

## Introspection Planning: Representing Metacognitive Experience

Ruediger Oehlmann    Pete Edwards    Derek Sleeman

Department of Computing Science

University of Aberdeen

Aberdeen AB9 2UE, SCOTLAND, UK

{oehlmann, pedwards, sleeman}@csd.abdn.ac.uk

### Abstract

Metacognition addresses the issues of knowledge about cognition and regulating cognition. We argue that the regulation process should be improved with growing experience. Therefore mental models are needed which facilitate the re-use of previous regulation processes. We will satisfy this requirement by describing a case-based approach to Introspection Planning which utilises previous experience obtained during reasoning at the meta-level and at the object level. The introspection plans used in this approach support various metacognitive tasks which are identified by the generation of self-questions. As an example of introspection planning, the metacognitive behaviour of our system, IULIAN, is described.

### Introduction

Experimental results from both psychologists and instructional scientists suggest that human cognitive performance can be improved by the use of metacognitive skills such as self-monitoring, prediction, and self-questioning (Wong & Jones, 1982; Wong, 1985). The term *metacognition* has been used in different ways (Gavelek & Raphael, 1985). However there is an increasing tendency to follow Flavell (1976), who related metacognition to *metamemory* (Brown, 1987; Campione, 1987). Metamemory has been defined by Flavell and Wellman (1977) as an agent's knowledge about his/her memory. Flavell (1976) viewed metacognition as referring to the "active monitoring and consequent regulation" of the agent's own cognitive processes. As a result, we can regard the term metacognition as describing two distinct but related issues: the issue of knowledge about cognition and the issue of regulating cognition. The first term includes awareness of the agent's resources with respect to the demands of the agent's thinking; for example, the availability of analogous knowledge during an analogical mapping task. The second issue involves self-regulating mechanisms such as planning and monitoring (Brown, 1975; Brown, 1987; Flavell & Wellman, 1977; Meacham, 1972; Patterson et al., 1980). Such meta-cognitive skills should be improved with growing experience. Reasoners should be able to re-use such experience when they maintain mental models about meta-cognition.

An important instance of metacognition is metacomprehension which addresses those issues of metacognition related to reading (Flavell, 1976; Baker & Brown, 1984). In metacomprehension, but also in other areas such as problem solving,

it has been demonstrated that the use of metacognitive skills such as self-questioning improve learning performance (Schewel & Waddell, 1986; Blank & Covington, 1965).

In Artificial Intelligence systems, experience related to a given domain and the ability to be reminded of previous experience have been modelled by the paradigm of Case-Based Reasoning, where previous experience is usually represented as cases (see Kolodner 1993 for an overview). Important components 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, 1989).

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 characterisation of the failure situation. The annotation is used as an index during future case retrieval. As a result, the prediction failure can be avoided in the future. This view of memory is referred to as *failure driven memory* (Schank, 1982).

Recent research in case-based reasoning has addressed the issue of guiding the reasoning process by introspection. Within this general area, research has mainly focused on the representation of knowledge about cognition.

For example, Ram & Cox (1994) exploited the concept of failure driven memory. When a failure occurs, the system generates a knowledge goal which drives the explanation process. This process utilises meta-explanation-patterns which are causal introspective explanation structures explaining how and why an agent reasons.

Cox and Freed (1994) presented examples of ways in which self-knowledge can be used during the learning process. Such knowledge supports the selection of diagnosis and repair strategies from among alternatives; it enables a system to distinguish between failure hypothesis candidates; and it supports the use of such knowledge across domain borders.

Models of device behaviour were used in the ROBBIE system to refine indexes and to repair reasoning failures (Fox & Leake, 1994; Fox & Leake, 1995). The system

monitors its own reasoning process and compares it with the "ideal" performance of its model.

Oehlmann et al. (1994) used introspective self-questioning to support reasoning from different perspectives. The reasoning perspectives are represented as explicit descriptions. Introspection mechanisms for improving memory search have been discussed by Kennedy (1995) and Leake (1995). Kennedy uses a domain independent representation scheme which supports the memory search during analogical reasoning. Leake views memory search as a reasoning task and describes a framework in which retrieval is guided by introspective reasoning. Moreover, introspective reasoning can be used to learn how to improve the memory search.

