Analogy is like Cognition: Dynamic, Emergent, and Context-Sensitive

Boicho Kokinov

Department of Cognitive Science
New Bulgarian University
21, Montevideo Str.
Sofia 1635, Bulgaria

Institute of Mathematics and Informatics
Bulgarian Academy of Science
Bl. 8 Acad. G. Bonchev Str.
Sofia 1113, Bulgaria

kokinov@cogs.nbu.acad.bg

Abstract

This paper presents several challenges to the models of analogy-making, namely the need for building integrated models, the need for using dynamic and emergent representations, the need for using dynamic and emergent computation, and the need to integrate analogy-making with other cognitive processes. Some experimental data are reviewed which substantiate these needs and the main ideas how the AMBR model of analogy-making could meet these challenges are presented.

1. From the Anatomy towards the Physiology of Analogy-Making: The Need for Integrated and Dynamic Models

For a long time now the research on analogy has concentrated on the anatomy of analogy-making, i.e. on decomposing it into pieces (representation building, retrieval, mapping, transfer, evaluation, learning) and trying to understand how each individual piece works. A number of successful models of various subprocesses (mainly of mapping and retrieval) have been built which account for most of the psychological data and make useful predictions: SME and MAC/FAC (Gentner, 1983, Falkenheiner, Forbus, Gentner, 1986, Forbus, Gentner, Law, 1995), ACME and ARCS (Holyoak, Thagard, 1989, Thagard, Holyoak, Nelson, Gochfeld, 1990, Holyoak, Thagard, 1995), IAM (Keane, Ledgeway, Duff, 1994), etc.

The big challenge in modeling analogy-making (and human cognition in general) is to move on from the atomistic and analytical approach of Democritus (469-370 BC) towards the holistic and interactionist approach of Heraclitus (544-481 BC), i.e. to start building integrated models of the phenomenon as a whole. These models should unite contraries and account for data arising from the interaction between subprocesses, which cannot be explained by an isolated model of a subprocess. Such models are gradually emerging. Thus the CopyCat and TableTop models (Hofstadter, 1995, Mitchell, 1993, French, 1995) integrate representation building with mapping and transfer, LISA (Hummel and Holyoak, 1997) integrates access, mapping, transfer, and learning, AMBR (Kokinov, 1988, 1994c) integrates access, mapping, and transfer.

Heraclitus took the view that “Everything flows, everything changes”, i.e. the dynamics of change is more important and informative than static objects and states. This is the next challenge to the current models: they should explain and predict not only the outcomes of the analogy-making process but also its dynamics. Unfortunately, only scare data is available on the dynamics of the process. This means that such data will have to be gathered by using experimental paradigms extensively used in other domains, for example, on-line experiments measuring reaction times, analysing thinking-aloud protocols, etc. These methods have already been used in analogy research but to a very limited extent (Ross and Sofka, 1986, Keane, Ledgeway, and Duff, 1994, Schunn and Dunbar, 1996).

There are already experimental data which support the existence of interaction effects between the subprocesses of analogy-making. Thus Keane, Ledgeway, and Duff (1994) have demonstrated a very strong ordering effect, i.e. effect of the order of presentation of the target problem elements on the response time for solving the problem. Thus in the "singleton-first" condition subjects found the mapping twice as fast as subjects in the "singleton-last" condition. These data can be considered as evidence for the interaction between perceptual and mapping processes. It would be even more interesting to find the reverse patterns: the mapping already established facilitating the perception of certain elements.

The analysis of thinking-aloud protocols done by Ross and Sofka (1986) revealed that the retrieval of various elements of the source domain is interrelated with the mapping between the two domains, i.e. the already established mappings guide the retrieval of specific source elements. These data cannot be explained by a serial model of analogy-making where first the source is being retrieved and then the source and target are mapped. An extensive discussion of this phenomenon and its modeling in AMBR as well as simulation data obtained with AMBR can be found in (Petrov, Kokinov, this volume). AMBR predicts also the reverse influence: the specific order of retrieval of elements of the source domain will facilitate certain mappings. As a result of these interactions, a pattern of retrieval has been demonstrated where initially one source domain looks more promising and is better retrieved based on the greater superficial similarity, but as soon as mapping starts (in parallel to the continuing retrieval of domain elements), the higher structural correspondence between a second source domain and the target and the established mappings make it possible for the second domain to be ultimately better retrieved and mapped which would be impossible if the retrieval and mapping were sequential isolated and irreversible processes.

