Blending of Non-Similar Episodes as a Result of Analogical Mapping with a Third One

Neda Zareva-Toncheva (nzareva@cogs.nbu.bg)[1]
Boicho Kokinov (bkokinov@nbu.bg)[12]
[1]Central and East European Center for Cognitive Science, Department of Cognitive Science and Psychology,
New Bulgarian University, 21 Montevideo Street
Sofia 1635, Bulgaria

[2]Institute of Mathematics and Informatics, Bulgarian Academy of Sciences, Acad. G. Bonchev Street, bl.8
Sofia 1113, Bulgaria

Abstract

This paper presents an attempt to establish a previously unknown phenomenon: blending of dissimilar episodes in long-term memory. In contrast with previous work which demonstrates that similar episodes can be blended, the current experiment shows that even highly dissimilar episodes can be blended if these episodes have simultaneously been mapped onto a third one. In this way analogy-making can produce memory distortions even between episodes that are neither superficially, nor structurally similar. This result confirms a prediction made by the AMBR model of analogy-making.

Introduction

Memory distortions have received recently a considerable amount of interest (Loftus, 1979/1996; Loftus, Feldman, & Dashiell, 1995; Neisser, 1998; Roediger & McDermott, 1995; 2000; Roediger, 1996; Moscovitch, 1995; Reinitz, Lammers, & Cochran, 1992; Schacter, 1995, 1999; Schacter, Koutstaal, Norman, 1997, Schacter, Norman, Koutstaal, 1998, Koriat, Goldsmith, Pansky, 2000; Goldsmith, Koriat, & Weinberg-Eliezer, 2002). This interest is provoked by the desire to find out the limitations of our memory system, but even more importantly, by the possibility to learn more about its underlying mechanisms by exploring the errors it produces.

The experimental research has been concentrated on memory distortions which are due to similarity between episodes. Thus the pioneering work of Elizabeth Loftus (1977, 1979/1996) and her colleagues (Loftus, Feldman, & Dashiell, 1995) has clearly demonstrated that when two episodes (e.g. an observed scene and a heard story, a real and an imagined event) differ in a few details, then our memory mixes them up. Thus the intrusions from one episode into the other causes the two episodes to be blended and people falsely report the participating objects (hammer vs. screwdriver) or their properties (blond hairs vs. dark hairs). James McClelland and his colleagues (Nystrom & McClelland, 1992; McClelland, 1995) demonstrated “trace synthesis errors” when subjects were presented with highly similar sentences (sharing most of the words except two nouns). As result they mixed them up and there were intruders from the alternative, but similar, sentence.

Sir Frederic Bartlett (1932) demonstrated how episodes are reconstructed based on their similarity with the schema, while Henry Roediger and Ken McDermott (1995) provided examples of errors based on associative strength and similarity with a semantic field (e.g. “sleep” is highly related to “bed”, “rest”, “tired”, “dream”).

Little or no attention has been paid to the structure of the episodes and its role in blending. Structure plays an important, even dominating, role in analogy-making and therefore it would be interesting to study the possibility of mixing up episodes which have similar structure even though they do not share superficial features. Kokinov (1998) has provided evidence for such blends. In the current paper a further step is attempted: the question is asked whether people do blend episodes that share few superficial and structural features. The experiment reported here extends a line of research initiated by the AMBR research group (Kokinov & Petrov, 2001; Kokinov & Zareva-Toncheva, 2001, Grinberg & Kokinov, 2003). A characteristic feature of this approach is the integration of memory and reasoning in the experimental tasks. Thus the reasoning process (analogy in this case) establishes direct and indirect structural correspondences which later on result in blending.

