In Proceedings of Cogsci-2005, the Annual Meetings of the Cognitive Science Society, July 21-23, Stresa, Utaly.

From Reflection to Interaction: Use of Memory in Interactive Knowledge Acquisition

Jihie Kim (jihie@isi.edu)
University of Southern California/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292 USA

Abstract

Most existing tools for interactive knowledge acquisition have limited understanding of how past knowledge authoring activities are related to current situation and provide limited assistance in organizing various knowledge acquisition tasks. Many tools do not keep track of how users perform KA tasks, how acquired knowledge is tested and used, and what needs to be improved. This paper describes our analysis of the literature on memory-inspired techniques developed in cognitive science and computer science research and presents a compilation of useful techniques that interactive knowledge acquisition tools can use in making the interactions more successful. The focus of our analysis is on how memory-inspired techniques could support reflection on various aspects of knowledge acquisition and knowledge use, and how reflection results can be used in providing better assistance to the user. Based on this analysis we are developing more proactive and effective acquisition tools.

Keywords: artificial intelligence, knowledge acquisition.

Introduction

Acquiring knowledge from end users who have no formal training in computer science remains a challenging task (Blythe et al 2001; Clark et al 2001; McGuinness et al., 2000; Eriksson et al., 1995). Existing knowledge acquisition (KA) tools use various approaches including graphical and structured editors, diagnosing errors and helping users to fix them, using existing knowledge to generate guidance for users, etc. However, users are mainly responsible for the process in terms of deciding when, what, how, and how well to enter knowledge. Most existing tools do not reflect on how users perform KA tasks, how acquired knowledge is tested and used, what needs to be improved, etc. Users themselves have to keep track of past mistakes, current status, potential new problems in order to decide the best options among alternatives. Users can easily become lost in the process of performing various tasks involved in knowledge authoring (Kim & Gil 2000; Pool et al., 2003).

When the user makes the same type of mistake in creating or modifying knowledge, the system could recognize repetitive problems and could provide help in preventing similar problems in the future. When the captured knowledge is used in problem solving, they could be applied based on the level of confidence assessed through their usage patterns. If confident knowledge were overridden or modified, the system could notice them as unexpected events and generate predictions that similar changes may occur in similar situations in the future. In this paper, we describe our analysis on various memory-inspired techniques and how they could be incorporated into interactive acquisition tools.

The paper starts with a brief discussion on the focus of our analysis, how existing techniques could be useful for KA tools. We then present a set of memory-inspired techniques that seem useful in improving interactions during knowledge acquisition and knowledge use. We also show how existing capabilities in KA tools can be mapped to these techniques and how KA tools can be enhanced by adopting the techniques thoroughly and widely. We conclude the paper with a brief introduction of the KA tool that we have developed based on the analysis.

Use of Memory in Cognitive Systems and in Interactive Knowledge Acquisition

In models of cognitive systems (both models of human cognition and other artificial intelligence systems), memories play critical role in learning and problem solving (Tulving 1983). Especially, metacognitive strategies that promote reflective thinking and self-assessment are known to increase the effectiveness of learning. Interactive knowledge acquisition tools can be seen as students learning new knowledge from the teacher (i.e., the user) (Kim and Gil 2003) and they could benefit from applying similar memory-inspired strategies. The systems could maintain memories of past episodes of knowledge acquisition and knowledge use (i.e., problem solving with acquired knowledge), and reflect on what worked and what needs improving. Based on these reflection results, the system could help the user organize their KA tasks in terms of achieving needed improvement.

The focus of our analysis is not on exhaustive compilation of memory techniques in cognitive science and computer science research but on how use of memory could improve the interactions during knowledge acquisition and knowledge use. How can they help the user prevent repetitive mistakes? How can they help users modify knowledge when there are dynamic changes in the world? How to apply acquired knowledge to the problem at hand? How the system can become self-reflective and provide more insightful assistance?

