

# Meta-Cognitive Attention: Reasoning about Strategy Selection

Ruediger Oehlmann

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

## Abstract

Both human learners and Case-Based Reasoning systems have applied metacognitive strategies such as self-questioning to improve the learning process. Whereas case-based reasoning systems do not allocate attention to reasoning strategies in order to facilitate strategy selection, previous work on attention in human thinking has focused on the selection of domain objects. We describe a computational model of metacognitive attention which integrates metacognitive approaches in case-based reasoning with the concept of attention which is applied to the reasoning process itself. An example of our implementation, IULIAN, will illustrate the process of allocating metacognitive attention.

## Introduction

People who are confronted with several reasoning tasks at the same time have to make a decision about the task they want to address first. For example, a surprising observation can lead to reasoning about its effect or explaining its cause. Reasoning about an agent's reasoning processes has been studied in metacognition. This term refers to the "active monitoring and consequent regulation" of an agent's cognitive processes (Flavell 1976). That is, the term metacognition describes two distinct but related issues: the issue of knowledge about cognition and the issue of regulating cognition. The notion of knowledge about cognition includes awareness of the resources available to the agent with respect to the demands of the agent's reasoning process; for example, the availability of analogous knowledge during an analogical mapping task. The process of regulating cognition involves self-regulating mechanisms such as planning, monitoring, and self-questioning (Brown, 1975; Wong, 1985). Such meta-cognitive skills should be improved with growing experience. Reasoners should be able to reuse such experience when they select reasoning strategies.

In Artificial Intelligence systems, experience related to a given domain and the ability to be reminded of previous experience have been modeled by the paradigm of Case-Based Reasoning, where previous experience is usually represented as cases (see Kolodner 1993 for an overview). Important stages of the case-based reasoning approach are the retrieval of a previous case which contains a previous problem and its solution, and the adaptation of this case to obtain a solution for the current problem. If this approach is applied to plans rather than to problem solution pairs, we

refer to it as case-based planning (Hammond et al., 1993). Often case-based reasoning systems use memories indexed in terms of prediction failures which occurred during the reasoning process. When the system generates a wrong prediction, the case on which the prediction was based is annotated with a characterization of the failure situation (Schank, 1986). The annotation is used as an index during future case retrieval. As a result, the prediction failure can be avoided in the future.

Recent research in case-based reasoning has addressed the issue of guiding the reasoning process by introspection (Fox & Leake, 1994; Ram & Cox, 1994; Oehlmann et al. 1995). Earlier research has stressed the importance of allocating attention for the human reasoning process. Whereas previous investigations of metacognitive aspects of case-based reasoning did not address the active selection of reasoning strategies by focusing attention, work on attention addresses the selection task. However, the selection process is often limited to the domain level. For example, Ohlsson (1984) proposes an attentional heuristic which selects objects about which no inference has been made at a given state of the reasoning process. In a similar spirit, Bacon (1995) characterizes meta-knowledge about the process of focusing attention on particular domain objects. In contrast to this domain related view, we will describe a computational model which we refer to as metacognitive attention. The model integrates the idea of metacognitive regulation of the reasoning process with the concept of attention. In particular, we will present an operational characterization of metacognitive attention in terms of *intention*, *situation*, and *resource*.

In the remainder of this paper, we present an example involving the chime of a mechanical clock followed by a top-level view of our implementation IULIAN. We then describe our approach to strategy selection based on metacognitive attention, and we explain the knowledge structures used by these processes. Finally, we discuss our approach and indicate options for future work.

## 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. The main relations of the chime mechanism

are indicated in Figure 1.

We assume that the learner attempts to understand how the hammer movement is initiated. In addition, we assume that several wheels have already been studied.




Figure 1: Relations in the Clock Example

The reasoner 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 neighboring wheel. The reasoner attempts to identify a chain of consequences responsible for the hammer movement and asks an appropriate question. At this point, the reasoner does not know the answer and performs an experiment. The reasoner 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. This process leads to 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 enables the reasoner to pursue two different reasoning strategies. One strategy involves a change of the viewpoint and subsequently a change of the reasoner'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. (Oehlmann, et al., 1994). This strategy can be viewed as reasoning about the *effect* of a surprising observation. The other strategy involves reasoning about the *cause* of the surprising observation: the reasoner attempts to explain *why* the lever moved into the orbit of the pin-wheel. Both strategies are based on the same surprising observation and have a high degree of interestingness; therefore the reasoner has to decide which strategy to apply first, i.e. how to direct the metacognitive attention. We will describe the details of this process in the over next section.

