

Changing the Viewpoint: Re-Indexing by Introspective Questioning

R. Oehlmann P. Edwards D. Sleeman

Department of Computing Science
University of Aberdeen
Aberdeen AB9 2UE, SCOTLAND, UK
(oehlmann, pedwards, sleeman)@csd.abdn.ac.uk

Abstract

Various cognitive and computational models have addressed the use of previous experience to understand a new domain. In particular, research in case-based reasoning has explored the ideas of retrieving and adapting previous experience in the form of cases, which can only be retrieved when they are appropriately indexed. In contrast to learning new indexes, re-indexing of existing cases has received little attention. The need for re-indexing a case arises when a previous situation has been incorrectly or incompletely understood. We describe a novel approach to re-indexing which integrates results from two different areas: multiple viewpoints used in intelligent tutoring systems and introspective questioning used in metacognitive activities. Furthermore, we apply ideas from Case-Based Reasoning to the re-indexing process itself. The revised index can be tested by active interaction with the agent's environment. An example of our implementation, IULIAN, will illustrate the re-indexing process.

Re-Indexing, Multiple Viewpoints, and Self-Questions

When people are exposed to an unfamiliar task such as understanding a new device, they are sometimes reminded of previous experience in other domains. Cognitive models of reminding have evolved in an attempt to explain this phenomenon (e.g., Ross 1989). The concept of re-use of previous experience has been exploited in the paradigm of case-based reasoning where previous experience is usually represented by cases (see Kolodner, 1993 for an overview). There is a strong mutual relationship between cognitive models of reminding and case-based reasoning. Perhaps the most crucial issues in case-based reasoning are the retrieval and modification of previous cases, both of which use indexes which describe the content of a case and reflect the current understanding of the reasoner. Various researchers have addressed the problem of learning indexes (Bhatta & Goel, 1993; Robinson & Kolodner, 1991; Osgood & Bareiss, 1993; Ram, 1993b). However, little attention has been paid to the revision of *existing* indexes. A rare exception is the Meta-AQUA system which addresses retrieval failures by specializing and generalizing indexes (Ram & Cox, 1994; Cox, 1994). The need for index revision arises when the understanding of a previous experience is regarded as faulty or incomplete in the light of a new situation. This need can be addressed by the use of metacognitive skills and multiple viewpoints.

Experiments have shown that human understanding is improved through the use of metacognitive skills which control cognitive processes (Wong, 1985), and multiple viewpoints which allow the learner to reason from different perspectives (Wenger, 1987).

Self-questioning is an important metacognitive skill which has been successfully used to improve reading comprehension and problem solving (Wong, 1985). Introspective questioning, i.e., the generation of questions about one's own reasoning as opposed to questions about a given domain, forms a part of the self-questioning approach. In addition to self-questioning, the understanding process can be guided by supporting multiple viewpoints (Wenger, 1987). This idea has been utilised in intelligent tutoring systems. A viewpoint describes certain aspects of a given object. For example, a watch might be regarded as a measuring instrument or as an item of jewellery. We view the idea of multiple viewpoints as related to the idea of extrapolative conceptual change (Ram, 1993a). This type of conceptual change occurs in unfamiliar settings, where existing concepts are no longer applicable. Therefore, they have to be adapted to accommodate the new setting. The idea of extrapolative conceptual change has been investigated in the ISAAC system (Ram, 1993a) which attempts to understand stories from the science fiction literature.

We argue that an integration of multiple viewpoints and self-questioning can be achieved by combining two methods: addressing viewpoints with respect to content-oriented indexing of previous cases, and through use of introspective question asking and answering as a tool to change a given viewpoint. We have implemented an exploratory discovery system, IULIAN, which addresses these issues. The main task of the system is the discovery of new explanations to revise an initial causal model¹. This task is accomplished by generating a sequence of questions, answers and experiments (Oehlmann, Sleeman, & Edwards, 1993). The entire process is influenced by the agency concept (Hammond, Marks, & Converse, 1993), i.e., questions, answers, and experiments are generated by case-based planning, and the principles of learning from failure and enforcing the stability of the environment are

1. The importance of self-explanations has been recognised by various researchers; see for example (Chi, *et al.*, 1989; Pirolli & Bielaczyc, 1989).

applied. An important component of this concept is the agent's interaction with the environment which is important for initialising changes of the viewpoint and for evaluating the correctness of the new viewpoint.

