Dynamic Extension of Episode Representation in Analogy-Making in AMBR

Boicho Kokinov (bkokinov@nbu.bg)
Central and East European Center for Cognitive Science
Department of Cognitive Science and Psychology
New Bulgarian University
21 Montevideo St, Sofia 1635, Bulgaria

Alexander Petrov (apetrov+@andrew.cmu.edu)
Central and East European Center for Cognitive Science
Department of Cognitive Science and Psychology
New Bulgarian University
21 Montevideo St, Sofia 1635, Bulgaria

Abstract

Models that rely exclusively on static representations cannot account fully for the flexibility of human analogy-making. More sophisticated models should provide mechanisms for dynamic extension, elaboration, and re-representation of episodes. One such mechanism—the *instantiation* mechanism—is described. It uses the target problem as a template for extending the source and vice versa. These extensions are driven and constrained by semantic knowledge about general regularities in the domain. The instantiation mechanism has been developed within a model of analogy-making called AMBR. It relies on AMBR's support for parallel, decentralized, and interactive computation. The instantiation mechanism runs in parallel with the mechanisms for analog access and mapping. Thus, these latter mechanisms guide the instantiation mechanism as to which facts in the large semantic memory are relevant to the specifics of the current situation.

Re-Representation in Problem-Solving

A substantial body of evidence suggests that people can change their mental representations dynamically during various cognitive tasks. Yet, despite the widespread agreement that human representations are dynamic and flexible, the mechanisms behind these re-representation abilities are not well explored and understood. This paper suggests some mechanisms that can serve that purpose.

In the context of problem solving, there are at least two complementary aspects of re-representation: re-representation of the target problem and re-representation of prior knowledge. These two aspects correspond to Jean Piaget's complementary and related processes of assimilation and accommodation. Re-representation of the target is a process of assimilation because the new information is transformed to comply with the existing knowledge. Conversely, re-representation of the existing knowledge in the face of new experience is a process of accommodation.

Re-representation of the target problem has received more attention although it is still not fully understood. Gestalt psychologists (Maier, 1931; Dunker, 1945) demonstrated the importance of dynamic changes of the target representation for successful problem solving. However, the mechanisms behind this remained unclear and somehow mysterious. Unfortunately, contemporary cognitive science has not been particularly successful to fill this gap, although some progress has been made. Douglas Hofstadter and his group have worked for many years on the integration of analogy-making and what they call high-level perception (Hofstadter, 1995; French, 1995; Mitchell, 1993; Chalmers, French & Hofstadter, 1992). They have proposed a number of mechanisms that work together to build several alternative representations simultaneously, settle gradually on one of them, and radically restructure the representation and settle on an alternative one if necessary. Lange & Wharton (1992) have worked on a similar problem—integrating language comprehension with analogical reminding. They suggested a mechanism for parallel processing of several possible interpretations of an ambiguous phrase. This mechanism would allow for re-interpretation of the phrase if necessary.

The issue of ***re-representation of existing knowledge*** during the problem solving process has been systematically ignored in the problem solving literature. Unfortunately, most classical models of problem solving take the existing knowledge as constant[1]. Long-term memory structures only have to be retrieved and applied (Ernst & Newell, 1969; Newell, 1990; Anderson, 1983; Anderson & Lebiere, 1998). The same is true for models of analogy-making: they retrieve a ready-made representation of an old episode when looking for a base for analogy (Gentner, 1989; Thagard, Holyoak, Nelson & Gochfeld, 1990; Kokinov, 1994a; Forbus, Gentner & Law, 1995; Hummel & Holyoak, 1997).

The idea that there are ready-made representations of episodes formed during the encoding stage has been challenged by researchers of human memory for quite some time. There is much evidence for the constructive nature of the "retrieved memory traces". Thus Loftus (1977, 1979) and Neisser and Harsch (1992) demonstrate that people can have vivid "memories" of nonexistent episodes which are clearly constructed during the "retrieval" process. Bartlett

[1] We do not discuss learning here—it involves gradual long-term knowledge change rather than short-term accommodation.

Petrov & Kokinov (2000). ... In "Proceedings of the 22nd Annual Conference of the Cognitive Science Society" (pp.247-279). Hillsdale, NJ: LEA. http://alexpetrov.com/pub/ambr00/

(1932) has shown that episode representations are distorted and enriched with information inherited from the schema for the typical event.