Finally, a study currently underway involves video recording of subjects solving a formatting task on a computer screen. The video protocols demonstrate a complex interaction between perceiving elements on the screen (including figure/background perception), retrieving elements from memory, mapping between these elements, and performing actions on the screen, the results of which are further perceived and mapped to expectations.

The explanation of all these data requires models which abandon the serial type of processing and which move on towards parallel processing which will allow the various subprocesses to interact dynamically with each other. AMBR is one such model that is based on the highly parallel cognitive architecture DUAL (Kokinov, 1994a, 1994b). All processes in AMBR are running in parallel and interacting with each other. Moreover, as described in section 3, each of these subprocesses emerges from the collective behavior of many micro-agents and thus is also inherently parallel. Since the micro-agents are taking part in various subprocesses there are no clear-cut boundaries between the various processes themselves.

Before the dynamics of computation in AMBR can be presented, the need for dynamic representations that will change in the course of analogy-making will be discussed in the next section.

2. From Printed Text towards Moving Picture: The Need for Dynamic and Emergent Representations

A printed text is a static representational object while a moving picture is a dynamic representation which emerges from the continuously changing frames. Moreover, this dynamic representation does not exist physically (only the static frames exist physically), it exists only in our consciousness. Analogously, memory traces may be considered either as physically existing static entities, or as emergent phenomena which are constructed in our consciousness.

From the very beginning of memory research the view of memory as consisting of stable representations has been under fire. Thus Bartlett (1932) has shown that episodes are grouped into schemas and their representations are systematically shifted or changed in order to fit these schemas. Research on autobiographical memory has provided evidence that people modify their memories by dropping elements (schematising), including new elements (filling in), replacing elements (distorting), etc. Loftus (1977, 1979) has convincingly demonstrated a number of interference effects. One example involves subjects looking at a movie where a blue car does not stop at the site of an accident. Later on in a questionnaire a number of questions are asked about a different green car. As a result, when asked about the color of the car which did not stop, subjects are quite confident that it was green. In another study subjects claim they have seen broken glass in a car crash whereas there was no broken glass in the movie shown to them.

Neisser and Harsch (1992) have demonstrated that the so-called "flash-bulb memory" does not exist but that descriptions constructed by human memory are so vivid that people strongly believe they are true. One day after the *Challenger* accident they asked subjects to tell them (and write down) how they learnt about the accident: whether they heard it on the radio, or saw it on TV, or learnt it on the street, in the supermarket, from friends. They asked further the subjects in the study what they were doing when they learnt about the accident, what their reactions were, etc. One year later the experimenters asked the same subjects whether they still remember the accident and how they learnt about it. People claimed they had very vivid ("flash-bulb") memories about every single detail and they started to tell the experimenters a very different story from the one they told before. Even after the experimenters showed them their own writings they could not believe that the new story they were telling the experimenters was not true.

Although it has long been demonstrated that human memory is a (re)constructive device rather than a store of stable memory traces from our past, models of analogy-making tend to ignore that fact. Typically these models would have a collection of representations of past episodes (prepared by the author of the model) "stored" in long-term memory (LTM), one or more of which would be "retrieved" during the problem solving process and would serve as a base (or source) for analogy. The very idea of having singular centralized and frozen representations of base episodes is at least questionable, but it underlies most analogy-making models, and certainly all case-based reasoning systems (Figure 1).

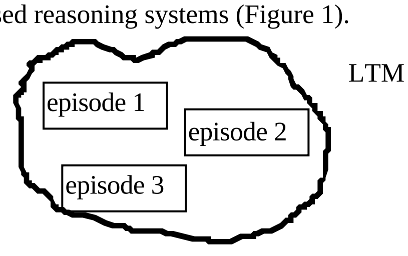

Figure 1. Centralized and frozen representations of episodes in LTM.

Research on retrieval in analogy-making has concentrated on how people select the most appropriate episode from the vast set of episodes in LTM. It has been established that the existence of similar objects, properties or relations in the two domains is the crucial factor for retrieval (Holyoak & Koh, 1987, Ross, 1989) and that is why remote analogies are very rare. On the other hand, structural similarities can also facilitate retrieval under certain circumstances, when there is a general similarity between the domains or story lines (Ross, 1989, Wharton, Holyoak, Lange, 1996). There is not much research either on the dynamics of the process of retrieving (or constructing), or on how complete the resulting descriptions of the episodes are.