Memory-inspired techniques	Literature/System	Related KA aspects
1. Collecting Experience		
Notice meaningful patterns of information	human experts (Bransford et al., 2000)	Predefined KA events (Clark et al 2001, McGuinness et al., 2000; Blythe et al 2001;…)
2. Organizing Experience		
Organize around common parts of similar episodes	MOP (Schank 1982), CBR systems (Kolodner 1993)	Organizing exemplars (Bareiss et al., 1990)
Organize in ways that reflect deep understanding	human learners (Tulving, 1983; Brown 1987), human experts(Bransford et al., 2000)	
3. Introspect with Experience		
Remember and predict failures & successes	failure driven learning (Cox & Ram 1999), CBR systems (Hammond 1989).	Detect errors (Davis, 1979; Clark et al 2001, McGuinness et al., 2000; Blythe et al 2001)
4. Retrieval of Experience		
Goal driven memory search	human learners(Brown 1987), goal-driven adaptation(Leake 1994)	Implicit KA goals (Davis, 1979; Clark et al 2001, McGuinness et al., 2000; Blythe et al 2001)
Retrieve knowledge based on context	Soar (Rosenbloom et al., 1993), human experts(Bransford et al., 2000)	Use context in assisting users (Marcus & McDermott 1989; Tallis & Gil 1997; Huffman & Laird, 1995; Witbrock et al., 2003)
5. Problem Solving with Experience		
Use similar experience in solving new problems	human learners(Ausubel 1968), CBR systems (Kolodner 1993, Hammond 1989; Veloso & Carbonell 1993).	Use similarity/differences of exemplars (Bareiss et al., 1990)
Use experience to perform similar tasks more efficiently	Prodigy (Veloso et al., 1995) Soar (Rosenbloom et al., 1993),	
Present varying levels of flexibility to new situations	human experts (Bransford et al., 2000), human learners (Piaget & Inhelder 1973), capturing adaptation cases (Craw et al., 2001), temporal learning(Oats 2002)	
6. Knowledge Refinement with Experience		
Generalize/specialize based on range of application	ACT-R(Anderson & Lebiere 1998), models of memory (Schank 1982; Kolodner 1993)	Knowledge refinement with example cases (Ginsberg et al, 1985; Bareiss et al., 1990)
Strengthen knowledge that is used often	ACT-R(Anderson. & Lebiere. 1998)	
7. Memory Management		
forget low utility knowledge	Memory management in CBR systems (Smyth & Keane 1995; Zhu & Yang 1999; Kira & Arkin 2004)	

Table 1: Some of the memory-inspired techniques that are useful for interactive knowledge acquisition

Some of the issues in memory-inspired techniques may not be directly related to our analysis. For example, approaches to understand how biological brains work and building models that are tuned to match biological functions are less essential in our analysis. Also our systems are not subject to cognitive limitations of typical human subjects and can exploit superb computational skills that are not common in models of biological systems.

Memory Inspired Techniques Useful for Interactive Knowledge Acquisition

The studies that we have used include 1) memory related techniques developed in the artificial intelligence (AI) field that support various problem solving and learning capabilities, 2) models of cognitive systems, and 3) how skillful experts use their memory such as how they approach and solve problems and how they are different from novices. Table 1 shows a summary of the techniques we found useful for interactive knowledge acquisition.

We divide the techniques into seven thematic phases, each with a different emphasis on how memory is used: collecting experience, organizing experience in memory, introspect, retrieval of experience from memory, use of experience during problem solving, knowledge refinement using experience, and memory management.

1. Collecting Experience

Notice meaningful patterns of information based on past experience

Skillful experts can make use of their experience in distinguishing meaningful features among the features they observe while novice experts may not recognize them easily (Bransford et al., 2000).

Many KA systems pre-define specific events that they want to recognize in assisting users. For example, systems define specific types of errors and user mistakes they want to check (Clark et al., 2001, Blythe et al., 2001, McGuinness et al., 2000). Some systems exploit prototypical sequences of user actions in recognizing the status of the KA process (Tallis & Gil 1997). Recognized problems or status are used in generating suggestions that are suitable for the situation.