### The IULIAN System

The IULIAN system uses the planning of self-questions, answers, and experiments to model reasoning about plans and actions. The main task of the system is the generation of new explanations to revise an initial theory. Figure 2 shows the main modules of the system: *question planner*, *answer planner*, *experiment planner*, *hypothesis formation*, and *introspection planner*. The figure also indicates that the IULIAN system represents an integration of case-based reasoning and case-based planning rather than a single case-based reasoner: the hypothesis formation module is a case-based reasoner whereas the other modules are implemented as case-based planners.

The Question Planner module accepts a problem description as input, generates a question about the problem, and transfers control to the Answer Planner. If a question cannot be answered, the Question Planner, the Answer Planner, and the Experiment Planner can be used to generate additional questions, answers, and experiments which help the IULIAN system to recover from this situation and to provide the knowledge needed to generate the answer. Before an experiment is performed, the Hypothesis Formation module hypothesizes the experimental result. When the actual result is generated, the Hypothesis Formation module determines an expectation failure as the difference between the hypothesis and the actual result. If an expectation failure has been detected, the exploration process is initiated. At its simplest, the process of question and answer generation is based on the Question Planner and the Answer Planner which generate a question about the problem and attempt to answer it. If a question cannot be answered, Question Planner, Answer Planner, and

Experiment Planner can again be used to generate additional questions, answers, and experiments in an attempt to provide the missing knowledge.,




Figure 2: The IULIAN System

During this process of question-based reasoning and experimentation-based activity, questions focus on objects of the domain to be investigated such as *pin-wheel* and *small-lever* (see Figure 1). In addition, the Question Planner generates Questions which focus on the system's reasoning process. If the system asks such a question, using an answer plan to generate a sentence is not sufficient because the answer planner needs additional meta-knowledge which is not available to it. Acquiring this knowledge is the task of the introspection planner (Oehlmann, 1992; Oehlmann, et al. 1995).

The basic knowledge structures of the IULIAN system are experiments and plans which are used as cases. An experiment consists of two components: an experimental setting (e.g. a description of a mechanical clock with wheels, springs, and hands) and the result of an experiment such as the statement that "the hammer of the chime is moving when the large hand is in upright position." Experiments are represented by objects and relations between objects. Objects are represented as *Memory Units*<sup>1</sup> (MU) which contain an *object frame* and a *content frame* (Figure 3). The *context frame* describes the context in which the object occurs represented by a set of relations. The *content frame* contains several sets of intentional descriptor values referred to as views, whereas

general information about the object is stored in the *object frame*.

Question plans are used to apply case-based planning techniques to the generation of single questions. For example, the question "What does the PIN-WHEEL turn?" can be built by combining the substructures "What", "does", "the OBJECT1", and "turn." OBJECT1 is a variable which is instantiated with the string "PIN-WHEEL" during plan execution. A question plan has two main parts: a set of descriptors used for indexing the plan and a sequence of steps. A plan is retrieved by matching its index with the current situation; this is characterized by the goals the system pursues in asking the question. If plan execution fails, the usual explanation-based repair mechanisms are employed, (see Hammond, 1989). 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 may have particular actions which retrieve knowledge from the library of experiments needed to form an answer. For example, an answer to the question "What does the PIN-WHEEL turn?" may be generated by executing the following steps: the first step retrieves the object *pin-wheel* and identifies an object which shares the relation *has-effect@turns* with the object *pin-wheel*. This object is *two-wing-wheel* and is stored in the answer plan. The following steps instantiate two variables with the objects *pin-wheel* and *two-wing-wheel* and combine these variables with the substrings "The", "turns", and "the." The resulting sentence is "The PIN-WHEEL turns the TWO-WING-WHEEL."

The case-based planning approach to generating questions and answers is highly flexible because it only depends on the current situation and the goals the system is attempting to pursue. Moreover, new plans can be generated by adapting existing plans to new situations.

An introspective answer is a sentence generated as response to a question about the reasoners internal knowledge and its internal processes. An answer plan which has to generate an introspective answer contains special steps. Executing such a step results in a call of the introspection planner. This planning process provides the information needed by the answer planner which can then complete the answer. An example of an introspection plan will be discussed in the next section. Introspection plans address different metacognitive tasks such as assessing goals, reasoning strategies, resources needed to perform a given reasoning strategy, failures which occurred during previous reasoning strategies, and conditions which have to be satisfied in order that a strategy can be executed.