A recently conducted experiment was designed as a replication of Holyoak and Koh's (1987) Experiment 1. However, a thinking-aloud method was used. Subjects discussed the solution of the radiation problem in a class on thinking within an introductory Cognitive Science course. From 3 to 7 days later they were invited by different experimenters to participate in a problem-solving session in an experimental lab. They had to solve the light bulb problem. Almost all subjects (except one who turned out not to have attended the class discussing the tumor problem) constructed the convergence solution and explicitly (in most cases) or implicitly made analogies with the radiation problem. We were interested how complete and correct their spontaneous descriptions of the tumor problem story were. It turned out that remembering the radiation problem is not an all-or-nothing case. Different statements from the story were recollected and used with varying frequency. Thus the application of several X-rays on the tumor was explicitly mentioned by 75% of the 16 subjects participating in the experiment, the statement that high intensity rays will destroy the healthy tissue was mentioned by 66% of the subjects, while the statement that low intensity rays will not destroy the tumor was mentioned only by 25%. Finally, no one mentioned that the patient would die if the tumor was not destroyed. All this demonstrates a partial retrieval of the base: which elements of the base will be retrieved depends on the pragmatically important aspects of the target problem.

On the other hand, there were some insertions, i.e. "recollections" of statements that were never made explicit in the source domain description. Thus one subject said that the doctor was an oncologist which was never explicated in the radiation problem description (nor should it be necessarily true). Another subject claimed that the tumor had to be burnt off by the rays, which was also never formulated in that way in the problem description.

Finally, there were borrowings from other possible bases in memory: thus one subject said that the tumor had to be "operated by laser beams" while in the base story the operation was even forbidden. Such blendings were very frequent between the base and the target, thus 7 out of the 11 subjects spontaneously re-telling the base (the radiation) story were mistakenly using laser beams instead of X-rays to destroy the tumor. This blending is evidently the result of the correspondence established between the two elements and their high similarity.

In summary, the experiment has shown that remindings about the base story are not all-or-nothing events and that subjects make omissions, insertions, and blendings with other episodes.

The representation of episodes in AMBR is de-centralized, which means that separate elements of the episode's description are represented by separate memory elements (called micro-agents in the DUAL cognitive architecture). Thus the episode as a whole is represented by a coalition of agents, but there is no guarantee that the whole coalition will be activated and become part of WM. Depending on the weights of the links between the agents the coalition might be looser or tighter. This makes it possible to model the above mentioned psychological effects. Thus very often only part of the agents in a coalition are being activated above the Working Memory (WM) threshold and thus the corresponding episode is only partially retrieved. Depending on the retrieval cues used various partial recollections will be produced.

Blendings also happen in AMBR. Thus agents representing aspects of several different episodes can be concurrently activated in WM. Mappings between elements of the target and elements of all partially retrieved episodes can be established in parallel and compete with each other. Typically the support that the agents in one coalition receive from each other is enough to achieve a global emergent "winner" episode. However, in some cases one or more aspects needed for the mapping (having counterparts in the target) are missing in the representation of an episode, or are not retrieved in WM, but instead corresponding elements from other episodes are retrieved. In such a case a blending between the episodes can happen, i.e. the target elements are partially mapped to elements of one base and partially to elements of another base (Figure 2).

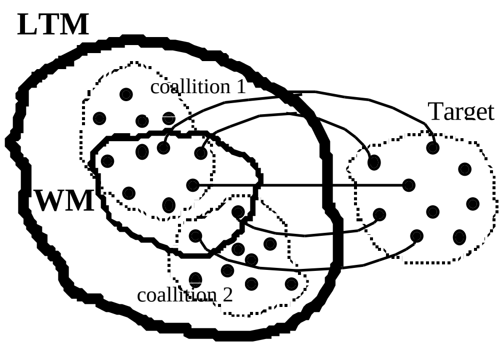

Figure 2. Blending of two episodes (represented by two coalitions) which are partially retrieved in WM and partially mapped on the target coalition. (The target coalition is also part of WM, but is depicted separately for simplicity of the diagram.)