The purpose of this paper is to show how episodic representations can be dynamically extended with information derived from general semantic knowledge. AMBR is among the few models of analogy-making which use both semantic and episodic knowledge (Kokinov, 1994a). However, the version reported in (Kokinov, 1994a) used them for different purposes. Episodic memory was a repository of ready-made representations of episodes that might be used as bases for analogy-making. Semantic memory served only to establish semantic similarity between the relations, attributes, and objects participating in the episodes. In the new version of AMBR described in this paper, the boundaries between semantic and episodic memory are blurred and general statements from semantic memory are instantiated to complement and extend the representations of episodes.

Main Assumptions of the AMBR Approach to Episode Re-representation

Flexible Dynamic Representations

The first assumption is that episodes and concepts are not represented by fixed and static complex memory structures such as schemas, lists of propositions or lists of rules. Rather, there are fuzzy and overlapping *coalitions* of simple memory elements. The key distinction is that complex memory structures are retrieved in an all-or-none fashion, while flexible dynamic representations are retrieved and/or constructed element by element. Each element has an activation level associated with it. Depending on the specific pattern of activation over the coalition, various partial representations of the same episode or concept can be retrieved. Some elements of the coalition might be strongly connected with each other and thus tend to be retrieved together, while others might be retrieved only rarely and under specific retrieval conditions. There is also a possibility for *blends* to emerge. This happens when elements belonging to more than one coalition become active together. All this flexibility is important for making re-representation possible.

Flexible Dynamic Computation

The second assumption is that computations are also flexible, parallel, and interactive. There are multiple processes running in parallel and interacting in complex ways that are not specified in advance. Thus, re-representation results from the interplay of many processes including: (*i*) retrieval of past episodes, (*ii*) retrieval of generic knowledge, (*iii*) instantiation of the generic rules or facts, (*iv*) attempt to build the representation of the past episode in a form that makes it alignable with the representation of the problem at hand, and (*v*) attempt to build the representation of the target problem in a form that makes it alignable with the representation of the past episode. As stated earlier, it is believed that all these processes should run in parallel and influence each other's work. In addition, each of the processes described above is quite complex in itself and in turn has to be considered as the result of the interplay of many simpler and more local processes. These requirements lead to a view that complex computation is an emergent interactionist phenomenon rather than pre-specified sequence of algorithmic steps.

Integrated Semantic and Episodic Memory

Semantic memory contains information about concepts and statements about classes of instances. Episodic memory contains information about instances, episodes, and statements about instances. There has been a long-lasting discussion for and against the distinction between these two memories (see, for example, Anderson & Ross, 1980, Herrmann & Harwood, 1980). The third assumption of AMBR is that semantic and episodic memories are integrated so as to allow for coordinated search in both memories. Whenever a cue is provided, both semantic and episodic memory elements can potentially be retrieved. Thus, when a past episode is *recalled* from memory, both specific and general knowledge is used in the recollection process. In this way semantic knowledge can extend the episodic knowledge.

Integration of Memory and Reasoning

The process of re-representation requires that the process of memory access, on one hand, and the processes of mapping and instantiation, on the other, run in parallel and interact with each other. Thus the fourth assumption is that memory and reasoning are highly integrated.

AMBR: An Analogy-Making Model Based on the Cognitive Architecture DUAL

An analogy-making model with re-representation capabilities needs the support of a full cognitive architecture that implements all the assumptions above. The cognitive architecture DUAL is specifically designed to support this decentralized and interactive style of computation (Kokinov, 1994b, 1994c, 1997). The AMBR model of analogy-making (Kokinov, 1994a) is based on this architecture. This paper describes the re-representation extensions that have been added to the model after the original publication. Before explaining how the assumptions are implemented and how they contribute to re-representation, a brief and more general description of DUAL and AMBR is needed.

DUAL is a cognitive architecture based on the *society of mind* idea (Minsky, 1986; see also Hofstadter, 1995). Every DUAL-based system consists of many micro-agents, each of which is quite simple. The micro-agents do not have goals and do not plan their activities; they are simple representation and computation devices. They can establish new links with other agents and some of them can construct new agents. DUAL-agents form coalitions that collectively represent an episode or a generalized concept, or dynamically form coalitions that collectively produce an emergent computation process. Each agent can participate in many coalitions to a various extent depending on the weights of the links connecting that agent to other agents in the coalition.