Theory

At the theoretical level blending effects are explained by the reconstructive nature of human memory. A few models exist that do replicate some of these findings, these include Hintzman’s (1988) multi-trace model, Metcalfe’s (1990) CHARM model, and McClelland’s (1995, McClelland, McNaughton & O’Reilly, 1995) PDP-type of model. All these models are based on distributed representation of the episodes or objects and the overlap between these representations causes the memory errors. This type of models can explain the cases when similar episodes are blended (i.e. when most of the features in the feature vector are the same), however, it would be difficult to explain why dissimilar episodes can be blended.

The AMBR model of analogical problem solving (Kokinov, 1988, 1994a, 1998; Kokinov & Petrov, 2000, 2001; Petrov & Kokinov, 1998, 1999) shares the view of decentralized representation of episodes, however, it also respects their structure. This model intends to explain analogy-making and therefore should be highly sensitive to structure. On the other hand, it aims also to provide a realistic explanation of the mechanisms of human memory, which will allow it to demonstrate those memory errors which are typical of human beings. Unfortunately, the most popular models of analogy-making SME-MAC/FAC (Forbus, Gentner, & Law, 1995), ACME-ARCS (Thagard, Holyoak, Nelson, & Gochfeld, 1990), LISA (Hummel & Holyoak, 1997) ignore the memory distortion data and assume "perfect" memory of episodes. They are based on centralized representation of episodes which means that episodes are either retrieved as a whole (and than mapped onto the target problem description) or they fail to be retrieved. No blending of episodes may occur, no false memories can arise, no partial retrieval can happen. There is no interaction between retrieval and mapping in these models.

Surprisingly, this is true even for the LISA model (Hummel & Holyoak, 1997) which is based on distributed representations, since the representations are truly distributed only in working memory, but the episode representation in LTM is highly centralized – a list of all units representing a single episode is assumed and if the episode wins the competition between episodes then all its units are switched from dormant to active state and mapping begins from that moment on.

The AMBR model combines decentralized representation of episodes with mechanisms for interaction between memory and reasoning. These two factors are crucial in explaining blending between dissimilar episodes: mapping produces new links and thus changes the representation of the old episodes. Thus AMBR predicts blending effects which are due to analogical mapping. The current experiment is designed to test these predictions. Such predictions have no parallel in any other model we know of. Although other analogy-making models could produce new mapping links, these links have no effect of the encapsulated representations of the episodes and thus no intruders from other episodes may appear.

AMBR Model and Blending

The AMBR model is based on the general cognitive architecture DUAL (Kokinov, 1994b,c). This architecture integrates symbolic representation of structure and connectionist representation of context and relevance. It is based on decentralized representations and emergent computations produced by a society of micro-agents (Minsky, 1986; Hofstadter, 1995). DUAL micro-agents are hybrid having both symbolic and connectionsit parts.

An episode is represented in the long-term memory of DUAL and AMBR by a coalition of micro-agents which are linked to each other and tend to be co-activated. Each of the agents represents symbolically a single aspect of the episode (e.g., a simple statement). The active agents at a particular moment of time comprise the WM. The elements of a coalition which are currently in WM represent the current partial view on the episode.

The process of episode recall is modeled by a process very similar to the analogy-making process: the memory cue or question is represented as a target and AMBR runs to find an episode which will be literally similar to the target and aspects of the old representation are mapped and transfered to the target which is considered to be the recall of the episode. The competition mechanisms in AMBR allows elements from two or more different coalitions (episodes) to be mapped and transfered to the target. This is typically avoided by the strong excitatory links within the coalitions and the inhibitory links between the alternative hypothesis, however, it may still happen under certain circumstances. Thus when the coaltion is not very well-connected and not very typical it invites intruders from other coalitions and thus blending between episodes take place. The mapping links established between analogous episodes additionally raise the probability of activating elements from alternative episodes and converting them into intruders. For this to happen, however, additional conditions should be met. Most importantly, the receiving coalition should be not active enough or important aspects of the corresponding episode should not have been encoded or activated at the moment. The intruders will fill in the blanks (the missing but needed information). The latter process is controlled by the analogical mapping process. Details about the simulation results can be found in (Grinberg & Kokinov, 2003). A simple example will be presented here.