During the processes of questioning, answering, and experimentation, an experimental result might establish a new view of a previous case; i.e., the index which describes the previous case is incomplete or incorrect and has to be revised. The revision can be achieved by adding a new view, i.e., a set of additional index descriptor values, to the index. The generation of the new view is based on the previous case, the new observation, and previous experience with indexing this type of observation.

We will describe our approach by means of an example involving a mechanical clock. The top level of the example is presented in Section 2. Changing the viewpoint is based on the generation of introspective questions; in Section 3, we describe a case-based planning approach to question generation. Section 4 elaborates on the clock example and explains the question-based re-indexing method. Finally, in Section 5, we will evaluate our approach with respect to the general problem of learning indexes.

Example

We will motivate the idea of changing the viewpoint by describing a problem involving a mechanical clock. In addition to the mechanism for moving the hands, the clock has a chime mechanism. The main relations of the chime mechanism are indicated in Figure 1. We assume that the discoverer is attempting to understand how the hammer movement is initiated. In addition, we assume that several wheels have already been studied.

The discoverer knows that the wheels transfer the force of the hammer spring wheel to the pin-wheel. The movement of one wheel drives the movement of the neighbouring wheel. The discoverer attempts to identify a chain of consequences responsible for the hammer movement and asks an appropriate question. At this point, the discoverer does not know the answer and performs an experiment. The discoverer attempts to enable the movement of the pin-wheel by moving the clock hands into the appropriate position and observes the movement of the two-wing wheel. During this process, the discoverer makes an important observation: a lever is moved into the orbit of the pin-wheel, touches the pin, and interrupts the movement of the pin-wheel. The unexpected observation leads to a change of the viewpoint and subsequently to a change of the discoverer's view. Rather than reasoning about the consequences of the *moving* pin-wheel, the reasoning process focuses now on the consequences of the *stopping* pin-wheel. This change has to be reflected by a change of the indexes because the index describes how an experiment has been viewed in the past. Before we describe the details of the re-indexing process, we have to explain how self-questions can be generated as the basis of this process.

Question Planning

IULIAN uses the planning of self-questions, answers and experiments to model reasoning about plans and actions (Oehlmann, *in preparation*). The main task of the system is the discovery of new explanations to revise an initial theory. The basic conceptual knowledge structures are cases and plans.

A case contains objects and relations between objects. Objects are represented as Memory Units (MU)² which