In addition, the IULIAN system uses experimentation plans to perform experiments. Experimentation plans describe the steps which have to be executed in order to perform an experiment. The experimental setting and the

1. MUs are similar to the Universal Index Frames (Schank & Osgood, 1990).

result of plan execution are stored as a new case. The same basic plan structure used for question and answer plans has been employed for experimentation plans, although the index vocabulary differs (Oehlmann et al., 1993).

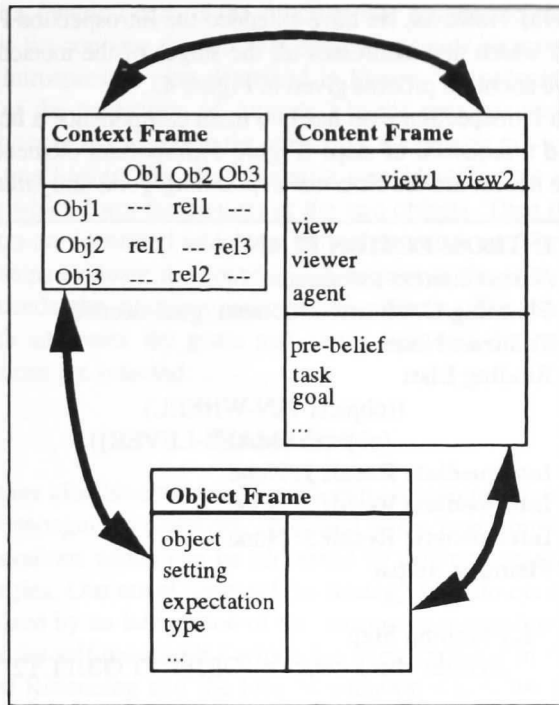


Figure 3: Memory Unit

### Meta-Cognitive Attention

In this section, we describe our approach to strategy selection in terms of metacognitive attention. We will exemplify our approach by elaborating the clock example described above. The strategies mentioned in the example are *changing the viewpoint based on a surprising observation* and *explaining a surprising observation*. We have argued elsewhere that similar to a physical viewer observing an object from a given viewpoint, the reasoner can consider objects from the perspective of different goals and beliefs (Oehlmann et al., 1994; Ohlsson, 1990; Schank & Osgood, 1990). The current viewpoint is characterized by goals associated with the pin-wheel. It is part of the initial reasoning strategy to identify the effect of the motion of the pin-wheel on the neighboring wheel. The strategy changing the viewpoint attempts to identify the effect caused by stopping the pin-wheel. In contrast, the strategy explaining a surprising observation focuses on the causes of the small-lever movement: the movements of the large lever and the *round lever* which control the movement of the *small lever* (Figure 1).

The process of attaching metacognitive attention to the reasoning strategies is based on elements of the index vocabulary such as *cognitive goal*, *cognitive need*, *context*, *strategy projection*, and *strategy resources*. The *cognitive goal* refers

to the objectives of the understanding process. The objective is related to strategies used in the understanding process and generated as a sequence of questions and answers. In contrast, the descriptor *cognitive need* is related to gaps in the current knowledge the reasoner attempts to fill. A reasoning strategy is focusing on one or more objects. Usually, some knowledge related to these objects is already available. This knowledge forms the *context*. The descriptor *strategy projection* indicates whether the result of the strategy can be already anticipated or whether it requires a substantial amount of inference. Finally, the descriptor *strategy resources* describes knowledge structures needed to perform a given reasoning strategy.

The reasoner can select a strategy by acquiring values for these index descriptors. The values are determined by generating a sequence of introspective self-questions and answers (Oehlmann, et al., 1995).

#### Intention Assessment:

1. Assesses the cognitive goals and select questions which may (partially) address these goals.
2. Assesses the cognitive need and select questions which may (partially) address this need.

#### Situation Assessment:

3. Assesses the context related to the focus of all questions under consideration and select the question related to the poorest context.
4. Attempts to project the outcome of the potential strategies.

#### Resource Assessment:

5. Assesses the resources needed to perform the potential strategies.

Figure 4: The Metacognitive Attention Process

The reasoner begins the process of attaching attention by identifying two questions. Each of these questions indicates the begin of a reasoning strategy. To identify these questions the following introspective question is asked:

**Question 0:** What are the questions I consider?

The question plans used to generate these questions can be identified on the basis of the question goals and reasoning goals the reasoner pursues.