Adopting the technique more closely, KA systems could dynamically learn what features are more meaningful based on their experiences such as situations where users needed more assistance. They also could learn to recognize situations where problem solving may dynamically change and help the user modify relevant problem solving knowledge. These additions could be used together with pre-defined event types.

2. Organizing Experience in Memory

Organize around common parts of similar episodes

Many models of memory organize experiences based on similar properties they share (Kolodner 1993; Schank 1982), often based on predictions of future retrieval tasks. Some of the KA tools actually use similar strategies in organizing and indexing exemplars based on similarities and differences in their domain features (Bareiss et al., 1990). These indices are used in refining category concepts that are being built with the user.

This technique could be more widely used in organizing other experiences in interacting with users. For example, episodes of system assistance that led to successful results (e.g., acquired knowledge leading to successful problem solving) could be distinguished from unsuccessful episodes, and could be organized based on similarities and differences of the situations and the interactions. By relating similar interactions that were successful (such as similar useful hints), the system may provide better assistance.

Organize experience in ways that reflect deep understanding of the problems

Human experts organize their experience around important ideas or concepts (Bransford et al., 2000). Likewise, competent learners organize their memory in ways that lead to conceptual understanding (Tulving, 1983;Brown, 1987).

As described above, existing KA systems identify specific events they keep track of in assisting users. Those events are defined in terms of the knowledge they acquire and the type of assistance provided to the user. In order to support deeper understanding of the problems, the system needs to capture the context where such events occur and whether the system responses actually improve the knowledge.

3. Introspect with Experience

Introspect on failure & success: remember and predict failures and successes

Some machine learning systems develop a set of typical failures and their causes, and use them in driving their learning activities (Cox & Ram 1999). Some case-based reasoning (CBR) systems use past failures in predicting new failures and apply remedy recipes to avoid similar problems (Hammond 1989).

Most existing KA systems use various techniques to detect errors and failures (Davis, 1979; Clark et al., 2001, Blythe et al., 2001, McGuinness et al., 2000). However, the systems have limited understanding on relations between failures (such as similarities among them) and how past failures could be used in preventing new ones. To be truly introspective, the system should be *self-aware*, accessing and reasoning on relevant aspects of knowledge acquisition and knowledge use, such as how failures are related to each other.

4. Retrieval of Experience

Goal driven memory search

Some case based reasoning approaches formulate goals (such as goals of adaptive problems solving) in performing memory search and learn how to perform goal-driven search (Leake 1994). The retrieved results support adaptation and problem solving. Metacognitive learners control cognitive activities to ensure that a cognitive goal has been met (Brown 1987).

Interactions in existing KA systems are driven by various implicit goals developed in the design of the tools. That is, most of their goals (such as find missing definitions, detect conflicting definitions, etc.) are buried in the design and they influence the interaction depending on how they are implemented in the underlying code. By adopting goal driven approaches and making their goals explicit, the systems could reason on how various events are related to the goals and drive the interactions towards achieving the goals.

Retrieve knowledge based on context

Human experts are good at retrieving important ideas and useful concepts in a given context (Bransford et al., 2000). Some of the cognitive systems are built to perform problem solving through problem spaces, and past experiences are saved and retrieved with respect to problem solving context (Rosenbloom et al., 1993).

Some of the existing KA systems exploit problem solving context in guiding the acquisition process. For example, given generic problem solver or inference structure defined for particular type of tasks (such as configuration design tasks), acquisition systems help user enter domain-specific knowledge that play specific roles during problem solving (Marcus & McDermott 1989). Typical KA tasks and their sequences also provide hints on the kinds of help needed by the user (Tallis & Gil 1997). Some acquisition systems exploit problem solving context in inducing proper representation of user input that fit into the situation at hand (Huffman & Laird, 1995). Some other systems divide knowledge bases into separate micro-theories to provide context boundaries (Witbrock et al., 2003). These existing techniques could be exploited in using memory as well as in assisting the user during knowledge authoring. Acquisition systems could use context in relating past KA events and build additional strategies for assisting users in certain problem solving context. For example, there could be common difficulties the user has in entering the same type of design parameters in configuration design domain, and the system could build a general strategy for them.