Double analogy

Here is an example of what we call double analogy. Let B1 and B2 be two old episodes which partially map onto the target episode T (Figure 1). Let A, B, I from B1 correspond to E, F, J in T, and the relation R2 corresponds to R3. On the other hand, let C, D, G in episode B2 correspond to F, E, H in the target T, and R1 to R3. This is a situation that we will call double analogy. This analogy makes implicit correspondences between B1 and B2. Thus both B and C correspond to F, and both A and D to E. In this way we may think of A corresponding in some indirect way to D, and B to C. After such a double analogy is established and all those correspondence links formed, what might be expected is that when later on the base episode B1 is recalled we may get intruders from B2. In particular, we might expect that A might be recalled as connected (related) to G. Such intrusion will come from the complex mapping established with T.

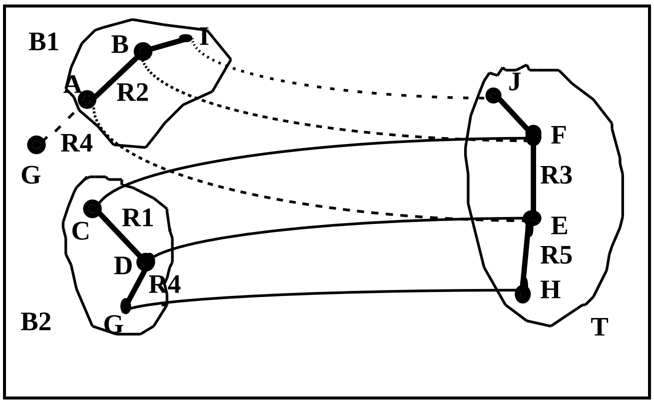

Figure 1. An example of double analogy between the base episodes B1 and B2, and the target episode T.

The experiment described below is designed to test this prediction of the model. The important point is that B1 and B2 do not need to be similar, neither superficially, nor structurally. It is enough that each of them is similar to different parts of T.

Experiment

In this experiment the participants solve two base problems (D and E) and as a side effect remember them in their long-term memory. Two weeks later half of the participants solve a target problem (T2) that requires the two base problems to be retrieved in WM and a double analogy with them to be made, and the other half solve a target problem (T1) that does not require a double analogy, but just a single analogy with one of the problems (D). Finally, the participants are asked to retell the base problems as accurate and complete as possible. We measure the blending between the base problems and expect that those who made the double analogy will tend to make more blends between D and E.

Method

Design

The experiment consists of three sessions:

- Session 1: solving two base problems (D and E);
- Session 2: solving one of two target problems (T1 – analogous to D, or T2 – partially analogous to D and partially analogous to E);
- Session 3: free recall of the problems from session 1.

The manipulated factor is the type of target problem solved during the second session (Figure 1). In the control group participants solve a target problem analogous only to D, while in the experimental group the participants solve a problem that requires a double analogy, i.e. it requires to make analogies both with D and E (the blending condition).

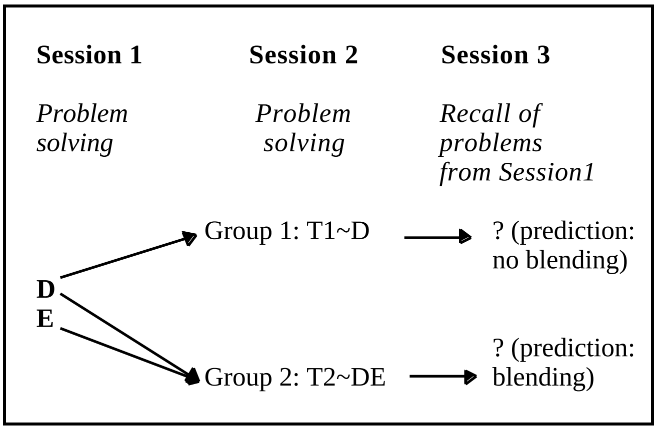

Figure 2. General design of the experiment.