It is a common feature of all these approaches that they represent particular types of control knowledge explicitly. However, a general computational model of metacognition requires knowledge structures which can be used in different metacognitive tasks. Such a set of knowledge structures has to address the representation of knowledge about cognition as well as the processes of regulating cognition.

In addition to the task of regulating the current cognitive process, an agent should be reminded of previous regulation processes. We will therefore describe an approach based on case-based introspection planning which satisfies this requirement. An introspection plan contains an index to retrieve the plan from memory and a sequence of actions. The actions are related to cognitive behaviour, e.g. an action might compare the result of a reasoning process with a reasoning goal as part of a monitoring task. The case-based planning approach enables an agent to be reminded of previous subprocesses.

We will describe the approach in the context of our implementation IULIAN which addresses a task involving discovery learning in terms of questions, answers and experiments, all generated by case-based planning. The example used involves electric circuits and systems of water pipes. A top level view of the example is presented in Section 2 followed by a description of introspection plans (Section 3). Section 4 elaborates on the circuit example and explains the question based method of introspection planning. Finally, in Section 5, we will discuss our approach with respect to the general problem of metacognition.

## Example

The approaches of case-based reasoning and case-based planning have been used in the IULIAN system to model discovery learning. We will use the discovery task to exemplify the issue of regulating cognition and representing the regulation process. It has been noted that human discoverers can be characterised as experimenters or theorists according to the discovery strategies employed (Klahr & Dunbar, 1988). Experimenters attempt to deduce regularities from experimental results. In contrast, theorists are able to generate a hypothesis by search processes without any experimen-

tation; the hypothesis is then tested experimentally.

We will use an example to motivate the idea of representing the regulation of cognitive processes. This example describes how an agent decides between experiment driven and theory driven behaviours. The example focuses on two strategies: a reasoning strategy and an activity strategy. The reasoning strategy is based on the generation of a sequence of self-questions and answers, whereas the activity strategy is based on the generation of a sequence of actions suitable to perform experiments. We note that a strategy emerges from the process of organising questions and answers into a sequence; strategies are not based on pre-stored sequences of questions and answers.

We assume that the IULIAN system receives as input the description of an electric circuit with a lamp and closed switch in parallel, and attempts to predict the behaviour of the circuit. The system is reminded of a previous experiment involving a serial circuit with a lamp and closed switch. This previous experience leads to the (wrong) prediction that the lamp in the parallel circuit is *on* because the lamp in the serial circuit is *on*. IULIAN tests the hypothesis by building the parallel circuit and observing its behaviour. The actual result indicates that the lamp is *off*; this constitutes an expectation failure which the system has to explain. The explanation can be generated by using the reasoning strategy *analogical mapping* or the activity strategy *experiment perturbation* (see Figure 1).

During analogical mapping (Oehlmann, et al., 1993), the IULIAN system is reminded of a previous experiment in the domain of water pipes. This experiment has the same topological structure as the parallel circuit experiment and includes a pump, a paddle wheel, and a valve. In addition, it is linked to an explanation which explains the behaviour of the water pipe system. This explanation can be mapped back to the domain of electric circuits. The adapted explanation is suitable to explain the behaviour of the initial parallel circuit.

Alternatively, the system may perturb the retrieved experiment involving a serial circuit. This circuit can gradually be modified until it is identical to the initial parallel circuit. After each modification, the explanation connected to the serial circuit experiment is adapted. Obviously, the circuit behaviour changes when an additional wire is inserted parallel to the lamp in the serial circuit. The system uses the generated circuit experiment to modify its explanation about the circuit behaviour and finally arrives at the explanation for the explanation failure.