**Answer 0:** I consider the questions: "What is the effect of the PIN-WHEEL being blocked?" and "Why does the SMALL-LEVER blocks the PIN-WHEEL?"

After identifying questions which lead to potential reasoning strategies, the reasoner attempts to evaluate these questions in terms of the descriptors given above. All the following questions refer to the set of questions identified in Answer 0. The next question focuses on the descriptor *cognitive goal*.

**Question 1:** What question is related to the current cognitive goal?

Obviously, a reasoning strategy related to the hammer serves the cognitive goal of understanding the hammer movement better than a reasoning strategy which explains the lever movement.

**Answer 1:** The question "What is the effect of the PIN-WHEEL being blocked?" is related to the current cognitive goal CONSEQUENCE-CHECKING.

The next question focuses on the cognitive need.

**Question 2:** What question is related to the current cognitive need?

**Answer 2:** The question "What is the effect of the PIN-WHEEL being blocked?" is related to the current cognitive need IDENTIFY-EFFECT::NEIGHBOURING-WHEELS.

Now the reasoner investigates the context of the objects *pin-wheel* and *small-lever*.

**Question 3:** What question focus is supported by a richer context?

**Answer 3:** The question "What is the effect of the PIN-WHEEL being blocked?" has the richest context.

This answer is given because the reasoner knows the relations between the *pin-wheel* and its neighboring wheels. However, the relations between the *small-lever* and another levers are unknown at this state. Now the reasoner attempts to anticipate the results of the two strategies in question. If the surprising observation would have supported the initial reasoning strategy of checking the consequences of the wheel movements, the result of this strategy could be easier hypothesized. In this situation the reasoner would prefer to explain the surprising observation, because it understands this issue less than the hammer movement.

**Question 4:** What can I project about the cause of the moving lever and the effect of the stopping pin-wheel. The current situation is characterized by a surprising observation which contradicts the initial reasoning strategy. Therefore an appropriate prediction cannot be made.

**Answer 4:** I do not have any knowledge which would allow me to make these predictions.

Finally, the reasoner attempts to evaluate the questions to asked in terms of the resources necessary to perform the strategies.

**Question 5:** What are the questions for which the necessary resources are available?

**Answer 5:** The previous case needed for the question: "What is the effect of the PIN-WHEEL being blocked?" and a previous explanation which can be adapted to answer the question "Why does the SMALL-LEVER blocks the PIN-WHEEL?" are available.

While the assessment of *context* and *projection* do not prefer any of the questions under consideration, the assessment of *cognitive goal* and *cognitive need* lead to the selection of a question which in turn leads to the strategy *changing the viewpoint*.

## Plans for Metacognitive Attention

The process of metacognitive attention described in the previous section can be supported by introspection plans which we have described elsewhere (Oehlmann et al., 1995). However, we have extended the Introspection Planner which now addresses all the stages of the metacognitive attention process given in Figure 4.

An introspection plan has two main components: a header and a sequence of steps (Figure 5). Important elements of the header are the slots *name*, *planning-goal*, and *failures*.

**INTROSPECTION PLAN**  
**Name:** context-assessment  
**Planning Goal:** assess-context::goal-needed  
**Failures:** None  
**Binding List:**  
    ((object1 PIN-WHEEL)  
    (object2 SMALL-LEVER))  
**Intermediate Result 1:** None  
**Intermediate Result 2:** None  
**Intermediate Result 3:** None  
**Planning Steps:**

- 1. Planning Step**  
    **Action:** check-relations OBJECT1 OBJECT2
- 2. Planning Step**  
    **Action:** check-relation-objects
- 3. Planning Step**  
    **Action:** compare-object-goals
- 4. Planning Step**  
    **Action:** compare-object-needs

Figure 5: Introspection Plan

The plan identifier is stored in the name slot. The slot *planning-goal* contains the goals the system attempts to satisfy by executing the plan; the slot *failures* characterizes planning failures which have occurred before. The two slots *planning-goal* and *failures* form the index of the introspection plan, i.e. these slots are used for plan retrieval. In addition, the header includes the slots *binding-list* and *intermediate-result*. The *binding-list* contains pairs of variable names and their values. If an action of a planning step contains variables, the *binding-list* is used to instantiate them. The *intermediate-result* slots are used to store a result which has been generated by a given planning step and which will be used by subsequent steps. A step has four slots: *name*, *precondition*, *goal*, and *action*. The *name* slot serves as an identifier for a given step. In the slot *precondition*, the conditions are described which have to be true before the action given in the *action* slot can be executed. The slot *goal* lists the specific goals

