it would be better to say ‘… which models some of the effects of order of presentation of the material that have been observed in humans’?
3. Can you supply a glossary defining complex terms for the layperson?