Knowledge representation in DUAL is highly decentralized. Each episode, concept, general theory, etc. is represented by a coalition of many agents, each of which represents just a small piece of knowledge. Thus a simple episode such as boiling water in the kitchen would be represented by a quite big coalition of agents: an agent for every *concept* related to the situation such as “water”, “kitchen”, “boiling”, “plate”, “pot”, “on”, “in”, “hot”, “cold”, “cause”; an agent for every *instance* of these concepts involved in the particular situation, i.e. “water-1”, “kitchen-3”, “boiling-2”, “plate-3”, “pot-3”; as well as for every single statement such as “on-1(pot-3,plate-3)”, “in-1(water-1,pot-3)”, “hot-1(plate-3)”, “red-1(pot-3)”, etc. However, it should not necessarily be the case that all elements of a coalition become members of the working memory (WM) at certain moment. On the contrary— typically only part of the coalition is activated. Thus each episode is almost always only partially available. Moreover, different subsets of the coalition are active in different contexts. The long-term memory (LTM) of DUAL is the population of all permanent agents, active or inactive. The working memory is simply the active part of LTM plus some newly created temporary agents.

Each agent is a DUAListic computational and representational device: it has a symbolic and a connectionist part. While the symbolic part represents a piece of knowledge (as described above), the connectionist part represents the relevance of this piece of knowledge to the current context. The relevance is represented by the graded activation level computed by the connectionist processor associated with the agent. All the inferences based on the knowledge represented by the agent are computed by the symbolic processor associated with the same agent. These computations are also based only on local interactions with neighboring agents. If necessary, the agents are able to establish new temporary links (and interactions) with other agents. The speed of symbolic processing of a given agent depends on the activation level. In this way the computations are faster if the corresponding knowledge structures are considered relevant to the context and slower or even impossible if they are less relevant or irrelevant.

AMBR is a model of analogy-making based on DUAL, which integrates memory and reasoning. The mechanisms for memory access, mapping, inference, re-representation, etc. are based on emergent computations implemented over a large set of DUAL agents. Memory access is based mainly on the spreading activation mechanism of the connectionist aspect of DUAL. Mapping is based on a number of mechanisms such as marker passing for establishing semantic correspondence, temporary-agent constructors for establishing hypotheses about possible correspondences, link constructors for establishing positive or negative links among hypotheses and existing long-term agents based on structure correspondence, etc. All these mechanisms are running in parallel and influence each other, thus giving rise to various interaction effects.

AMBR Mechanisms for Dynamic Extension of Episode Representation

Episode representation is dynamically extended in AMBR by the interplay of three processes running in parallel: (*i*) gradual and partial retrieval of episodic and semantic memory elements, (*ii*) gradual and partial mapping the retrieved episode and semantic elements onto the target elements, and (*iii*) gradual and partial instantiation of general statements from semantic memory.

The gradual retrieval process is based on the spreading activation over the links between the neighboring agents. When the activation level of certain agent exceeds a given threshold, the agent becomes part of the working memory[2]. It is possible that only part of the coalition passes the threshold, which means that it is possible that only part of the encoded episode elements are retrieved. Thus different representations of the past episode are “constructed” or “retrieved” in different contexts. This differs from other analogy models. Most of them use centralized episode representations (Forbus, Gentner, & Law, 1995, Thagard et al., 1990). Even in LISA (Hummel & Holyoak, 1997) where the episodes are represented in a decentralized way and where the retrieval process is a gradual one, there is a final decision about which episode has won the competition. This decision is done centrally and all elements of the winner are switched from “dormant” into “active” state. Therefore, no partial retrieval of episodes is possible. In AMBR there is even no in-principle possibility to do this form of forced retrieval of whole episodes because the system does not keep any central registry of rosters enumerating the affiliation of elements to episodes.