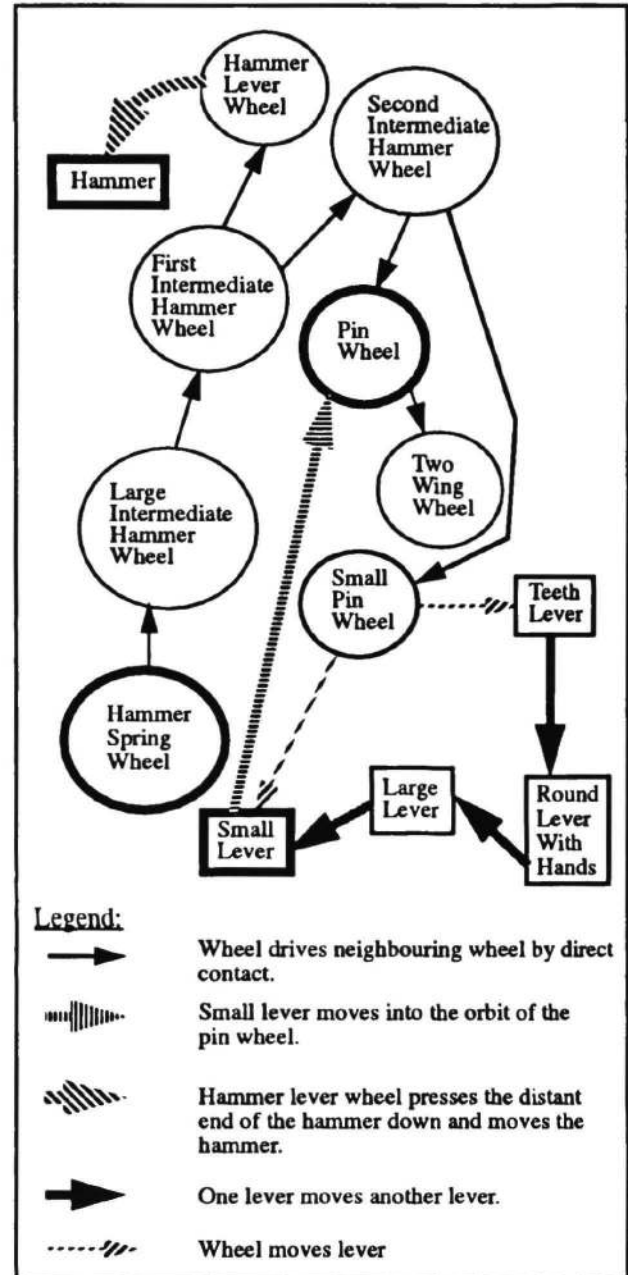


Figure 1: Relations in the Clock Example

2. MUs are similar to the Universal Index Frame (Schank & Osgood, 1990).

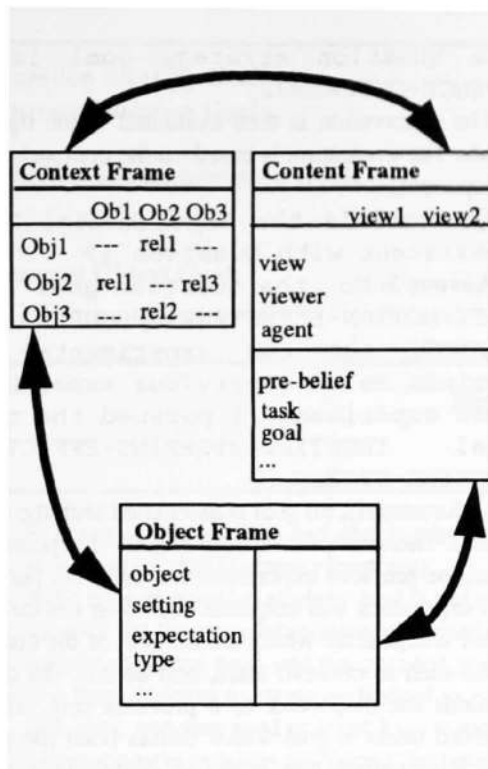


Figure 2: Memory Unit

```
(plan :name identify-effect
:collector "What does the PIN-WHEEL turn?"
:collector1
:collector2
:collector3
:bindings ((focus-object1 pin-wheel))
:question-goals (identify-motion-effect@neighbouring-object)
:question-strategy-goals (consequence-checking)
:recovery-goals
:question-failures
:expectation-failure
:experimentation-goal (enable-movement@focus-object
observe-movement@neighbouring-object)
:failures
:steps
:s-name add-question-word
:action (add-substructure "What " collector1)
:s-name add-question-mark
:action (add-end-substructure "?" collector1)
:s-name add-auxillary-verb
:action (add-substructure "does " collector2)
:s-name add-verb
:action (add-end-substructure "turn " collector2)
:s-name transfer-predicate
:action (transfer-before collector2 "?" collector1)
:s-name add-article
:action (add-substructure "the " collector2)
:s-name add-noun
:action (add-end FOCUS-OBJECT1 collector2)
:s-name transfer-and-insert-subject
:action (transfer-insert collector2 "turn " collector1))
```

Figure 3: Question Plan

contain an object frame, a context frame and a content frame (Figure 2). The context frame describes the context in which the object occurs represented by a set of relations. The content frame comprises several sets of intentional descriptor values referred to as *views*. The object frame comprises general information about the object.

Question plans are used to apply case-based planning techniques to the generation of single questions. The elementary actions the planner executes combine question substructures to form a complete question. For example, the question "What does the PIN-WHEEL turn?" can be built by combining the substructures "What ", "does", "turn " , "the", and OBJECT1. OBJECT1 is a variable which can be instantiated with the string "pin wheel". A question plan has two main parts: the set of descriptors used for indexing the plan and a sequence of steps, see Figure 3. The plan is retrieved by matching its index with the current situation; this is characterised by the goals the discoverer pursues in asking the question. Additional slots in the head of each question plan contain a list of variable instantiations referred to as *bindings* and a set of *collector* slots. The bindings are used to instantiate variables in the step actions. In the collector slots, intermediate results are stored during the question formation process. Each planning step has *precondition*, *goal*, and *action* slots to ensure correct plan execution³. If plan execution fails, the usual explanation-based repair

3. Note that this is a reduced version of the original plan where each planning step contains additional slots for preconditions and goals.

mechanisms are employed, see Hammond, Marks, & Converse (1993). These mechanisms use pre-stored repair rules and are based on preconditions and goals.