Thus we use a between subject design, the independent variable is the type of target problem solved in session 2 and how it relates to the base problems. The dependent variables measure the blending found in session 3 in the recalled versions of the stories from session 1. We use a qualitative and a quantitative measure:

- the qualitative measure comes from expert judgment: experts decide whether or not a blend of D and E has been produced;
- the quantitative measure comes from formal protocol encoding and enumerating the number of intruders from the alternative story in the protocol – in this way a continuous measure is formed which we call degree of blending (it takes values between 0 and 1).

Material

The target problems presented to the participants in the second session are designed in such a way that they could not be solved without an analogy with the base problems, i.e. they are difficult problems based on an unknown (fictitious) principle. The fictitious principle is provided when solving one of the base problems in session 1. This design ensures that the target problems will be solved by analogy. Here are the fictitious principles:

- the fictitious principle in base problem D is that the genes of an imagined extraterresial population are not carried over from generation to generation, but are conveyed to every other generation, i.e. the kids inherit the genes of their grandparents, not of their parents;
- the fictitious principle in base problem E is that external force can cause a change in an organism producing new chemical elements in it in such a way that the organism starts to process the ingredients in a new way, thus normally used food may become dangerous being involved in a chemical reaction with the new chemical elements.

Problem D: The population of the planet Tokay is separated into two races. One of them could live only in warm and dry places, while the other one could live only in cold and humid places. Their genotype was such that if anybody happened to be in the opposite condition, he or she would die. An important feature of the members of both races was that their genes passed over every other generation. One year the planet was affected by a passing by radioactive planet, which caused anomalies among the young dwellers, so that members of one of the races acquired the characteristic of the other one and vice versa. In order to survive the affected members of both races had to swap places. The question is what can we expect to happen in the next decades.

The correct answer is that since the genes passed over every second generation the next generations will always have to be a swap places in order to survive, since the children of the current generation will belong to the race of their grandparents and thus will have to go back to their lands, then their children will belong to the alternative race and will have to change the homeland again, etc.

Problem E: A tribe lived in a waterless village and had to use water from the caves in the mountains nearby. The water that the members of the tribe were using for ages had a high concentration of acids, but evidently their organisms adapted to it. However, suddenly, an earthquake blocked the water springs in the caves and the people had to search for new water sources. Finally they found one and greedily drank from it. However, the people who drank water from the new source got stomachaches. To play it safe, the doctors offered them traditional cave water from a concealed reserve. Few hours later they died in convulsions. The question is what caused the shuttering death of these people.

The correct solution involves the fact that the new spring water had different ingredients which caused a sudden change in people's metabolism. The highly acid water they drank afterwards triggered a chemical reaction that converted the normal water into poison.

*Target problem **1*** was about a planet, where two types of cows were grown up – one of them was a meat breed and the other – a milk breed. Once the animals were exposed to a radiation impact and immediately after that the meat-breed cows became milk-breed and vice versa. Fortunately, the next generation of the affected cows had characteristics corresponding to the original breeds. The following generation, however, was born with the opposite characteristics. The question was why did the last change happen. Thus the participants had to make analogy with base problem D only and use the principle that the genes of the cows passed over every second generation.

Target problem 2 was also about the meat-breed and milk-breed cows grown up on a planet. The meat-breed cows had a smooth and soft coat, while milk-bread cows had a rough and firm coat. All cows drank water from well-established springs, until one of the springs run short of water and had to be replaced with another one. After some time, the cows that drank from the new spring started changing their coats – the meat-bread cows with smooth and soft coat got rough and firm coat and vice versa. Fortunately the calves of the affected cows were born with normal coats. The following generation, however, was born again with reversed coats. The questions were why the coats of the cows have changed initially, and why the next generation was normal, but the generation after that was reversed again.