Finally, insertions (analogous to the doctor-oncologist case) are also possible in AMBR. Semantic knowledge is represented in a similar decentralized fashion, i.e. different aspects of a concept are represented by different agents. Suppose, for example, that there is a general rule saying that liquids are typically held in containers. Suppose now that an episode is being retrieved in which water is heated by an immersion heater. It might well be the case that the fact that the water was in a glass was either not encoded at all, or was not retrieved under the current circumstances. At the same time the target situation involves tea being heated in a pot on a plate. The agent representing the fact that the tea is in the pot will activate many agents representing similar facts and in particular the one representing liquids being in containers. If during the mapping process a correspondence is attempted between those agents: IN(TEA1, POT1) and IN(LIQUID, CONTAINER), then instead of building a correspondence hypothesis between them, a new agent is being built which represents a skolemized version of the general statement, namely IN(WATER1, CONTAINER1) and a correspondence hypothesis between it and IN(TEA1, POT1) will be formed. In this way the mapping process guided the process of extending the representation of the old episode, thus producing a new richer representation with inclusions, such as IN(WATER1, CONTAINER1).

In summary, AMBR dynamically forms the representation of old episodes by selecting only some of the encoded aspects of the episode (hopefully the relevant ones), and by adding new aspects which have not been explicitly encoded from beforehand – this is done either as skolemized versions of more general facts, or by borrowing facts from other episodes (blending).

The specific mechanisms proposed in AMBR for re-representation of old episodes might be psychologically valid or not, but the very fact that such dynamic re-representations are being made by humans has been shown to be valid above. Another important aspect is that this re-representation in AMBR is a result of the interplay of memory retrieval (determining which agents will be brought into WM), mapping (determining which agents are unpaired), and deductive reasoning (skolemization) and could not be realised if they were not running in parallel and interacting with each other. Finally, as I will discuss in the next section, all these complicated processes of re-representation and mapping are performed using only local information, i.e. each individual agent decides which links to establish, which new agents to form, etc.

3. From Centralized Planning towards Free Market: The Need for Dynamic and Emergent Computation

Adam Smith is not only the most famous economist who introduced the theory of the free market as a regulator of the economy and was against any form of governmental control over the market. In his book "An Inquiry into the Nature and Causes of the Wealth of Nations" (Smith, 1776) he also introduced the idea of emergent phenomena in the social sciences. He wrote about "the invisible hand by which man is led to promote an end which was not part of his intention". Thus when someone decides to start the production of certain goods in an area where the rate of profit is very high he/she does it in order to gain this high profit, however, since many will do the same, this will result in declining prices and eventually decreasing the rate of profit in this area which was in no way a goal of the producers, but they have achieved it by their actions. Von Hayek (1967), another famous economist, proclaimed that finding an explanation of the mechanisms of these emergent phenomena is the main task of the social sciences: "those unintended patterns and regularities which we find to exist in human society and which it is the task of social theory to explain".

Some human societies were tempted to find a more direct and faster way to achieve a balance in their economy – why wait till the free market regulates prices and production when the government could calculate the desired prices and amounts of production in every economic area and directly postulate them. These attempts have recently collapsed completely. Why? The problem is that economic systems are too complex to be directly controlled and what seems to be "the more efficient direct way" is actually a very rigid way that cannot be flexible enough to reflect dynamic changes in the environment.

Cognitive scientists are gradually learning the same lesson. The attempts to build a model of human cognition based on a centralized control system are doomed to failure. No such system could be flexible enough to adapt to all dynamic changes in the environment and to reflect all possible human goals. Such a system is inherently rigid as it reflects the tasks and circumstances envisaged by its designer. An alternative approach has been proposed by Marvin Minsky (1983) which is based exactly on the analogy with human societies and has been called "the society of mind". Another alternative is the connectionist approach based on the analogy with human neural networks.

Nevertheless, we are still trying to build models of analogy-making based on the assumption that the solution of a problem is determined by its formulation and the knowledge background (including previous solutions to other problems) the subject has. Several examples of *context effects* are presented here which demonstrate that analogy-making is not that simple and predictable.