5. Problem Solving with Experience

Use similar experience in solving new problems

Like human learners use familiar experiences in learning and problem solving (Ausubel 1968), given a problem to solve, case-based reasoners find most similar cases from the memory, predict possible directions from the retrieved cases, and generate solutions using various adaptation techniques (Kolodner 1993, Hammond 1989; Veloso & Carbonell 1993).

Interactive KA tools could be improved by adopting these techniques by retrieving similar interactions in the past, forming predictions on possible directions from retrieved experiences, and generating effective ways to guide the user based on the similarities and differences between the past situations and the current situation at hand.

Use experience to perform similar tasks more efficiently

Some cognitive systems use memory for improving efficiency of performing similar tasks (Rosenbloom et al., 1993; Veloso et al., 1995). The same result can be produced with stored productions instead of going through multiple problem solving steps.

Improving speed of interaction has not been a focus of KA tools. However, if acquisition systems could compile the steps that the user has gone through to reach a desired state (i.e., correct/useful knowledge) or an undesirable state, then the systems may use them to distinguish interaction paths, which ones are better than the others.

Present varying levels of flexibility in their approach to new situations

Adaptation is a key element of cognitive development (Piaget & Inhelder 1973). Given a new situation, skilled experts can apply their skills to the new situation and adapt rapidly to new demands (Bransford et al., 2000). Some of the temporal reasoning approaches can learn how to detect situation patterns where things may change (Oates 2002). Also some CBR systems capture task-dependent adaptation knowledge from past cases (Craw et al., 2001).

Adopting these principles, interactive KA systems could learn how to recognize situation changes and develop adaptation strategies for them. The acquisition systems could explicitly capture the situations where problem solving changed in the past and also learn the kinds of modification that improved the knowledge in such situations. These could be used in helping users modify knowledge when similar problem solving changes may occur.

6. Knowledge Refinement with Experience

Generalize or specialize knowledge based on range of application observed

As a part of learning, some cognitive systems define generalization and discrimination of knowledge in memory as the knowledge becomes broader or narrow in its range of application (Anderson & Lebiere 1998). Some models of memory (Kolodner 1993; Schank 1982) support mechanisms to generalize or specialize learned knowledge depending their uses and reliability observed.

Some of the existing KA systems perform similar knowledge refinement process by generalizing or specializing rules to correctly classify a suite of test cases (Ginsberg et al, 1985) or finding correct categories based on differences and similarity of the exemplars (Bareiss et al., 1990).

The knowledge refinement approaches could be further improved by incorporating adaptation strategies. When the problem solving changes (e.g. changes in how to classify examples due to new findings), the acquisition systems will need to distinguish invalid or irrelevant experiences, avoiding modifications due to obsolete examples.

Strengthen knowledge that is used more often

Another interesting aspect of some cognitive systems is that they can adjust strength of knowledge (such as rules) based on their associative strength measured from related concepts (Anderson & Lebiere 1998). Each rule can be used differently depending on the degree of the strength.

By adopting this technique, interactive KA systems could measure strength of user entered knowledge based on how they were created and how they were exercised in problem solving. When knowledge has been successfully tested or used multiple times, it can be more confidently applied than it has not been. Likewise, acquisition systems could evaluate usefulness of suggestions in a situation based on how they were followed by the user and how they improved the knowledge. More useful suggestions could be more confidently used.

7. Memory Management

Forget low utility knowledge

There have been approaches to assess utility of knowledge in the memory in terms of competence level changes due to the knowledge, and use the assessment in discarding unnecessary information in the memory (Smyth & Keane 1995; Zhu & Yang 1999).