Both strategies described in the example lead to a correct explanation. Therefore the IULIAN system has to decide between theory driven analogical mapping and experiment driven perturbation behaviours. This decision can be supported by introspection planning. Furthermore, introspection planning can be used to monitor reasoning and experimentation processes. We will discuss examples of metacognitive decision and monitoring processes in the

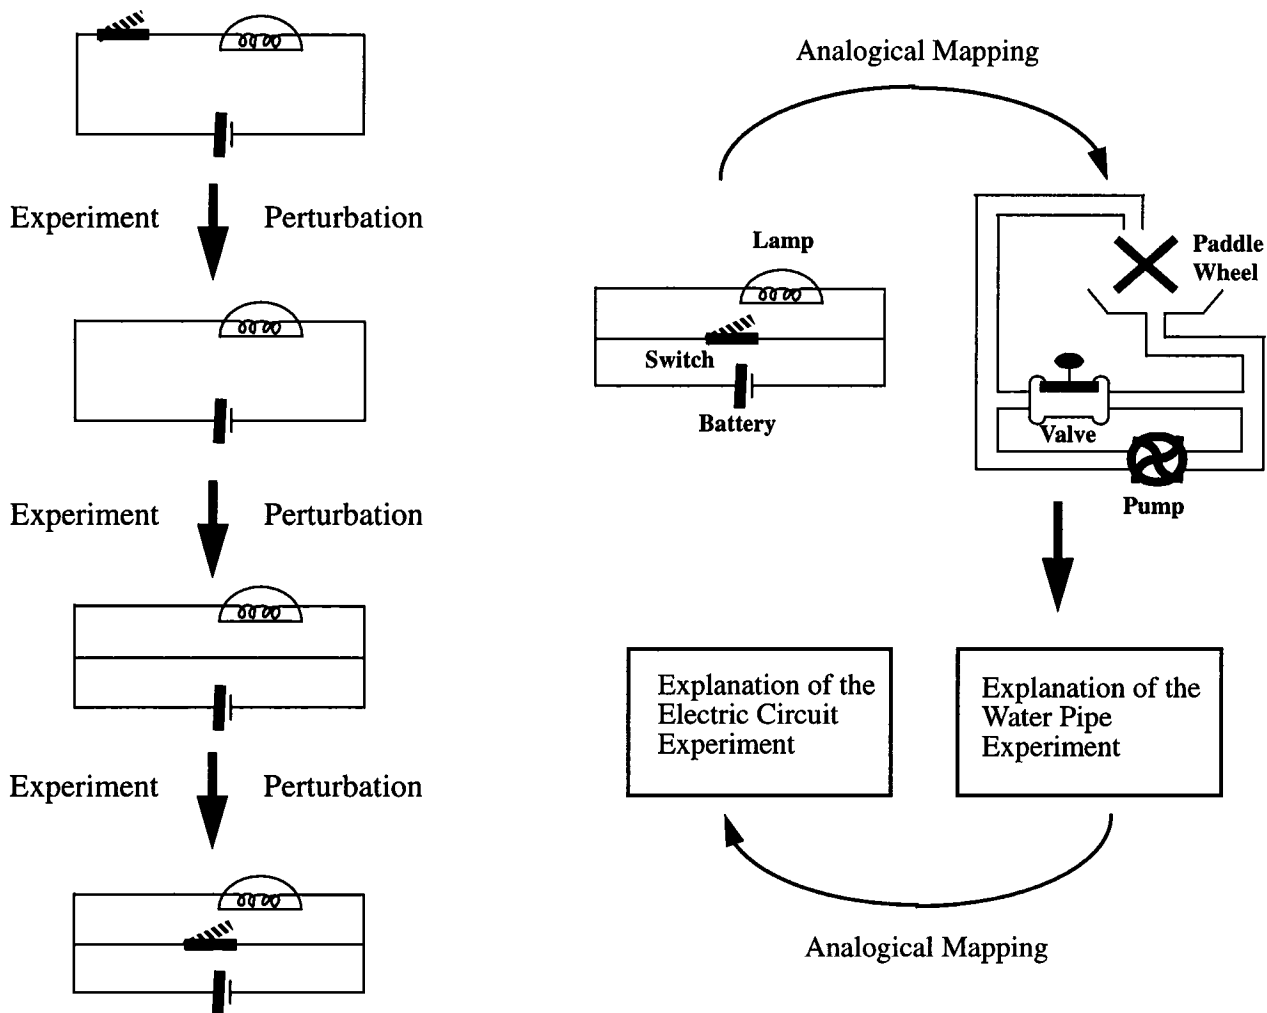


Figure 1: Experiment Perturbation and Analogical Mapping

next section.