Kokinov and Yoveva (1996) conducted an experiment on problem solving where seemingly irrelevant elements of the problem solver's environment were manipulated. The material manipulated consisted of drawings accompanying other problems which happened to be printed on the same sheet of paper. There was no relation between the problems and the subject did not have to solve the second problem. However, these seemingly irrelevant pictures proved to play a role in the problem solving process as we obtained different results with different drawings. We used Clement's spring problem:
"Two springs are made of the same steel wire and have the same number of coils. They differ only in the diameters of the coils. Which spring would stretch further down if we hang the same weights on both of them?"

The problem description was accompanied by the following picture.

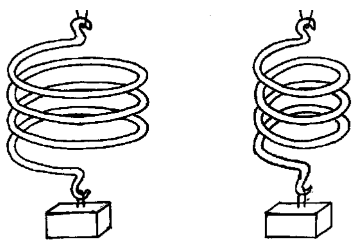

Figure 3. Illustration accompanying the target problem.

In different experimental conditions the drawings used as accompanying a second unrelated problem on the same sheet of paper were different: a comb, a bent comb, and a beam (Figure 4).

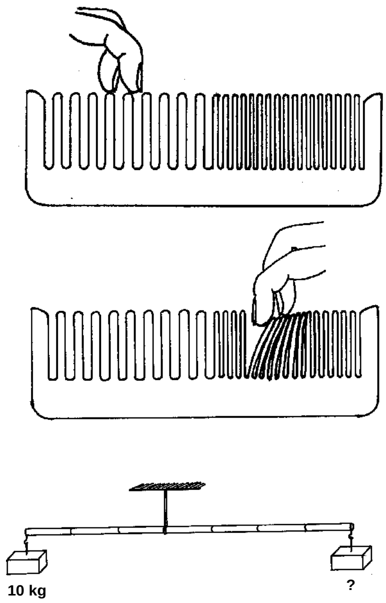

Figure 4. Illustrations accompanying the irrelevant problems in the various experimental conditions.

The results obtained in these experimental conditions differed significantly (at the 0.01 and 0.001 levels): in the control condition (no second picture on the same sheet of paper) about half of the subjects decided that the first spring will stretch more and the other half 'voted' for the second one, with only a few saying they will stretch equally. In the comb condition considerably more subjects suggested that the first spring will stretch more, in the bent comb condition considerably more subjects preferred the second spring, and in the beam condition more subjects than usual decided that both springs will stretch equally (Figure 5).

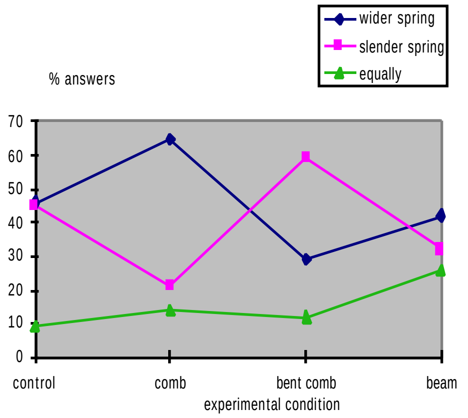

Figure 5. Percentage of proposed answer in all the experimental conditions.

In a more recent study (the thinking-aloud experiment described in section 2) the subjects who had to solve the lightbulb problem were divided into two groups. In the control group there were no other problems on the sheet of paper, in the context group the following problem was presented on the same sheet.

"The voting results from the parliamentary elections in a faraway country have been depicted in the following pie-chart. Would it be possible for the largest and the smallest parties to form a coalition which will have more than 2/3 of the seats?"

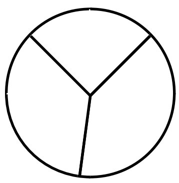

Figure 6. Illustration accompanying the context problem.

The results are the following: in the context group *all* 7 subjects who produced the convergence solution to the lightbulb problem used *three* laser beams (7:0), while in the control group two subjects said they would use *two or three* beams and the rest said they would use either *two* or *several* beams (2:5). The difference is significant at the 0.01 level.

The results from both experiments demonstrate that sometimes small changes of a seemingly arbitrary element of the environment can radically change the outcomes of the problem solving process (can block it, or guide it into a specific direction).[1] Such phenomena are called "catastrophes". It would be very difficult to account for such effects by a model based on centralized control because in order to do so the centralized processor would have to process all possible stimuli in the perceptual field and to check whether they can be involved in the problem solving process, which would be inefficient and time-demanding to such an extent as to make it impossible.