Since there are tight links between the elements of semantic and episodic memory, activated agents do not necessarily represent elements of an episode. They can also represent pieces of semantic knowledge. Thus, contrary to other models, the retrieval process in AMBR brings both elements of episodic and semantic memory into the WM. Since semantic knowledge is also represented in a decentralized manner, it has the same degree of flexibility. Two scenarios are worth mentioning. The spreading activation mechanism can retrieve (*i*) a coalition representing schematic knowledge about a typical situation (e.g. “boiling water in the kitchen”) or (*ii*) single generic statements (such as “a pot is made of metal”). Because the process of instantiation of a schema is much more traditional and well studied, we will focus on the instantiation of single generic statements.

The gradual mapping process starts as soon as the first elements from episodic or semantic memory pass the working-memory threshold. An attempt is made to map them onto elements from the target description. An external observer monitoring the behavior of the system as a whole can ascribe different labels to this process depending on the particular kind of prior knowledge that the system happens to use in each particular case. If a past episode is retrieved and mapped to the target, this could be labelled "analogy". If a general schema is retrieved and used as a source for the

[2] After entering the WM, the graded activation continues to play an important role since the speed of symbolic processing performed by the agent depends on its activation level.

mapping, the "analogy-making" mechanism produces an inference that we might prefer to call deductive[3]. The prevailing number of cases will be mixed, however: both episodic and semantic elements will be mapped. These are the cases considered in more detail in this paper. The process of analogy-making is an emergent process. What actually happens in the system at the micro-level is that individual elements of the descriptions try to find their "mates", i.e. to form correspondence hypotheses between target elements and retrieved elements regardless of whether these elements are originating from episodic or semantic memory. At the same time all the agents participating in this process establish temporary links among themselves in order to cooperate in finding a structurally consistent mapping (Gentner, 1983).

The details of how retrieval and mapping are performed in a decentralized and emergent way will not be presented here because of lack of space. Interested readers can find such descriptions elsewhere (Kokinov, 1994a, Kokinov, Nikolov, & Petrov, 1996). The focus here is on the processing that takes place after a mapping between elements of semantic memory and target elements is established. For example, when an isolated generic statement from semantic memory, such as "made-of-1(teapot,metal)" or "is-hot-1(plate)", is retrieved it can be mapped onto elements of the target description such as "made-of-2(vessel-1,wood-3)" or "is-burning-2(fire-1)", respectively.

After the initial correspondence is established, which might be based on the semantic similarity between the predicates (established by the marker-passing mechanism), a generic hypothesis is formed (i.e. a new agent is created) which puts the target proposition—"made-of-2(vessel-1, wood-3)"—in correspondence with the general statement coming from semantic memory that teapots are typically made of metal: "made-of-1(teapot,metal)". In case that the retrieved episode representation already contains a statement "made-of-3(teapot-1,metal-1)" then most probably it will win the competition and the generic hypothesis will be rejected. However, if such statement is not encoded in the long-term memory since the material of the teapot was not important at the time of experiencing the event, or it was encoded but for some reason it is now not retrieved in WM and therefore does not exist in the current representation of the episode, then it might happen that the generic hypothesis wins the competition or at least is strong enough to start an instantiation process.

The instantiation process builds up a new proposition where all the universally quantified variables will be replaced by specific instances-constants, e.g. "made-of-1(teapot,metal)" goes into "made-of-4(teapot-1,metal-prototype)". How are the constants chosen? If there is a constant (object) of the same type in the retrieved episode, it should be used. In the example above, the episode representation involves such an instance—"teapot-1". Then the new proposition will use it as an argument. If the episode contains no specific instance of that type then a new instance is constructed which has the properties of the prototype of the corresponding class. In this example, there is no instance of metal in the episode representation and therefore a new instance is formed—"metal-prototype". Thus the instantiation mechanism tries to reuse existing instances whenever possible. The DUAL marker-passing mechanism provides information about which instances of the concept "teapot", if any, are active in the current representation of the episode and hence are available for instantiation.

Instantiation has been used in analogy-making models so far only for adding new objects and propositions to the target problem, i.e. for making analogical inferences (Holyoak & Thagard, 1989, Falkenhainer, Forbus, & Gentner, 1989). In AMBR it is used for extending episode representation and relies heavily on the semantic knowledge of the system.