Most existing KA systems do not concern management of unnecessary knowledge. However, similar assessment will be very useful when there are dynamic changes in the problem solving and associated knowledge needs to be either modified or discarded in order to maintain or improve the overall competence level.

Summary of Analysis: Developing Reflection Capabilities for Interactive Knowledge Acquisition

Table 1 shows a summary of the memory-inspired techniques we found useful for interactive knowledge acquisition. The table also shows related aspects in knowledge acquisition approaches in the 'Related KA aspects' column. First of all, most existing tools either ignore past interactions or use them in limited ways (such as to refine concept definitions with a suite of past examples). Empty cells in the column indicate lack of related capabilities in existing acquisition tools. Most tools do not directly support memory-based reasoning.

However, we notice that many of the techniques in existing acquisition systems could be exploited in adopting memory-inspired techniques. For example, most systems diagnose and detect errors in user entered knowledge and have strategies to help users fix them. By remembering how different user mistakes were handled in the past and relating mistakes based on their similarities and differences of the situations where they occur, the systems could help users avoid similar mistakes in similar situations. Other existing approaches to assist users could be exploited in a similar fashion as described above.

One of the key issues in supporting most of the above memory-inspired techniques is being *self-aware*, accessing and reasoning on interesting aspects of knowledge acquisition and knowledge use. If the systems want to be more proactive in preventing errors, they should be able to keep track of past user mistakes and relations between them. Another common requirement in supporting abovementioned KA improvements is being able to *relate* current situation to similar situations in the past. The relations could be used in evaluating alternative options (such as finding effectiveness of the options in handling similar problems) to guide the user. Finally, the systems should provide approaches to recognize potential changes in the relations and build strategies to cope with the changes. As indicated in several places above (phase 1, 5, 6, and 7), adaptation is a key element of memory-inspired techniques, and in providing assistance during knowledge acquisition, the system should recognize what knowledge needs to be modified and how knowledge authoring tasks should be changed.

ECHO: Reflection Patterns for Interactive Knowledge Acquisition

Based on our observation of memory based strategies described in the previous section, we have developed a knowledge acquisition framework called Echo (mEta-Cognitive History analysis and Organization). Echo supports the following capabilities.

- being self-aware, accessing and reasoning on selected aspects of knowledge acquisition and knowledge use in assisting users.
- relating the current situation to similar situations in the past and assessing the levels of confidence in pursuing alternative options based on the relations.
- recognizing dynamic changes in the problem solving and deciding how to guide users in modifying and using relevant knowledge.

Echo adds an additional layer to existing tools and explicitly keeps track of knowledge acquisition and knowledge use episodes through a set of declarative *reflection patterns*. Reflection patterns define a set of abstractions of knowledge acquisition and knowledge use episodes that the system makes use of in assisting users. Each episode is a sequence of basic knowledge acquisition and knowledge use events such as the user ignored a suggestion and then the problem solving failed. The system assesses the levels of confidence in providing a suggestion based on its *supporting* and *opposing* episodes captured in the reflection patterns. Any changes that are noticed (e.g. confident knowledge became inconsistent with problem solving results) and associated modifications are explicitly captured in the reflection patterns and are used in guiding the user. A prototype system has been developed for a domain of interactive scheduling where the user incrementally builds scheduling constraints and the user entered constraints assist users during scheduling. Since scheduling constraints can change over time, the system should be able to assist users in making associated modifications. The details of the system are described in (Kim 2005).