The AMBR model adopts the following approach to accounting for context effects. It assumes that different micro-agents process different aspects of the problem and the environment. If it happens that one agent processing an arbitrary and seemingly irrelevant visual stimulus enters an interaction with a second agent processing a relevant problem aspect, then the first agent will be additionally activated and become more relevant and thus involved in the collective process of problem solving performed by the society of agents. This is a very brief and simplified description of what happens in the model, a detailed description would be based on the specific mechanisms of spreading activation, marker passing, link establishment and between-agent communication which are too complicated to be outlined in the limited space of this article.

Another important aspect of analogy-making which makes it difficult to predict whether the

---

[1] This is analogous to the following phenomenon in economy – the bankruptcy of a single bank can trigger off a chain of bankruptcies and eventually a global financial crisis.

subject will be able to spontaneously find an analogous base (which we know he/she knows) is that this process depends on his/her preliminary internal state which is typically not related to the current problem, but is related to recently performed activities. Thus Kokinov (1990) demonstrated *priming effects* on analogical problem solving (as well as on other types of reasoning) which have a very dynamic nature, namely they are very powerful immediately after the priming event and decrease in the course of time and eventually disappear after a short period of time (in this particular study within a period of about 25 minutes). These priming effects have been qualitatively reproduced by a previous version of the AMBR model based on the pre-activation of certain agents and the decay of their activation in the course of time (Kokinov, 1994c). We plan to reproduce these priming effects with the new version by running it continuously thus solving various problems one after another.

The main conclusion from the considerations in this section is that in order to build adequate models of analogy-making, we need to base them on massively parallel architectures allowing the parallel work and interaction of many small processing entities. In addition the architecture should allow for dynamic short-term changes in the structure of interactions between these entities, something that current connectionist models do not allow.

AMBR and the underlying cognitive architecture DUAL are definitely not the best solution to these requirements. For example, top-down pressure ("the invisible hand" of the context) is limited to the current distribution of activation over agents which facilitates the local communication between agents in one direction and inhibits it in another, supports certain coalitions of agents and suppresses others. It is doubtful that this would be enough to explain all context and priming effects. On the other hand, CopyCat and TableTop have one additional top-down pressure which is called "temperature" and reflects an internal evaluation of the mental state and how close the system is to the solution of the problem. A problem with this approach is that it assumes the existence of a centralized agent watching the whole situation, computing the temperature and then communicating it back to all agents – this resembles again centralized "government" control, although it is weak control – it does not specify what the agents should do, but only changes their biases and thresholds.

The next question to be discussed in the last section is whether the mechanisms performing analogy-making can be considered domain-specific and thus form something that several researchers have called an analogy-making engine.

4. From a Specialised Engine towards an Emergent Phenomenon: Integrating Analogy with other Cognitive Processes

If analogy-making is modeled within a highly parallel architecture of "the society of mind" type, then there is no need to assume that there are mechanisms or agents which are so specific that are solely used for analogy-making. On the contrary, the analogy-making process would be considered as an emergent phenomenon, i.e. that is how we describe certain types of emergent behavior produced by the society of agents. AMBR, for example, uses mechanisms like spreading activation, marker passing, etc. which in no way may be considered as specific for analogy-making. Spreading activation, in particular, is involved in all memory processes; marker passing is involved in the processes of evaluating semantic similarity, categorization, directed search, property inheritance, etc. A process that might seem more specific for analogy-making is the ability of agents to establish hypotheses for structure correspondence (i.e.

correspondence between substructures), however, this process seems so fundamental that it is doubtful that it is specifically designed for analogy-making – all processes of perception would need some structure correspondence abilities, all relational processing would also require this ability.

If we subscribe to the “emergent phenomenon” view on analogy, then it would be natural to integrate it with all other cognitive processes – simply they are emerging from the collective behavior of the same micro-agents. Then the boundaries between analogy-making, perception, memory, deductive reasoning, etc. can be described as conventional – as classification of various types of collective behavior of the same set of agents and produced by the same mechanisms (probably in different proportions). Thus Kokinov (1988, 1990, 1994c) has argued that the boundaries between analogy, deduction and generalization are a convention and that these processes are implemented by the same mechanisms. Of course, this is yet only one unsubstantiated hypothesis.