### The IULIAN System

The decision process characterised above has been explored in the IULIAN system which uses the planning of self-questions, answers and experiments to model reasoning about plans and actions. The main task of the system is the discovery 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, question strategy planner, and introspection planner. It 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 case-based planners.

The Question Planner Module accepts a problem descrip-

tion as input, generates a question about the problem, and transfers control to the Answer Planner. If a question cannot be answered, the Question Planner and the Experiment Planner can be used to generate additional questions and experiments which helps the IULIAN system 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. The returned answer should be wrong because if the correct information

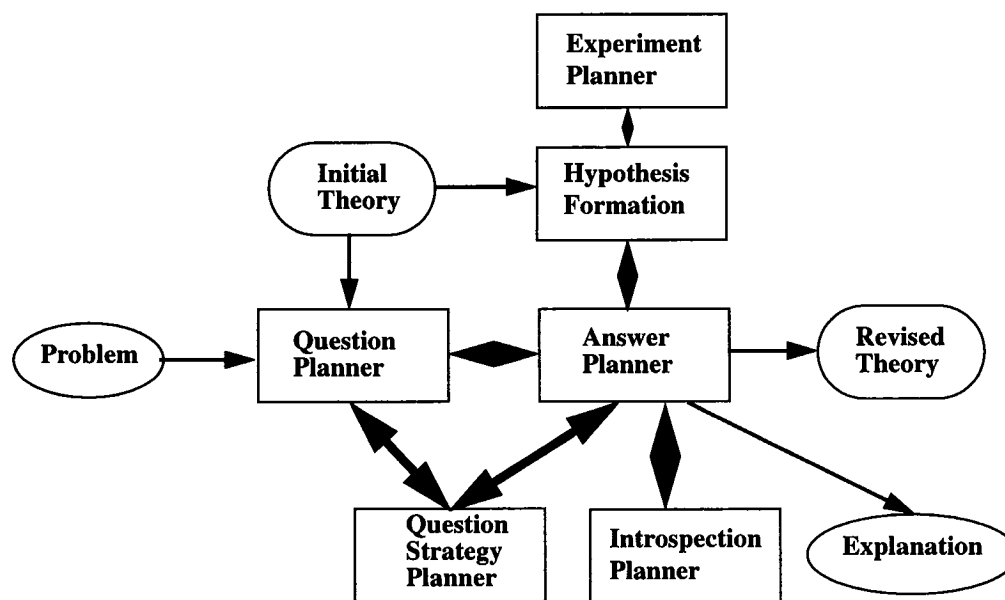


Figure 2: The IULIAN System

were known, IULIAN would have generated the correct hypothesis and the expectation failure would not have occurred. If a question cannot be answered, the Question Planner and the Experiment Planner can again be used to generate additional questions and experiments in an attempt to provide the missing knowledge.

During this process of question-based reasoning and experimentation-based activity, questions focus on objects of the domain to be investigated such as *lamp* and *battery*. 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.

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 an electric circuit with battery, lamp, and switch) and the result of an experiment such as the statement that "the lamp is on when the battery is switched on." Experiments are represented by objects and relations between objects. Objects are represented as Memory Units (MU) which contain an object frame and a content frame. 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. In addition to experiments,

causal models are used to explain experimental results. Causal models have a similar representation to experiments. However, they are stored in a separate library and their objects are viewed as abstract concepts, e.g. the concept "lamp" rather than an actual lamp used in a given experiment (Oehlmann, 1992). In addition, causal models use particular relations between concepts to represent causal links. A causal model is linked to the experiment used to generate that model. Question plans are used to apply case-based planning techniques to the generation of single questions. For example, the question "What is the state of the LAMP?" can be built by combining the substructures "What", "is" "the state of", and "the OBJECT1". OBJECT1 is a variable which is instantiated with the string "LAMP" during plan execution. A question plan has two main parts: the set of descriptors used for indexing the plan and a sequence of steps.

