Kokinov, B. and **French, R. M.** (2003) Computational Models of Analogy-making. In Nadel, L. (Ed.) *Encyclopedia of Cognitive Science. Vol. 1*, pp.113 - 118. London: Nature Publishing Group.

Computational Models of Analogy-Making

Boicho Kokinov
Central and East European Center for Cognitive Science
Department of Cognitive Science and Psychology
New Bulgarian University
Sofia, Bulgaria
email: bkokinov@nbu.bg

Robert M. French
Quantitative Psychology and
Cognitive Science
Psychology Department
University of Liege,
Liege, Belgium
email: rfrench@ulg.ac.be

Introduction

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, translating poetry from one language to another (Gentner, Holyoak, & Kokinov, 2001). All these cases 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 referred to as the target.

Analogy-making is a very basic cognitive ability which 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, ultimately, 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 and explains the importance of developing computational models of analogy-making.

Analogy-making involves at least the following sub-processes: representation-building, retrieval of a "base" for the analogy, mapping this base onto the "target," transferring of 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 schemas. There are, at present, no models that incorporate all these sub-processes, although individual models focus on one or, in some cases, several of these sub-processes.

Representation-building. This process is absent in most models of analogy-making. Typically hand-made representations are fed into the model. However, there are some models (ANALOGY, CopyCat, TableTop, MetaCat) that do produce their own high-level representations based on essentially unprocessed input. These latter models (Mitchell, 1993; Hofstadter, 1995; French, 1995) attempt to build flexible, context-sensitive representations during the course of the mapping phase. Other models, such as AMBR (Kokinov & Petrov, 2001), perform re-representation of old episodes.

Retrieval. This has been extensively studied experimentally and it is now clear that superficial similarity plays the major role in analogical retrieval, i.e. the retrieval of a base for analogy is easier if it shares similar objects, similar properties, similar general theme 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 schemas also facilitate retrieval. Most models of retrieval are based on exhaustive search of 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. This is unquestionably the core of analogy-making and, therefore, all computer 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 onto 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 individual 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 the relational structure. For example, we could decide that it was important to preserve the giver-receiver relationship in the first domain with 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 receiver was wearing a red shirt.

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 – people try to find mappings that involve the most important elements in the target. Searching for the appropriate correspondences between the base and target is a computationally complex task that can lead to combinatorial explosion if the search is unconstrained.

Transfer. This is the process of inserting new knowledge into the target domain based on the mapping. For example, assume a new type of car appears on the market and it turns out that this car maps well onto another type of car that is small, fast, and handles well on tight curves. But you also know that this latter type of car is frequently in need of repair. Transfer is when you wonder whether the analogous new model of car will also be in the garage often for repairs.

Transfer is present in one form or another in most models of analogy-making and is typically integrated with mapping. Transfer is considered by some authors as an extension of the mapping already established, thus adding new elements to the target in such a way that the mapping can be extended.

Evaluation. This is the process of establishing the likelihood that the transferred knowledge will turn out to 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 frequently be in need of repair. Evaluation is often implicit in the mechanisms of mapping and transfer.

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

In what follows we review a number of important computational models of analogy-making. An attempt is made to present models 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, episodes (“John”, “chair”, “run”, “greater-than”, “John ate fish”, “My birthday party last year”, etc.).Then the so-called “connectionist” models are presented. Here the objects, relations, and episodes are represented as overlapping patterns of activation in a neural network. Finally, a third class of models are presented which are of hybrid type – they 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.

Classical Symbolic Models

ANALOGY

The earliest computational model of analogy-making, ANALOGY, was developed by Thomas Evans (1964). This program solves geometric-analogy problems of the form A:B::C:? taken from IQ tests and college entrance exams.

An important feature of this program is that the input is not a hand-coded, high-level description of the problem, but, rather, 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 (e.g. ((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, e.g. ((MATCH P2 P4) (MATCH P1 P5) (REMOVE P3)) which means that the figure P2 from A corresponds to figure P4 from B, P1 to P5, and the figure P3 does not have a correspondent 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 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 aspects of the program is its ability to represent the target problem on its own — a feature that has unfortunately been dropped in most recent models.

Structure-Mapping Theory

Without question, 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 sound and convincing analogies. There are several important assumptions underlying the computational implementation of SMT called SME (Falkenhainer, Forbus & Gentner, 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 identical relations in both domains can be put into correspondence; 4) relations that are arguments of higher-order relations that can also be mapped have priority since they implement the "systematicity principle" that favors systems of relations over isolated relations; and 5) construction of two or three interpretations 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-universal nature of the mapping process. Recent versions of SME have explored 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 (Forbus, Gentner, & Law, 1995) of analogical retrieval was developed to be coupled with SME. This model assumes that episodes are encapsulated representations of past events which 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 frequency of predicates are used in the detailed representation). These representations are used in two sequential stages in the retrieval process. The first stage makes use of 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 then 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 simulate the dominance of superficial similarity on retrieval, but also the fact that retrieval is influenced by the structural similarity.

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

Other Symbolic Models

A number of other symbolic models have played a role in the advance of analogy-making understanding. 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. an analogy with the way of reaching up the solution is made, 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 that learned from multiple analogies combining several bases. Mark Keane and his colleagues developed an incremental model of mapping, IAM, which would explain the order effects in presentation of the material. These symbolic models, as well as a number of other early symbolic models of analogy-making are described in detail in Hall (1989).