This paper is probably too general and full of speculation, however, its purpose has been neither to describe AMBR in details (which is not possible because of space limitations), nor to defend its basic principles. I am fully aware of the fact that these principles express only one possible point of view on modeling analogy-making. The purpose is to present some challenges to current models of analogy-making as seen by the author and to suggest possible ways of meeting them hoping to combine these ideas with other views expressed during the workshop.

References

Bartlett, F. (1932). Remembering. Cambridge: Cambridge Univ. Press.

Falkenhainer, B., Forbus, K., and Gentner, D. (1986). The structure-mapping engine. *Proceedings of the Fifth Annual Conference on Artificial Intelligence*. Los Altos, CA: Morgan Kaufman.

Forbus K., Gentner D., and Law, K (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19*, 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*, 155-170.

Hofstadter, D. and the Fluid Analogies Research Group (1995). *Fluid concepts and creative analogies: Comuter models of the fundamental mechanisms of thought.* New York: Basic Books.

Holyoak K. and Koh K. (1987). Surface and structural similarity in analogical transfer. *Memory and Cognition, 15* (4), 332-340.

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

Holyoak, K. and Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press.

Hummel, J. and Holyoak, K. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review, 104*, 427-466.

Keane, M., Ledgeway, K., and Duff, S. (1994). Constraints on analogical mapping: A comparison of three models. *Cognitive Science, 18*, 387-438.

Kokinov, B. (1988). Associative memory-based reasoning: How to represent and retrieve cases. In T. O'Shea and V. Sgurev (Eds.), *Artificial intelligence III: Methodology, systems, applications.* Amsterdam: Elsevier.

Kokinov, B. (1990). Associative memory-based reasoning: Some experimental results. *Proceedings of the Twelfth Annual Conference of the Cognitive Science*

Society. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kokinov, B. (1994a). The context-sensitive cognitive architecture DUAL. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale,NJ: Lawrence Erlbaum Associates.

Kokinov, B. (1994b). The DUAL cognitive architecture: A hybrid multi-agent approach. *Proceedings of the Eleventh European Conference of Artificial Intelligence.* London: John Wiley & Sons, Ltd.

Kokinov, B. (1994c). A hybrid model of reasoning by Analogy. In K. Holyoak and J. Barnden (Eds.), *Advances in Connectionist and Neural Computation Theory*. *Vol. 2: Analogical Connections*. Norwood, NJ: Ablex Publishing Corp.

Kokinov, B., Yoveva, M. (1996). Context Effects on Problem Solving. In: *Proceedings of the 18th Annual Conference of the Cognitive Science Society*. Erlbaum, Hillsdale, NJ.

Kokinov,B., Nikolov,V., and Petrov,A. (1996). Dynamics of emergent computation in DUAL. In A. Ramsay (Ed.), *Artificial Intelligence: Methodology, Systems, Applications*. Amsterdam: IOS Press.

Loftus, E. (1977). Shifting Human Color Memory. Memory and Cognition, vol. 5, 696-699.

Loftus, E. (1979). Eyewitness Testimony. Cambridge, MA: Harvard Univ. Press.

Minsky, M. (1986). *The society of mind*. New York: Simon and Schuster.

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

Neisser, U. & Harsch, N. (1992). Phantom Flashbulbs: False Recollections of Hearing the News about the Challenger. In: Winograd, E. & Neisser, U. (eds.) Affect and Accuracy in Recall. NY: Cambridge Univ. Press

Ross, B. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 456-468.

Ross, B. and Sofka, M. (1986). [Remindings: Noticing, remembering, and using specific knowledge of earlier problems]. Unpublished manuscript.

Schunn, C. and Dunbar, K. (1996). Priming, analogy, and awareness in complex reasoning. *Memory and Cognition, 24*, 271-284.

Smith, A. (1776). Inquiry into the Nature and Causes of the Wealth of Nations.

Thagard, P., Holyoak, K., Nelson, G., and Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence, 46*, 259-310.

Von Hayek, F. (1967). Studies in Philosophy, Politics, and Economics. London: Routledge & Regan Paul.

Wharton, C., Holyoak, K., and Lange, T. (1996). Remote analogical reminding. *Memory and Cognition, 24* (5), 629-643.