The plan is retrieved by matching its index with the current situation; this is characterised by the goals the system 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 execution. 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 questions 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 is the state of the LAMP?" may be generated by executing the following steps: the first step retrieves the object LAMP and identifies an object which shares the relation HAS-STATE with the object LAMP. This object is ON and is stored in the binding list of the answer plan. The following steps instantiate two variables with the objects LAMP and ON and combine these variable values with the substrings "The", "has", and "the state". The resulting sentence is "The LAMP has the state ON."

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.

During the generation of questions and answers, the system notes the names of the question and answer plans used. These names are stored in a reasoning strategy plan. Reasoning strategy plans have a similar structure to questions and answer plans. However, their actions contain simple calls to the Question and Answer Planners; i.e. after executing a sequence of questions and answers for the first time, a reasoning strategy plan is built which can be retrieved in future situations and executed directly.

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

## The Introspection Process

In this section, we describe our approach for regulating the cognition process based on introspection planning by extending the example described in Section 2.

The system supports metacognitive processes by generating questions about the categories introduced in the last section: goals, strategies, failures, resources, and conditions. The system begins by attempting to explain the expectation failure described in the example. In this situation, the two strategies *experiment-perturbation* and *analogical-mapping* could be performed. We now assume that in the source of the analogical-mapping strategy, the system cannot retrieve an experiment which is sufficiently analogous to the initial experiment. Introspection planning can assess this situation and identify an appropriate strategy to generate an explanation of the initial problem. In addition, it has to monitor the execution of the questions which realise the strategy to ensure goal satisfaction. The stages of the process are summarised in Figure 3.

The identification of candidate questions initiating the *experiment-perturbation* strategy or the *analogical-mapping* strategy requires the system to focus on the reasoning goals to be satisfied by executing a given sequence of questions. Therefore the system generates the following question:

**Question 1:** What are the reasoning goals I attempt to achieve?

In this situation only one reasoning goal is found which leads to Answer 1:

**Answer 1:** The reasoning goal is EXPLAIN-EXPECTATION-FAILURE.

The reasoning goal enables the system to identify the candidate questions.

**Question 2:** What are the question plans which are expected to satisfy the reasoning goals?

**Answer 2:** The question plans are EXPERIMENT-PERTURBATION and ANALOGICAL-MAPPING.

The identified question plans now have to be evaluated with respect to the resources needed, failures which occurred during previous executions, and initial conditions which have to be satisfied before a question plan can be executed. Question 3 focuses on the resources.

**Question 3:** What are the resources needed to accomplish these question plans?

The answer to this question indicates that the analogical mapping strategy requires an analogous previous experiment, whereas the perturbation strategy requires appropriate modification-structures. In these structures, the necessary knowledge is stored to modify a given experimentation plan

**Answer 3:** The resources needed for the question plan ANALOGICAL-MAPPING are ANALOGOUS-EXPERIMENT and the resources needed

**Metacognitive Selection Process:**

1. Identify reasoning goals
2. Use the reasoning goals to identify reasoning strategies.
3. Identify resources. Eliminate those strategies which cannot be executed, because the necessary resources are not available.
4. Identify previous failures. Eliminate those strategies which are likely to fail because they failed before in a similar situation.
5. Identify initial conditions. Eliminate those strategies which cannot be executed under the given conditions.
6. If more than one strategy can be applied, prefer reasoning strategies over experimentation strategies.

**Metacognitive Monitoring Process:**

1. Evaluate modification result. Determine differences between the initial experiment and the experiment which results from the modification.
2. Evaluate understanding. The system attempts to generate an explanation for the new experimental result.

Figure 3: Meta-cognitive processes

for the question plan EXPERIMENT-PERTURBATION are MODIFICATION-STRUCTURES.

After the relation between question plans and resources is established, the system has to evaluate the question plans with respect to the available resources. In particular, it has to eliminate those question plans which require unavailable resources.

**Question 4:** What are the question plans for which the necessary resources are available?

The answer to this question needs only consider the questions plans mentioned in Answer 3, because Question 4 is asked in the context of the previous questions and answers. The answer to Question 4 prefers the plan *experiment-perturbation* because the modification-structures for this plan

are available. In contrast, the analogous experiment needed for the question plan *analogical-mapping* is not available.