Procedure

The participants were tested individually. During the first session they had to solve five problems (D, E, and three fillers). The problems were given one by one without an explicit time restriction. After they produced a written solution and if this was not the targeted one the experimenter encouraged them to find an alternative solution, if this did not help, a hint was given, and finally if the target solution was not produced it was provided by the experimenter. All participants in session 1 had to solve the base problems correctly and to acquire the basic principles to be used further in the later sessions. There were no instructions to remember the problems.

The second session followed after a period of 14 days. During this second session the group was split and half of the participants solved problem T1, and the other half – T2. The second session was run again individually and the thinking aloud method was used, the speech of the participants was recorded. No hints were provided here.

The third session followed immediately after the second one. The participants were asked to retell the problems from the first session as accurate and complete as possible. The stories were reproduced orally and tape-recorded.

Participants

48 workers (from 20 to 35 years old, most of them with an university degree) from two companies participated in the experiment for payment. Only 31 were able to solve the target problems in session 2 (13 in group 1 solved target problem 1, and 18 in group 2 solved target problem 2). 15 were female and 16 – male.

Results and Discussion

The records were transcribed and the protocols of the third session were used as the main data set. The texts of the problems D and E were segmented into separate statements and their appearance in the body of the protocol was encoded. We separated the text of problem D into 29 statements marked as D1-D29 and whenever a phrase (or its semantic equivalent) occurred in the narration of the subject the corresponding Dk was inserted in the protocol encoding, the text of problem E was separated into 41 statements in the same way, they were marked E1-E41.

When reproducing one of the problems participants sometimes inserted statements from the other problem. We were counting these intrusions and used their number in calculating the degree of blending. It was calculated as the ratio: the number of Ds over the number of Es (when the Ds were less than the Es), or reverse – the number of Es over the number of Ds (when the Ds were more than the Es). Thus if the number of Ds or Es are zero than no blending has occurred (degree of blending is 0), and when the number of Ds and Es are equal then an absolute blend has been produced (with degree of blending equal to 1). The results are shown in Figure 3. As we can see the results are coherent with our hypothesis: we have 4 times higher degree of DE blending in group 2 (where T2 required double analogy with D and E) than in group 1. The difference in the degree of DE blending is significant at p-level <0.01 ($F(1,29)=7.817$, $p=0.009$). Only intrusions from the other base problem (D or E) was counted, intrusions from the target were ignored in this calculation.

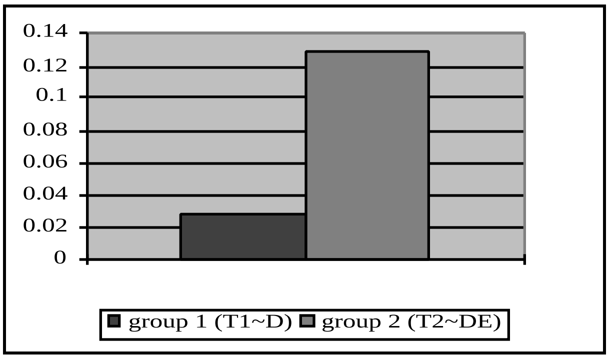

Figure 3: Degree of blending between the base problems depending on which target problem has been solved in session 2.

Since we used a very formal method of measuring blending – the number of intrusions as registered in the protocols – we were curious to compare this to a more qualitative judgment done by human experts who may recognize whether there is a real blending or just some general knowledge intrusions or superficial mixture. Two independent judges had to read each protocol (without knowing neither about our hypothesis nor which group this protocol comes from). The experts had to judge whether there was a blending between D and E. There was a high degree of agreement between the experts (about 10% disagreement where a third expert was called for arbitration). The percentage of DE blending is presented in Figure 4.

The results from the expert judgments are coherent with the protocol encoding results of the degree of blending and again far more DE blends are observed in group 2 than in group 1. The difference is significant measured by the Chi-square test (Chi-square = 5.134, p=0.02).