It is an important advantage of the case-based planning approach that new questions can be learned by modifying previous question plans. Answers are generated in a similar way; however, steps in *answer plans* might have particular actions which retrieve knowledge from the case base needed to form an answer. *Question strategies* are higher level plans which organise the execution of single question and answer plans to generate questions in a particular sequence. The same basic plan structure used for question and answer plans has been employed for *experimentation plans*, although the index vocabulary differs (Oehlmann, Sleeman, & Edwards, 1993). Experimentation plans describe the steps which have to be executed in order to perform an experiment. The experimental setting and the result of plan execution are stored as a new case.

Re-indexing by Changing the Viewpoint

In this section, we describe our approach to case re-indexing in terms of multiple viewpoints which are expressed by elements of the index vocabulary such as goals, beliefs, and themes (Schank, & Abelson, 1977). Similar to a physical viewer observing an object from a given viewpoint, the discoverer can consider objects from the perspective of different goals or beliefs. The same object viewed from the perspective of different goals would be interpreted in terms of different tasks, plans, or

side effects. Therefore, we define a *view* as an interpretation of a set of objects from a given *viewpoint*.

The discoverer supports re-indexing by generating questions about index components. We will exemplify our approach by elaborating the clock example described above. The underlying question strategy is to identify a chain of consequences responsible for the hammer movement. The current viewpoint is characterised by goals associated with the pin-wheel, the question about the pin-wheel, the current question strategy, and the current experimentation plan. It is part of the question strategy to identify the effect of the motion of the pin-wheel on the neighbouring wheel. Therefore, this objective determines the current question goal and the current question.

Question 1: What does the pin-wheel turn when the hammer moves?

At this point, the discoverer does not know the answer and performs an experiment. The experiment attempts first to enable the movement of the pin-wheel by moving the clock hands into the right position. Then it attempts to observe the movement of the neighbouring wheel.

Experiment 1: Focus on the pin-wheel and check which component is turned by the pin-wheel.

Hypothesis 1: The pin-wheel turns and therefore moves the neighbouring wheel.

Experimental Result 1: First the pin-wheel turns and moves the wing-wheel. A lever is then moved into the orbit of the pin-wheel which interrupts the movement of the wheel.

The experimental setting and the experimental result are stored as a new case. The index of this case includes the goal associated with the pin-wheel, i.e. to enable the movement of the neighbouring wheel. The initial situation is expressed by the goals described in Figure 4.

Part of the experimental result, i.e., the movement of the lever and the subsequent stopping of the wheel, is unexpected. The unexpected observation leads to a change of the viewpoint⁴ and subsequently to a change of the discoverer's view according to the steps given in Figure 5. Our approach to changing the viewpoint is realised as a sequence of questions and answers. First the discoverer attempts to identify the current question and question strategy goals⁵.

Question 2: What are the goals I pursue by using the current question and question strategy?

Answer 2: The question goal is IDENTIFY-MOTION-EFFECT@NEIGHBOURING-OBJECT and

4. In Figure 4, the viewpoint is expressed in terms of goals. We note that viewpoints can also be established by themes or beliefs.

5. The distinction between the question which has already been generated and the question goal under which the question plan has been stored is related to the idea that it is sometimes easier to recall a question than the goals or beliefs related to that question.

the question strategy goal is CONSEQUENCE-CHECKING.

The observation is then evaluated in the light of these goals. The evaluation is based on the comparison of index components.

Question 3: Is the Experimental Result 1 consistent with Question 1?

Answer 3: No, the question goal is IDENTIFY-MOTION-EFFECT@NEIGHBOURING-OBJECT, but the experimental result reminds me of a previous experiment. In this experiment, I pursued the question goal IDENTIFY-STOPPING-EFFECT@NEIGHBOURING-OBJECT.

Unfortunately, no goal is associated with the new observation. Therefore, the change of goal-viewpoint should be based on previous experience and the goals pursued when this experience was acquired. Based on similar values for index components which are not part of the current viewpoint such as context, tasks, and beliefs, the observation reminds the discoverer of a previous case. This case is indexed under a goal which differs from the goal under which the current case is stored. Additionally, in this previous situation, a question was asked, a question strategy was used, and an experimentation plan was executed. Therefore, all the indexes under which these knowledge structures are stored are known and can be used to extend the index of the current case.

Previous Experiment: A bicycle is in an upside-down position. The back wheel is moved by turning the pedal wheel and is therefore turning freely. A wooden rod is inserted into the orbit of the pedals: the movement of the wheel is stopped.