**Answer 4:** All the necessary MODIFICATION-STRUCTURES are available for the question plan EXPERIMENT-PERTURBATION.

Although the candidate question plans have been reduced to a single plan, even this strategy might not be suitable. Therefore, the next question focuses on failures which occurred during previous executions of the question plan.

**Question 5:** What are the failures which occurred during previous applications of the question plan EXPERIMENT-PERTURBATION?

**Answer 5:** No failure occurred during a previous application of the question plan EXPERIMENT-PERTURBATION.

Finally, the system checks if the current situation matches the initial conditions necessary to execute the question plan. This evaluation step has two parts: identifying the initial conditions and matching the initial conditions with the current situation. Question 6 addresses the first part:

**Question 6:** What are the initial conditions necessary to execute the question plan?

The answer to this question states that an initial experiment and a previously retrieved experiment are necessary to perform the question plan. In addition, the strategy needs to access the plans which were used to generate these experiments.

**Answer 6:** The initial conditions for the question plan EXPERIMENT-PERTURBATION are INITIAL-EXPERIMENT, INITIAL-PLAN, RETRIEVED-EXPERIMENT, and RETRIEVED-PLAN.

The identified conditions are then used to determine whether the conditions are satisfied by the current situation.

**Question 7:** Are these conditions satisfied for the question plan EXPERIMENT-PERTURBATION?

**Answer 7:** Yes.

This answer concludes the process of selecting between the two candidate question plans and enables the system to execute the question plan *experiment-perturbation*. If at this stage the system were still considering more than one plan, it would attempt to perform a reasoning strategy before an experimentation strategy<sup>1</sup>. However, in such a situation the system would have a high confidence in all the selected strategies, because it has considered the goals, resources, and failures involved.

During its execution, the reasoning strategy has to be

1. This strategy is consistent with psychological results which have suggested that subjects perform experiments if the generation of hypotheses is not successful (Klahr & Dunbar, 1988).

monitored. This process is again supported by the interaction between self-questions and introspection plans. First the result of the experiment modification is compared with the initial experiment.

**Question 8:** Is the perturbed experimentation plan sufficiently close to the initial experimentation plan?

**Answer 8:** No, there are components in the target experiment which do not have equivalents in the source experiment.

The answer indicates that an additional experiment modification has to be performed. However, the system has first to ensure that it “understands” the perturbed experiment. The understanding criterion is given by the ability to explain the experimental result in the context of the experimental setting (see Schank, 1986).

**Question 9:** Do I understand this behaviour?

**Answer 9:** Yes the LAMP is ON, because there is a wire between the BATTERY and the LAMP and there is a CURRENT-FLOW through the LAMP.

The sub-processes of generating experiment modifications and evaluating them are iterated until the resulting experiment is identical with the initial experiment. The final explanation is used to explain the initial expectation failure (see Section 2).

## Introspection Plans

An introspection plan has two main components: a header and a sequence of steps (Figure 4). Important elements of the header are the slots *name*, *planning-goal*, and *failures*. 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* characterises planning failures which have occurred before. If a plan execution fails, the system attempts to explain and to repair the failure. The repair mechanism uses repair rules similar to those described by Hammond (1989). 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*<sup>2</sup>. 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.

<p><b>Plan</b>  <b>Name:</b> resource-assessmnt@modification-struct  <b>Planning-Goal:</b> assess-modification-structure  <b>Failures:</b> None  <b>Binding-List:</b>              ((initial-plan serial@switch-lamp)              (final-plan parallel@switch-lamp))  <b>Intermediate-Result1:</b>              mod-struct@remove-object  <b>Intermediate-Result2:</b> None  <b>Intermediate-Result3:</b> None  <b>Planning Steps:</b></p> <p><b>1. Planning Step</b>  <b>Step-Name:</b> check-deleted-features  <b>Action:</b> (identify-features-to-be-deleted              initial-plan final-plan)</p> <p><b>2. Planning Step</b>  <b>Step Name:</b> check-added-features  <b>Action:</b> (identify-features-to-be-added              initial-plan final-plan)</p> <p><b>3. Planning Step</b>  <b>Step Name:</b> identify-mod  <b>Action:</b> (identify-modification-structures              initial-plan final-plan)</p> <p><b>4. Planning Step</b>  <b>Step Name:</b> evaluate-mod</p>
---