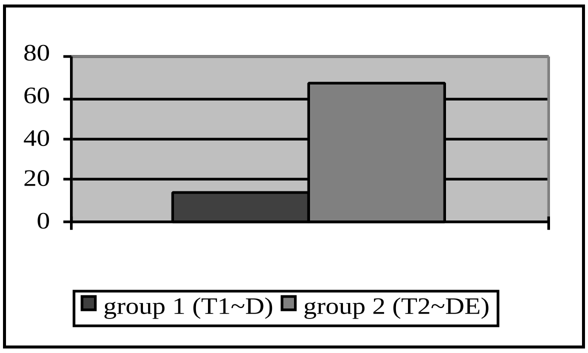

Figure 4: Percentage of blended memories as judged by experts. Participants who made a double analogy in the second session demonstrate 5 times higher percentage of blending than those who made single analogies.

Conclusions

AMBR predicts (Kokinov & Petrov, 2001) that in the case of double analogy-making additional links are established between the two base episodes. This allows afterwards these two base episodes to be blended even when they are not similar to each other as the simulations demonstrated (Grinberg & Kokinov, 2003).

The experiment described in this paper tests this prediction. Two quite dissimilar stories were designed (one involving extraterresial races whose genes are inherited every other generation, and the other one is about a tribe which happened to drink new type of water and as result of that a metabolism changed their stomach juices and the old water became a poison). These two stories were mapped onto a single target story in the experimental condition and only one of them was mapped to the target in the control condition. We obtained about 4 times higher degree of blending in the experimental condition.

These results are coherent with a previous study which used different design and stories (Kokinov & Zareva-Toncheva, 2001) and we can gladly acknowledge the phenomenon of blending between dissimilar episodes. A number of open issues remain: whether these blending effects are temporal as result of residual activation, or they are due to permanent changes in long-term memory; whether double analogy-making is a necessary condition or simple co-activation of the two base stories will be enough to produce blending; whether blending of dissimilar episodes which are presented during different sessions can also happen or the initial temporal relation is crucial. We are currently running experiments to answer these and other questions. From the results so far it seems that the changes are permanent and co-activation is not enough for producing blending. This makes the interplay between analogy-making and memory even more interesting and important to study.

Acknowledgments

We would like to thank the AMBR research team for their continuous support and stimulating environment. We would like to thank Georgi and Penka Petkovi, Boika Bratanova and Irina Gerdjikova for helping in running the experiment and encoding the data.

References

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

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

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

Goldsmith, M., Koriat, A., Weinberg-Eliezer, A. (2002). Strategic Regulation of Grain Size Memory Reporting. *Journal of Experimental Psychology: General*, 131:73-95.

Grinberg, M. & Kokinov, B. (2003). Simulating Episode Blending in the AMBR Model. In: *Proceedings of the European Cognitive Science Conference*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Hintzman, D. (1988). Judgements of frequency and recognition memory in a multiple-trace model. *Psychological Review, 95*, 528-551.

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

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

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. (1994a). A hybrid model of reasoning by analogy. In K. Holyoak & J. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections* (pp. 247-318). Norwood, NJ: Ablex

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

Kokinov, B. (1998). Analogy is like cognition: Dynamic, emergent, and context-sensitive. In K. Holyoak, D.

Gentner, & B. Kokinov (Eds.), *Advances in analogy research.* Sofia, Bulgaria: NBU Press.

Kokinov, B. & Petrov, A. (2000). Dynamic Extension of Episode Representation in Analogy-Making in AMBR. In: *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

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: Perspectives from Cognitive Science*, Cambridge, MA: MIT Press.

Kokinov, B. & Zareva-Toncheva, N. (2001). Episode Blending as Result of Analogical Problem Solving. In: *Proceedings of the 23nd Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Koriat, A., Goldsmith, M. (1996). Monitoring and Control Processes in the Strategic Regulation of Memory Accuracy. *Psychological Review*, 103: 490-517.