In summary, the process of extending the representation of the episode emerges from the interaction of several processes that are themselves emergent: the retrieval process which continuously brings up new episodic and semantic memory elements into WM, the mapping process which continuously builds hypotheses about possible correspondences between the retrieved elements and elements from the target description, and the instantiation process which continuously constructs new specific propositions based on generic propositions retrieved from semantic memory.

Why is the continuous interplay between these three processes important? The interactions guide each of the processes and therefore make each of them more effective. They preclude the model from doing exhaustive search. The influence of retrieval on mapping and instantiation is obvious since nothing can be mapped or instantiated if it is not activated (retrieved). The role of mapping is unusual compared to other models of analogy-making. Since the retrieval process in AMBR is a piece-by-piece process that runs continuously and in parallel with the mapping, the latter can influence the former. It is always the case that the retrieved elements of the episode send out activation to the rest of the elements of the episode representation and thus constantly try to activate the whole coalition. However, if the coalition is not tight enough (which is the typical case) they would be able to retrieve only some of their coalition partners. Exactly which elements will be retrieved depends not only on the initial set of elements but also on their mapping status, i.e. which of them are mapped onto target elements and which are not. Elements that are mapped receive abundant activation from the target and therefore will play important role in any further retrieval. In this way the mapping influences the retrieval process.

The importance of this interaction between processes can be demonstrated by contrasting two runs of the system: one with parallelism and interaction and one without. Figure 1 presents such a comparison from a simulation experiment. In the "parallel condition" (thick lines in the figure) all processes are running in the way they have been described so far. In the "retrieval only condition" the mapping process has been intentionally switched off. The important result is that, although the target and the background knowledge were exactly the same in the two runs, two different episodes are retrieved—a more structurally similar one in the first case and a more superficially similar one in the second case.

[3] AMBR has been proposed as a unifying mechanism for deductive, inductive, and analogical reasoning (Kokinov, 1988, 1990, 1992, 1994a).

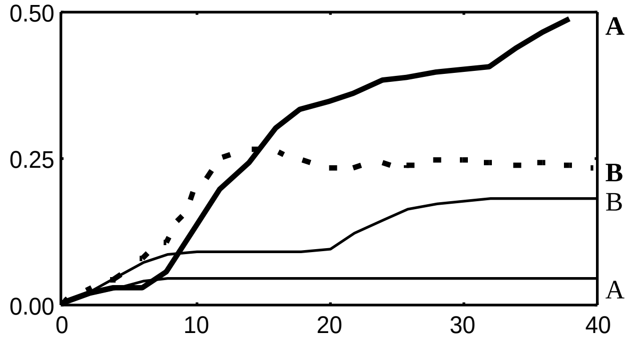

Figure 1. Retrieval indices for two episodes, A and B, in two different conditions as a function of time. The thick lines correspond to the parallel condition in which mapping influences retrieval, the thin lines show 'pure' retrieval.

The mapping influences the instantiation process as well. If there were no such influence, the model would have to build up unrealistically many instantiations—one or more for each generic proposition retrieved from semantic memory. However, the instantiation process is guided by the mapping process—only general propositions that are mapped onto target propositions will be instantiated. On the other hand, once an instantiation is build it supports the mapping and helps in further retrieval of memory elements.

Conclusions

The mechanisms described above allow for dynamic re-representation of the episodes by: retrieving additional information from episodic memory based on the established mappings; by constructing new memory elements and integrating them into the episode representation based on instantiation of generic statements retrieved from semantic memory and mapped onto the target description; and by retrieving elements from other episodes thus producing a blending between episodes.

In this way AMBR makes the following predictions which can be tested experimentally. The first prediction is that the partial mapping established up to a point influences the further retrieval process. This prediction can be tested by analysing thinking-aloud protocols. Actually, such results have been obtained by Ross and Sofka (1986) as a side effect in a thinking-aloud study. They are summarized in (Ross, 1989) as follows: "... other work (Ross & Sofka, 1986) suggests the possibility that the retrieval may be greatly affected by the use. In particular, we found that subjects, whose task was to recall the details of an earlier example that the current test problem reminded them of, used the test problem not only as an initial reminder but throughout the recall. For instance, the test problem was used to probe for similar objects, and relations and to prompt recall of particular numbers from the earlier example. The retrieval of the earlier example appeared to be interleaved with its use because subjects were setting up correspondences between the earlier example and the test problem during the retrieval." The simulation data described here are obtained absolutely independently and are based only on the theoretical assumptions of DUAL and AMBR and exhibit exactly the same pattern of interaction.