Figure 4: Introspection Plan

The introspection plan given in Figure 4 is executed as part of the generation of Answer 4 in the previous section. The answer has to assess the resources related to the question plan *experiment-perturbation*. These resources are modification structures which are used to transform an experimentation plan. The first step in the introspection plan identifies steps in the initial experimentation plan which are to be deleted. The second step attempts to identify features which are to be added. The results of the first two steps are used to identify a modification structure which can be used to transform the experimentation plan. The final step attempts to evaluate the modification structure in terms of the modification goal to be achieved.

2. Note that Figure 4 shows a reduced version of the original plan. In particular, the slots *precondition* and *goal* in the planning steps are omitted.

## Discussion

We have identified a need to represent knowledge about metacognitive processes. Our novel approach to metacognitive experience is characterised by the integration of introspective self-questions and introspection plans. We have argued that introspection plans can be used to represent mental models about meta-cognitive processes. Moreover, they can be used to realise these processes.

In the example described in this paper, an introspection plan has been used to decide between the reasoning strategy *analogical-mapping* and the activity strategy *experiment-perturbation*. Other reasoning strategies which the IULIAN system has addressed are: *changing the viewpoint* (Oehlmann et al., 1994), *changing the focus*, and *consequence checking*. An example of an additional activity strategy is enforcing the stability of the environment (Hammond, et. al. 1993; Oehlmann, 1995).

Introspection plans can be re-used because they represent metacognitive experience which can be adapted to new situations. We have shown that this approach enables a system to decide between experiment driven and theory driven discovery strategies based on predictions about these strategies. Moreover, we have demonstrated that introspection plans can be used to monitor the execution of such strategies. We expect that introspection plans could support other strategies such as reasoning about an issue from different perspectives (Oehlmann et al., 1994).

This paper focuses on metacognitive processes; future work has to address the representation of knowledge about metacognitive states and events in an integrated model (see Leake, 1995; Cox, 1995; Cox & Freed, 1994). For example, in the current approach, assessing the availability of experiments which serve as a resource for a given reasoning strategy has been achieved by searching the memory of experiments. The result of the assessment could be represented as meta-knowledge in the form of a case which would make this particular search unnecessary in future.

In addition, we expect that the results described in this paper could be applied to educational technology. Kass (1992) used the idea of "coaxing case-based reasoning" in the design of educational software. Introspection plans could be used in a similar way to motivate a student's metacognitive processes.