Koriat, A., Goldsmith, M., Pansky, A. (2000). Toward a Psychology of Memory Accuracy. *Annual Review of Psychology*, 51: 481-537.

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

Loftus, E. (1979/1996). *Eyewitness testimony*. Cambridge, MA: Harvard University Press

Loftus, E., Feldman, J., & Dashiell, R. (1995). The reality of illusory memories. In D. Schacter (ed.), *Memory distortions: How minds, brains, and societies reconstruct the past*. Cambridge, MA: Harvard University Press.

McClelland, J. (1995). Constructive memory and memory distortions: A parallel distributed processing approach. In D. Schacter (Ed.), *Memory distortions: How minds, brains, and societies reconstruct the past*. Cambridge, MA: Harvard University Press.

McClelland, J., McNaughton, B. & O'Reilly, R. (1995). Why there are Complementary Learning Systems in the Hypocampus and Neocortex: Insights from the Successes and Failures of Connectionist Models of Learning and Memory. *Psychological Review*, 102: 419-457.

Metcalfe, J. (1990). Composite holographic associative recall model (CHARM) and blended memories in eyewitness testimony. *Journal of Experimental Psychology: General, 119*, 145-160.

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

Moscovitch, M. (1995). Confabulation. In D. Schacter (Ed.), *Memory distortions: How minds, brains, and societies reconstruct the past*. Cambridge, MA: Harvard University Press.

Neisser, U. (1998). Stories, selves, and schemata: A review of ecological findings. In M. Conway, S. Gathercole, & C. Cornoldi (Eds.), *Theories of memory* (Vol. 2). Hove, UK: Psychology Press.

Nystrom, L. & McClelland, J. (1992). Trace synthesis in cued recall.*Journal of Memory and Language,31*:591-614

O'Reilly, R. & Norman, K. (2002). Hippocampal and Neocortical Contributions to Memory: Advances in the Complementary Learning Systems Framework. *Trends in Cognitive Science*, 6: 505-510.

Petrov, A. & Kokinov, B. (1998). Mapping and access in analogy-making: Independent or interactive? A simulation experiment with AMBR. In K. Holyoak, D. Gentner, & B. Kokinov (Eds.), *Advances in analogy research*. (pp. 124-134) Sofia, Bulgaria: NBU Press.

Petrov, A., & Kokinov, B. (1999). Processing Symbols at Variable Speed in DUAL: Connectionist Activation as Power Supply. In: Dean, T. (ed.) *Proceedings of the 16th International Joint Conference on Artificial Intelligence.* (pp. 846-851) San Francisco, CA: Morgan Kaufman.

Reinitz, M., Lammers, W., & Cochran, B. (1992). Memory-conjunction errors: miscombination of stored stimulus features can produce illusions of memory. *Memory & Cognition, 20*, 1-11.

Roediger. H., McDermott, K. (1995). Creating False Memories: Remembering Words not Presented in Lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21: 803-814.

Roediger. H., McDermott, K. (2000). Tricks of Memory. *Current Directions in Psychological Science*, 9: 123-127.

Roediger, H. (1996). Memory Illusions. *Journal of Memory and Language*, 35: 76-100.

Schacter, D., Ed. (1995). *Memory distortions: How minds, brains, and societies reconstruct the past*. Cambridge, MA: Harvard University Press.

Schacter, D. (1999). The Seven Sins of Memory. *American Psychologist*, 54: 182-203.

Schacter, D., Koutstaal, W., Norman, K. (1997). False Memory and Aging. *Trends in Cognitive Science*, 1: 229-236.

Schacter, D., Norman, K., Koutstaal, W. (1998). The Cognitive Neuroscience of Constructive Memory. *Annual Review of Psychology*, 49: 289-318.

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