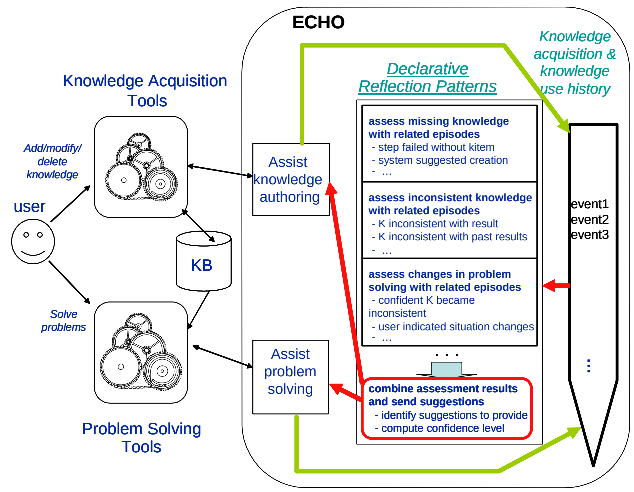

Figure 1: Knowledge Acquisition and Knowledge Use with Echo

Summary and Future Work

We have presented an analysis of memory-inspired techniques in cognitive science and computers science research in terms of how the techniques could be useful in the context of developing interactive KA tools. We have noticed that although most existing tools either ignore past interactions or use them in limited ways, many aspects of existing acquisition approaches can be related to memory-based reflection and the related KA approaches could be exploited in adopting memory-inspired techniques. We believe that the resulting reflective capability will play central role in making the systems truly proactive assistants.

We have developed a novel extension to existing KA tools where the system organizes memory with as set of declarative reflection patterns and uses them to recognize selected knowledge authoring and knowledge use episodes. The reflection patterns are also used in assessing how the knowledge acquisition tasks should be done and how to guide the user. They allow the system to 1) be aware of interesting knowledge acquisition and knowledge use episodes 2) relate current episodes to past similar episodes and generate suggestions based on related episodes, and 3) assess dynamic changes in the problem solving.

We plan to investigate how Echo's reflection patterns can be used in combination with the KA strategies used by existing KA tools. For example, existing knowledge refinement algorithms can be associated with the dynamic changes noticed by Echo. There has been work on developing a dialogue tool for interactive knowledge acquisition (Kim & Gil 2003). The tool incorporates the dynamics of tutor-student interactions in order to support users in their role of tutors of computers, making acquisition tools better students. Assessment of user built knowledge and their progresses over time in Echo could be combined with other dialogue strategies and be used in structuring the front-end interactions for knowledge authoring. We also plan to perform intensive evaluation of user interactions with Echo in terms of assessing its reflective capabilities and effectiveness of assistance provided.

Acknowledgments

We thank Yolanda Gil, Jim Blythe, and Tim Chklovski for helpful comment on earlier drafts. We gratefully acknowledge funding for this work by DARPA under contract no. NBCHD030010.

References

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

Ausubel, D. (1968). *Educational psychology: A cognitive approach*. New York, Holt, Rinehart and Winston.

Blythe, J. Kim, J., Ramachandran, S., and Gil, Y. (2001). An Integrated Environment for Knowledge Acquisition. *Proceedings of IUI-2001*.

Bransford, J., Brown. A., & Cocking, R. (2000). *How people Learn: Brain, Mind, Experience, and School: Expanded Edition*, National Academies Press.

Brown, A. L. (1987). Metacognition, executive control, self-regulation, and other more mysterious mechanisms. In F. E. Weinert & R. H. Kluwe (Eds.), *Metacognition, motivation, and understanding* (pp. 65-116). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Clark, P., Thompson, J., Barker, K., Porter, B., Chaudhri, V., Rodriguez, A., Thomere, J., Mishra, S., Gil, Y., Hayes, P. and Reichherzer, T. (2001). Knowledge Entry as the Graphical Assembly of Components. *Proceedings of K-Cap-2001*.

Cox, M. T., & Ram, A. (1999). Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence, 112*, 1-55.

Craw, S., Jarmulak, J., & Rowe, R. (2001) Learning and applying case-based adaptation knowledge., *ICCBR-2001*.

Davis, R. (1979). Interactive Transfer of Expertise: Acquisition of New Inference Rules. *Artif. Intell*. 12(2),.