A second prediction is that people would instantiate generic knowledge in cases where there is missing information from the episode representation and where this information is needed for the mapping, i.e. there is a corresponding piece of information in the target which needs to be mapped onto something from the base. An experiment is currently being prepared to test this hypothesis. McKoon and Ratcliff (1981) demonstrated that people make inferences and extend episode representation during the encoding process, e.g. after listening to a sentence such as "Alice pounded in the nail until the board was safely secured." listeners would infer and encode that "Alice used a hammer." Our prediction is that they would further extend the representation during the recall process when they use that episode in order to map it to the target.

Finally, a third prediction is that people will tend to blend episodic information if the information needed for mapping is missing in the best retrieved episode, but is present in another episode that is also partially retrieved. Another experiment is under development to test this prediction.

References

Anderson, J. (1983). *The architecture of cognition*. Cambridge, MA, Harvard University Press.

Anderson, J. & Lebiere, C. (1998) *The atomic components of thought*. Mahwah, NJ:. Erlbaum

Anderson, J. & Ross, B. (1980). Evidence Against a Semantic-Episodic Distinction. *Journal of Experimental Psychology: Human Learning and Memory, 6* (5), 441-466.

Bartlett, F. (1932). *Remembering*. Cambridge: Cambridge University Press.

Charlmers, D., French, R., & Hofstadter, D. (1992). High-Level Perception, Representation, and Analogy: A Critique of Artificial Intelligence Methodology. *Jornal of Experimental and Theoretical AI, 4* (3), 185-211.

Dunker, K. (1945) *On problem solving*. In: Psychological Monographs, vol. 58 (whole No 270).

Ernst, G. & Newell, A. (1969). *GPS: A case study in generality and problem solving*. NY: Academic Press.

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

Forbus, K., Gentner, D., & 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. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.). *Similarity and analogical reasoning*. New York, NY: Cambridge University Press.

Herrmann, D. & Harwood, J. (1980). More evidence for the existence of separate semantic and episodic stores in long-term memory. *Journal of Experimental Psychology: Human Learning and Memory, 6* (5), 467-478.

Hofstadter, D. & the Fluid Analogies Research Group (1995). *Fluid concepts and creative analogies.* New York: Basic Books.

Holyoak K. & Thagard P. (1989b). 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* (3), 427-466.

Kokinov, B. (1994a). A hybrid model of reasoning by analogy. In K. Holyoak and J. Barnden (Ed). *Advances in connectionist and neural computation theory. Vol. 2: Analogical connections*. Norwood, NJ: Ablex

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

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

Kokinov, B. (1997). Micro-level hybridization in the cognitive architecture DUAL. In R. Sun & F. Alexander (Eds). *Connectionist-symbolic integration: From unified to hybrid architectures.* Hillsdale, NJ: Erlbaum

Kokinov, B., Nikolov, V., & Petrov, A. (1996). Dynamics of emergent computation in DUAL. In A. Ramsey (Ed.). *Artificial intelligence: Methodology, systems, applications*. Amsterdam: IOS Press

Kolodner, J. (1983). Reconstructive memory: A computer model. *Cognitive Science*, 7, 281-328.

Lange, T. & Wharton, C. (1992). Dynamic memories: Analysis of an integrated comprehension and episodic memory retrieval model. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum

Loftus, E. (1977). Shifting human color memory. *Memory and Cognition, 5*, 696-699.

Loftus, E. (1979). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.

McKoon, G. & Ratcliff, R. (1981). The comprehension process and memory structures involved in instrumental inference. *Journal of Verbal Learning and Verbal Behavior, 20*, 671-682.

Maier, N. (1931). Reasoning in humans II: The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology, 12*, 181-194

Minsky, M. (1986). *The society of mind*. New York: Simon & 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 E. Winograd & U. Neisser (Eds.) *Affect and Accuracy in Recall.* NY: Cambridge University Press.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University 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. & Sofka, M. (1986). Remindings: Noticing, remembering, and using specific knowledge of earlier problems. Unpublished manuscript.

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