Connectionist Models

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 addition, 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 latter ability is crucial to analogy-making and has proved hard for symbolic models to implement.

Multiple Constraints Theory

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 a representation of the target and one of the base and proceeds to build a localist constraint-satisfaction connectionist network where each node corresponds to a possible pairing hypothesis for each element of the base and each element of the target. So, 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 reaches 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 a model of retrieval that is coupled with ACME and operates in a similar fashion. However, while mapping is dominated by structural similarity, retrieval is dominated by semantic similarity.

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: MOTHER-OF⊗CAT⊗KITTEN. The tensor product in this case is a three-dimensional array of numbers where the number in each cell is the multiplication 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: MOTHER-OF ⊗ CAT • MOTHER-OF ⊗ CAT ⊗ KITTEN = 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 the model STAR-1 solves proportional analogies like CAT:KITTEN::MARE:?.

STAR-2 (Wilson, et al., 2001) maps complex analogies by sequentially focusing on various parts of the domains (simple propositions with no more than 4 dimensions) and finding the best map for the arguments of these propositions by parallel processing in the constraint satisfaction network (similarly to ACME). Since the number of units required for a tensor-product representation increases exponentially with the number of arguments of a predicate, this implies processing constraints in the model. The authors of the model 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.

LISA

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). In this way 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 it with synchronous oscillation of the patterns for “Mary” and “Employee”. Alternating the activation of the two pairs periodically in time makes it possible to represent 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 different pairs of synchronously firing concepts cannot exceed six. Representations in LISA’s Working Memory are distributed over the network of semantic primitives, but are localist in Long Term Memory – there are separate units representing the episode, the propositions, their subparts, 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

Two groups of researchers 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 both the symbolic and connectionist approaches. Semantic knowledge is incorporated in order to compute the similarity between elements of both domains in a context-sensitive way.

COPYCAT, TABLETOP, etc.

The family of COPYCAT and TABLETOP architectures (Mitchell, 1993; Hofstadter, 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:ABD::KLM:? and gives probabilistically possible answers like KLN, KLD, etc. 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 set of consistent set of structures that will determine the mapping between the base and the target.

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. In this way, the partial mapping can have an impact on 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, this being crucial to creative analogy-making.

AMBR

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

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 and/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 speed proportional to their computed relevance thus making the behavior of the system highly context-sensitive. The AMBR model implements the interactive parallel work of recollection, mapping and transfer which emerge from the collective behavior of the agents and whose work produces the analogy. Recollection in AMBR-2 (Kokinov & Petrov, 2001) is reconstruction of the base episode in WM 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 predicts illusory memories, including insertions from general knowledge and blending with other episodes as well as context and priming effects. A number of these predictions have been experimentally confirmed.

Conclusions

The field of computer-modeling of analogy-making has moved from the early models which were intended mainly as existence proofs to demonstrate that computers could, in fact, be programmed to do analogy-making to complex models which 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 building 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 (cf. Gentner, Holyoak, Kokinov, 2001). However, many issues still remain to be explored in the endeavor to model analogy-making, a capacity which is, without question, one of humans' most important cognitive abilities.

References

Evans, T. (1964). A Heuristic Program to Solve Geometric-Analogy Problems. In: *Spring Joint Computer Conference*, vol. 25. Reprinted in: M. Fischler & O. Firschein (eds.) *Readings in Computer Vision.* Morgan Kaufman Publ., 1987

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

Forbus, K., Gentner, D. & Law, K. (1995). MAC/FAC: A Model of Similarity-Based Retrieval. *Cognitive Science,* vol. 19 (2).

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., & Kokinov, B., eds. (2001) *The Analogical Mind: Perspectives from Cognitive Science.* Cambridge, MA: MIT Press.

Halford, G., Wilson, W., Guo, J., Gayler, R., Wiles, J., Stewart, J. (1994). Connectionist implications for processing capacity limitations in analogies. In: K. Holyoak & J.

Barnden (eds.) *Advances in connectionist and neural computation theory, vol. 2, Analogical Connections*, pp. 363--415. Norwood, NJ: Ablex.

Hall, R. (1989). Computational Approaches to Analogical Reasoning: A Comparative Analysis. *Artificial Intelligence*, vol. 39, pp. 39-120.

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

Holyoak K. & Thagard P. (1989). Analogical Mapping by Constraint Satisfaction. *Cognitive Science*, 13, 295-355

Hummel, J. & 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. & Barnden J. (eds.) *Advances in connectionist and neural computation theory, vol. 2, Analogical Connections*, pp. 247-318, Norwood, NJ.: Ablex

Kokinov, B., Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: Gentner, D., Holyoak, K., Kokinov, B. (eds.) *The Analogical Mind*, Cambridge, MA: MIT Press.

Mitchell, M. (1993) *Analogy-making as Perception: A computer model.* Cambridge, MA: MIT Press.

Wilson, W., Halford, G. Gray, B., Phillips, S. (2001). The STAR-2 Model for Mapping Hierarchically Structured Analogs. In: D. Gentner, K. Holyoak, B. Kokinov (eds.) *The Analogical Mind.* Cambridge, MA: MIT Press.