Eriksson, H. Shahar, Y., Tu S. , Puerta, A., & Musen, M. (1995). Task modeling with reusable problem-solving methods. *Artificial Intelligence* 79:293-326.

Gil, Y. & Kim, J. (2002). Interactive Knowledge Acqusition Tools: A Tutoring Perspective. *In Proceedings of CogSci-2002*.

Ginsberg, A., Weiss, S., Politakis, P. (1985). SEEK2: A Generalized Approach to Automatic Knowledge Base Refinement. *Proceedings of IJCAI-1985*: 367-374.

Hammond, K. (1990). Case-based planning: a framework for planning from experience, *Cognitive Science* 14.

Huffman S. & Laird, J. (1995). Flexibly instructable agents. *Journal of Artificial Intelligence Research*, 3:271-324.

Kim, J. (2005). Reflection Patterns for Interactive Knowledge Capture, internal report, USC-ISI.

Kim, J. & Gil, Y. (2000). Acquiring Problem-Solving Knowledge from End Users: Putting Interdependency Models to the Test. *Proceedings of AAAI-2000*.

Kim, J. & Gil, Y. (2003). Proactive Acquisition from Tutoring and Learning Principles, *Proc. of AIEd*- 2003.

Kolodner, J. (1993) *Case-Based Reasoning*, Morgan Kaufmann Publishers.

Leake, D. B. (1994). Towards a computer model of memory search strategy learning. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*.

Marcus , S. & McDermott, J. (1989). SALT: A knowledge acquisition language for propose-and-revise systems. *Artificial Intelligence* 39(1):1-37.

McGuinnes, D., Fikes, R., Rice, J., & Wilde, S. (2000) An environment for merging and testing large ontologies. *Proceedings of KR-2000*.

Oates T.(2002). Peruse: An Unsupervised Algorithm for Finding Recurring Patterns in Time Series., *Proceedings of ICML-2002*.

Piaget, J. & Inhelder, B. (1973). *Memory and intelligence*. NY: Basic Books.

Pool, M., Murray, K.. Fitzgerald, J,. Mehrotra, M., Schrag, R., Blythe, J., Kim, J., Chalupsky, H., Miraglia, P., Russ, T., & Schneider D. (2003), Evaluating SME-Authored COA Critiquing Knowledge, *Proc. of K-CAP 2003*.

Rosenbloom, P., Laird, J., & Newell, A. (1993). *The Soar papers: research on integrated intelligence*., MIT Press.

Schank, R.C. (1982). Dynamic Memory: *A Theory of Reminding and Learning in Computers and People*. Cambridge University Press.

Smyth, B. & Keane, M. (1995). Remembering to Forget: A Competence-Preserving Deletion Policy for Case-Based Reasoning". *In Proceedings IJCAI-1995*.

Tallis, M. & Gil, Y. (1997). A Script-Based Approach to Modifying Knowledge Bases, *Proc. of AAAI-1997*.

Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.

Veloso, M. & Carbonell, J.(1993). Derivational Analogy in PRODIGY: Automating Case Acquisition, Storage, and Utilization. *Machine Learning*, 10(3):249-278.

Veloso, M., Carbonell, J., Perez, A., Borrajo, D., Fink E., & Blythe, J.(1995), Integrating Planning and Learning: The PRODIGY Architecture, *Journal of Theoretical and Experimental Artificial Intelligence*, 7(1), 1995.

Witbrock, M., Baxter,D., Curtis,J., Schneider,D., Kahlert, R., Miraglia, P., Wagner,P., Panton,K., Matthews, G., Vizedom A. (2003) An Interactive Dialogue System for Knowledge Acquisition in Cyc.. *IJCAI-03 Workshop on Mixed Initiative Intelligent Systems*.

Zhu, J. & Yang Q. (1999).Remembering to Add: Competence-preserving Case-Addition Policies for Case Base Maintenance. *Proceedings of IJCAI-1999*.