## References

- Baker, L. & Brown, A. (1984). Metacognitive Skills of Reading. In D. Pearson (Ed.), *A Handbook of Reading Research*. New York: Longman.
- Blank, S. & Covington, M (1965). Inducing Children to Ask Questions in Solving Problems, *Journal of Educational Research*, 59:21-27.
- 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.
- Brown, A. (1987). Metacognition, Executive Control, Self-Regulation and other more Mysterious Mechanisms. In F. Weinert & R. Kluwe (Eds.), *Metacognition, Motivation, and Human Performance*, (pp. 65-116) Hillsdale, NJ: Lawrence Erlbaum.
- Campione, J. (1987). Metacognitive Components of Instructional Research with Problem Learners. In F. Weinert & R. Kluwe (Eds.), *Metacognition, Motivation, and Human Performance*, (pp. 117-139) Hillsdale, NJ: Lawrence Erlbaum.
- Cox, M. (1995). Representing Mental Events (or the Lack Thereof). Position Paper. In M. Cox & M. Freed (Eds.) *Proceedings of the AAAI Spring Symposium on Representing Mental States and Mechanisms*, in print.
- Cox, M. & Freed, M (1994). Using Knowledge of Cognitive Behavior to Learn from Failure. In *Proceedings of the Seventh International Conference on Systems Research, Informatics and Cybernetics* (pp. 142-147).
- Flavell, J. (1976). Metacognitive Aspects of Problem Solving. In L. Resnick (Ed.), *The Nature of Intelligence*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Flavell, J. & Wellman, H. (1977). Metamemory. In R. Kail & J. Hagen (Eds.), *Perspectives of the Development of Memory and Cognition*. 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.
- Fox, S. & Leake, D. (1995). Modelling Case-Based Planning for Repairing Reasoning Failures. In M. Cox & M. Freed (Eds.), *Proceedings of the AAAI Spring Symposium on Representing Mental States and Mechanisms*, in print.
- Gavelek, J. & Raphael, T. (1985). Metacognition, Instruction, and the Role of Questioning Activities. In D. Forrest-Pressley, G. MacKinnon, T. Waller (Eds.), *Metacognition, Cognition, and Human Performance*. New York: Academic Press.
- Hammond, K. (1989). *Case-Based Planning: Viewing Planning as a Memory Task*. New York: Academic Press.
- Hammond, K., Converse, T., & Marks, M. (1993). Towards a Theory of Agency. In S. Minton (Ed.), *Machine Learning Methods for Planning*, (pp. 351-396).
- Kass, A. (1992). Question Asking, Artificial Intelligence and Human Creativity. In T. Lauer, E. Peacock, A. Graesser (Eds.), *Question and Information Systems*, (pp. 303 - 360). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kennedy, A. (1995). Using a Domain Independent Introspection Mechanism to Improve Memory Search. In M. Cox & M. Freed (Eds.), *Proceedings of the AAAI Spring Symposium on Representing Mental States and Mechanisms*, in print.



- Klahr, D. & Dunbar, K. (1988). Dual Space Search during Scientific Reasoning. *Cognitive Science*, 12, 1-48.
- Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.
- Leake, D (1995). Representing Self-Knowledge for Introspection about Memory Search. Position Paper. In M. Cox & M. Freed (Eds.), *AAAI Spring Symposium on Representing Mental States and Mechanisms*.
- Meacham, J. (1972) The Development of Memory Abilities in the Individual and in Society. *Human Development*, 15:205-228.
- 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. (submitted). *Exploratory Discovery: A Case-Based Planning Approach to Learning from Self-Questioning and Experimentation*. Ph.D. Thesis, University of Aberdeen.
- Oehlmann, R., Edwards, P., & Sleeman, D. (1994). Changing the Viewpoint: Re-Indexing by Introspective Questioning. *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, (pp. 675 - 680). Hillsdale, NJ: Lawrence Erlbaum Associates
- Oehlmann, R., Sleeman, D., & Edwards, P. (1993). Learning Plan Transformations from Self-Questions: A Memory-Based Approach. *Proceedings of the 11th National Conference on Artificial Intelligence*, (pp. 520-525). Cambridge, MA: AAAI-Press.
- Patterson, C., Cosgrove, J., & O'Brien, R. (1980). Nonverbal Indicators of Comprehension and Noncomprehension in Children. *Developmental Psychology*, 16:38-48.
- Ram, A. & Cox, M. (1994). Introspective Reasoning Using Meta-Explanations for Multistrategy Learning. In: Michalski, R., & Tecuci, G. (Eds.) *Machine Learning: A Multistrategy Approach*, 4, (pp. 349 - 377). San Mateo, CA: Morgan Kaufmann.
- Schank, R. (1986). *Explanation Patterns. Understanding Mechanically and Creatively*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schank, R. (1982). *Dynamic Memory. A Theory of Reminding and Learning in Computers and People*. Cambridge, UK: Cambridge University Press.
- Schewel, R. & Waddell, J. (1986). Metacognitive Skills: Practical Strategies. *Academic Therapy*, 22(1):19-25.
- Wong, B. & Jones, W. (1982). Increasing Metacomprehension in Learning Disabled and Normally Achieving Students Through Self-Questioning Training. *Learning Disability Quarterly* 5, (pp. 228 - 240).
- Wong, B. (1985). Self-Questioning Instructional Research. A Review. *Review of Educational Research*, 55, 227-268.