the system attempts to satisfy by executing the action described in the *action* slot. The value of the *action* slot is a list with a function name as first element. The remaining list elements are the arguments which, together with the name, form a function call. If the function has no arguments, the action list contains the function name as a single element. The introspection plan described in Figure 5 is executed as part of the generation of Answer 3 in the previous section. The answer has to assess the context of the objects *pin-wheel* and *small-lever*. The first planning step identifies the relations which form the context of the two objects, then those objects are identified which are part of the relations. The last two steps evaluate the context with respect to the goals and the needs the reasoner attempts to address. The context which addresses the goals and needs best and the related questions are selected.

### Discussion

We have identified the need for selecting reasoning strategies by metacognitive attention. This need arises from surprising observations which can be addressed by different reasoning strategies. Our novel approach to strategy selection is characterized by an integration of the concept of regulating cognition by self-questioning which has been explored in Case-Based Reasoning and the idea of attention which has been explored in domains of human thinking. In addition, we have presented an operational characterization of metacognitive attention in terms of assessing intention, situation, and strategy resources. The criteria considered in every group of assessments are sufficiently abstract to be applied to different strategy selections. Currently, metacognitive attention is only used to select a strategy. Using introspection plans enables the reasoner to store the process of allocating attention as experience. However, the result of this process should be stored as metacognitive knowledge which could be used during metacognitive reasoning. Future work will therefore focus on the representation of the attention focus itself rather than the process of allocating attention.

### References

Bacon, W. (1995). What Everyone Knows about Attention. In *Proceedings of the AAAI Spring Symposium on Representing Mental Models and Mechanisms*, (pp. 1-6).

Brown, A. (1975). The Development of Memory: Knowing, Knowing about Knowing, and Knowing How to Know. In H. W. Reese (Ed.), *Advances in Child Development and Behavior* (Vol. 10). New York: Academic Press.

Flavell, J. (1976). Metacognitive Aspects of Problem Solving. In L. Resnick (Ed.), *The Nature of Intelligence*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Fox, S. & Leake, D. (1994). Using Introspective Reasoning to guide Index Refinement in Case-Based Reasoning. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, (pp. 324-329). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hammond, K., Converse, T., & Marks, M. (1993). Towards a Theory of Agency. In S. Minton (Ed.), *Machine Learning Methods for Planning*, (pp. 351-396).

Hammond, K. (1989). *Case-Based Planning: Viewing Planning as a Memory Task*. New York: Academic Press.

Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.

Oehlmann, R. (1992). Learning Causal Models by Self-Questioning and Experimentation. In *Proceedings of the AAAI Workshop on Communicating Scientific and Technical Knowledge*, (pp. 73-80).

Oehlmann, R., Edwards, P. & Sleeman, D. (1994). Changing the Viewpoint: Re-Indexing by Introspective Questioning. In *Proceedings of the Annual Conference of the Cognitive Science Society*, (pp. 675-680).

Oehlmann, R., Edwards, P. & Sleeman, D. (1995). Introspection Planning: Representing Metacognitive Experience. In *Proceedings of the AAAI Spring Symposium on Representing Mental Models and Mechanisms*, (pp. 102-110).

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.

Ohlsson, S. (1984). Attentional Heuristics in Human Thinking. In *Proceedings of the 6th Annual Conference of the Cognitive Science Society*, (pp. 273 - 276).

Ohlsson, S. (1990). The Mechanism of Restructuring in Geometry. In: *Proceedings of the 12th Annual Conference of the Cognitive Science Society*, (pp. 237-244).

Ram, A. & Cox, M. (1994). Introspective Reasoning Using Meta-Explanations for Multistrategy Learning. In: R. Michalski & G. Tecuci (Eds.) *Machine Learning: A Multistrategy Approach*, 4 (pp. 348-377). San Mateo, CA: Morgan Kaufmann.

Schank, R. (1986). *Dynamic Memory. A Theory of Reminding and Learning in Computers and People*. Cambridge, UK: Cambridge University Press.

Schank, R. & Osgood, R. (1990). A Content Theory of Memory Indexing, Technical Report No. 2, Northwestern University, Institute for the Learning Sciences.

Wong, B. (1985). Self-Questioning Instructional Research. A Review. *Review of Educational Research*, 55, 227-268.

2. Note that Figure 5 shows a reduced version of the original plan. In particular, the slots *precondition* and *goal* in the planning steps are omitted.