This (bicycle) experiment has been performed to support the same question strategy used in the clock experiment. In contrast to the question goals discussed in the clock experiment, the question which initiated the bicycle experiment focuses on the effect on the rear wheel of stopping the pedal wheel⁶. When the (bicycle) experiment was stored as a new case, the goal associated with the pedal wheel was to stop the motion of the neighbouring wheel and the active question goal was to identify the effect of stopping this wheel. Now, a sequence of questions and answers can identify differences between the various goals in both experiments. The answers are able to use the (different) index information in the bicycle experiment to re-index the clock case. The additional indexes created during re-indexing establish a new viewpoint. This new viewpoint enables the discoverer to interpret the task addressed in the clock case as a task to stop the hammer movement.

6. Note that we use the *neighbour* of relation in an abstract sense which includes two wheels which are in direct physical contact as well as two wheels connected with a chain.

Question Strategy Goal:	CONSEQUENCE-CHECKING
Current Question Goal:	IDENTIFY-MOTION-EFFECT@- NEIGHBOURING-OBJECT
Current Experimentation Goal:	ENABLE-MOVEMENT@OBJECT1 OBSERVE-MOVEMENT@- NEIGHBOURING-OBJECT2
Current Object Goal:	SUPPORTING-MOTION@- NEIGHBOURING-OBJECT

Figure 4: Initial Goals

1. Evaluate an observation and decide whether it supports the current question goal and the current question strategy goal (current viewpoint).
2. If the current question strategy goal is not supported, use the observation to retrieve a previous case.
3. Identify for the experimentation plan used to generate the previous case the experimentation goal, the question strategy goal, and the question goal pursued in executing the experimentation plan.
4. Use these indexes to create additional indexes for the current question, question strategy, experimentation plan, and case used in stage 1 (new viewpoint, re-indexing).
5. Retrieve a new question plan using the newly generated index.
6. If necessary, identify a new experimentation plan to answer the new question (new view).
7. Collect the abstract indexes used to retrieve the appropriate questions and answers during re-indexing and build a new question strategy.

Figure 5: Steps for Changing the Viewpoint

This approach to re-indexing can be regarded as adapting the index rather than the case. The indexes used for retrieving question and answer plans are collected and used to build a new question strategy which can be executed when a similar problem arises. Therefore, the reasoner has learned a new re-indexing *method*.

As a result of re-indexing, planning goals related to the previous (bicycle) case are now considered for the current (clock) case. This leads to the retrieval of new plans. Previously, the discoverer had to identify causes which supported the hammer movement, now the discovery task is to identify causes for stopping the hammer movement. In particular, experimentation has to focus on issues related to the small lever which moved into the orbit of the pin-wheel.

Question 8: What are the consequences of interrupting the movement of the neighbouring wheel?

Answering this question and building the subsequent question-answer chain reveal that the hammer movement is interrupted because the lever moved into the orbit of the pin-wheel.

Discussion

We have identified the need for re-indexing cases which arises from incomplete or faulty understanding of previous episodes. Our novel approach to re-indexing is characterised by the integration of introspective questioning and multiple viewpoints, and views re-indexing as reflection of improved understanding of previous concepts rather than as creation of new concepts. Therefore, we view re-indexing by changing the viewpoint as *interpolative conceptual change* which uses existing indexes as boundaries for interpolation. In contrast to Ram's (1993a) description of extrapolative conceptual change in the ISAAC system, we test revised indexes by active interaction with the environment.

Our method of re-indexing represents a bridge between adaptation and indexing. Similar to adaptation, it uses case-based plans which are guided by interactions with the environment and goals derived from the reasoning process (Leake, 1993; Oehlmann, Sleeman, & Edwards, 1993). Similar to indexing, our method attempts to identify index values which represent the content and the context of a case (Schank & Osgood, 1990; Cox, 1994). Common to all three approaches is the re-use of prior concrete knowl-

edge based on memory search. This re-use can be supported by introspective questioning and evaluation of the adapted case or index based on interactions with the environment. These commonalities between adaptation, generation of indexes, and re-indexing point to a common theory of adaptation and indexing on which we will focus our future work. Specifically, we plan to address the following issues:

1. We have developed methods of decomposing the processes of re-indexing (described in this paper) and adaptation (Oehlmann, Sleeman, & Edwards, 1993). We represent these processes as a sequence of interacting plans rather than as a single plan. This approach raises the question of the appropriate level of granularity of the decomposition.
2. Currently, the level of complexity of our examples is comparatively low. Therefore, we will investigate more complex devices; in addition, we hope to move from modelling the behaviour of devices to modelling the behaviour of other agents. We expect that our approach will scale up very well, because complex adaptation tasks are decomposed to the same level of question and answer plans as simple adaptation tasks.
3. Interacting with the environment enables the agent to evaluate indexes based on environmental feedback. Currently, the feedback is obtained from experiments which are not performed for that specific reason. In future, we will investigate how goals derived from the re-indexing process can be used to guide the active experimental evaluation of a new viewpoint.

References

- Bhatta, S., & Goel, A. (1993). Model-Based Learning of Structural Indices to Design Cases. In: *Proceedings of the IJCAI-93 Workshop on Reuse of Design: An Interdisciplinary Cognitive Approach* (pp. A1-A13).
- Chi, M., Bassok, M., Lewis, M., Reiman, P. & Glaser, R. (1989). Self-Explanations: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science*, 13, 145-182.
- Cox, M. (1994). Machines that Forget: Learning from retrieval failure of mis-indexed explanations. In: *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, accepted for publication.
- Hammond, K., Converse, T. & Marks, M. (1993). Toward a Theory of Agency. In S. Minton (Ed.) *Machine Learning Methods for Planning* (pp. 351-396). San Mateo, CA: Morgan Kaufmann.
- Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.
- Leake, D (1993). Learning Adaptation Strategies by Introspective Reasoning about Memory Search. In: *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning* (pp. 57-63). Cambridge, MA: AAAI-MIT Press.
- Oehlmann, R. (in preparation), Exploratory Discovery: A Case-Based Planning Approach to Learning from Self-Questioning and Experimentation, Ph.D. Thesis, University of Aberdeen.
- Oehlmann, R., Sleeman, D. & Edwards, P. (1993). Learning Plan Transformations from Self-Questions: A Memory-Based Approach. In: *Proceedings of the 11th National Conference on Artificial Intelligence* (pp. 520-525). Cambridge, MA: AAAI-MIT Press.
- Osgood, R. & Bareiss, R. (1993). Automated Index Generation for Constructing Large-scale Conversational Hypermedia Systems. In: *Proceedings of the 11th National Conference on Artificial Intelligence* (pp. 309-314). Cambridge, MA: AAAI-MIT Press.
- Pirolli, P. & Bielaczyc, K. (1989). Empirical Analyses of Self-Explanation and Transfer in Learning to Program. In: *Proceedings of the 11th Annual Conference of the Cognitive Science Society* (pp. 450-457). Hillsdale, NJ: Lawrence Erlbaum.
- Ram, A. (1993a). Creative Conceptual Change. In: *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 17-26). Hillsdale, NJ: Lawrence Erlbaum.
- Ram, A. (1993b). Indexing, Elaboration and Refinement: Incremental Learning of Explanatory Cases. *Machine Learning*, 10, 201-248.
- Ram, A. & Cox, M. (1994). Introspective Reasoning Using Meta-Explanations for Multistrategy Learning. In R. Michalski & G. Tecuci (Eds.), *Machine Learning: A Multistrategy Approach*, Vol. IV (pp. 349-377). San Mateo, CA: Morgan Kaufman.
- Robinson, S. & Kolodner, J. (1991). Indexing Cases for Planning and Acting in Dynamic Environments: Exploiting Hierarchical Goal Structures. In: *Proceedings of the 13th Annual Conference of the Cognitive Science Society* (pp. 882-886). Hillsdale, NJ: Lawrence Erlbaum.
- Ross, B. (1989). Reminders in Learning and Instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning*. Cambridge, UK: Cambridge University Press.
- Schank, R. & Abelson, R. (1977) *Scripts, Plans, Goals, and Understanding. An Inquiry into Human Knowledge Structures*. Hillsdale, NJ: Lawrence Erlbaum.
- Schank, R. & Osgood, R. (1990), A Content Theory of Memory Indexing, Technical Report No. 2, Northwestern University, Institute for the Learning Sciences.
- Wenger, E. (1987). *Artificial Intelligence and Tutoring Systems*. Los Altos, CA: Morgan Kaufmann.
- Wong, B. (1985). Self-Questioning Instructional Research. A Review. *Review of Educational Research*, 55, 227-268.