Combining Case-based Planning and Introspective Reasoning*

Susan Fox **David B. Leake**
Computer Science Department
Indiana University

Abstract

There is much current interest in introspective reasoning, reasoning about reasoning processes themselves. One application of introspective reasoning is to detect flaws in a system's own reasoning, and to refine its reasoning methods to correct those flaws. We propose a framework for performing such introspective refinement, and describe its implementation in a system which combines introspective learning with case-based planning. We describe empirical tests performed to evaluate the effect of introspective reasoning for this system.

1 Introduction

AI systems which can learn new knowledge about the domain in which they operate can learn to respond to novel situations. They are often restricted, however, by the fixed reasoning methods they must use to manipulate that knowledge. The ability to reason about and refine those reasoning methods in response to new situations should improve the system's ability to adapt to unanticipated and changing circumstances.

Recent research on introspective reasoning has addressed the problem of representing information about reasoning methods, and applying that knowledge to the task of detecting and repairing flaws in the system's reasoning (Freed & Collins, 1994; Stroulia & Goel, 1994; Oehlmann et al., 1994; Ram & Cox, 1994). We have developed a framework for representing reasoning methods and have implemented this approach in ROBBIE, combining an introspective reasoner with an underlying case-based planner.

The ROBBIE[1] system includes learning of domain knowledge through the addition of cases and learning which refines the methods used to retrieve stored cases. ROBBIE has two main parts, as shown in Figure 1: a planner which develops plans through case-based reasoning (Hammond, 1989; Kolodner, 1993) and applies them through simulated execution,

*This work was supported in part by the National Science Foundation under Grant No. IRI-9409348. This paper appears in *Proceedings of the Sixth Midwest Artificial Intelligence and Cognitve Science Society Conference,* Carbondale, IL, pp. 32-36.

[1]Re-Organization of Behavior By Introspective Evaluation

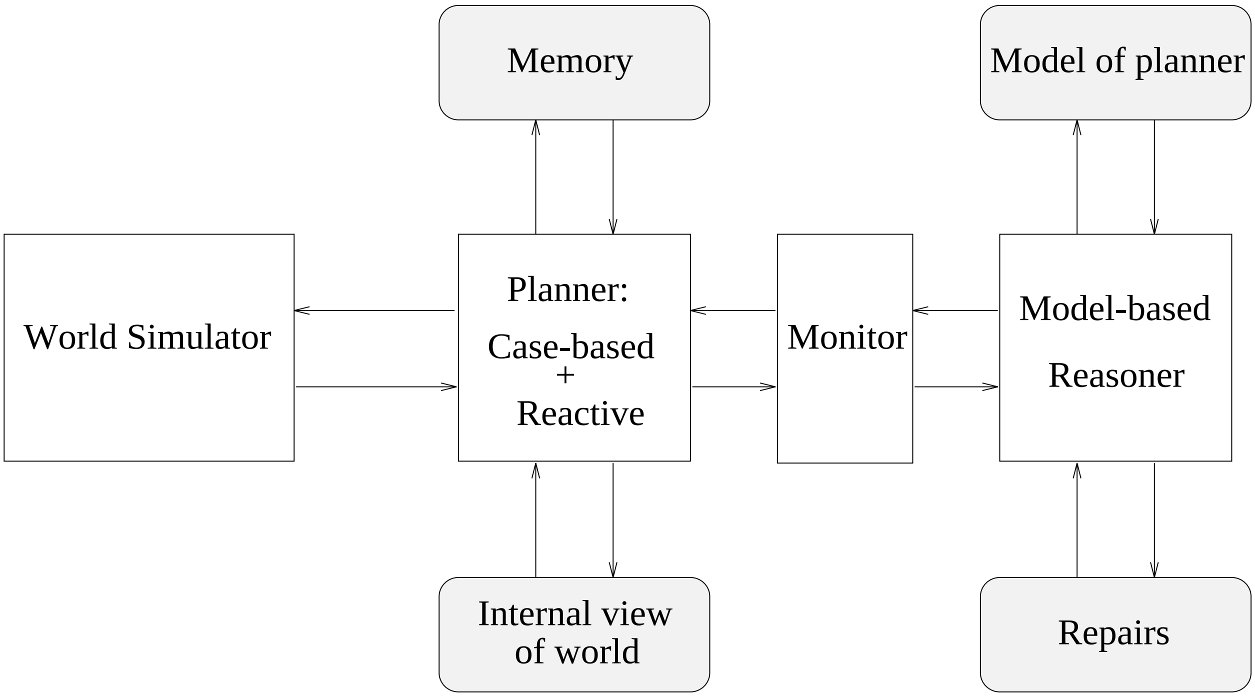

Figure 1: ROBBIE: Combining CBR with Introspective Reasoning

and an introspective model-based reasoner. The introspective reasoner contains a declarative model of the ideal reasoning behavior of the system. Assertions in the model about states of the case-based reasoning process provide expectations for reasoning performance: expectations which are compared to the actual reasoning performance to detect discrepancies. These expectation failures provide an opportunity to learn: to improve the reasoning method by explaining the failure and suggesting a repair to the reasoning method. ROBBIE currently learns new indexing criteria to guide retrieval of cases from memory.

We have performed a set of experiments examining ROBBIE's performance on a wide range of situations, finding that introspective reasoning combined with case learning outperforms case learning alone. In this paper we will present the ROBBIE system and describe its mechanisms for case-based and introspective learning, providing an example of its learning to illustrate the process. We will briefly discuss the empirical tests of ROBBIE's performance, and describe what we conclude from this research.

2 The Performance Task: Pedestrian Navigation

ROBBIE's performance task is navigation around a simulated set of city streets as a pedestrian.ROBBIE is given a starting location and a goal location and must create a plan for getting from one to the other. Because ROBBIE's world model is not guaranteed to be perfect, generation of a plan through case-based reasoning does not assure successful execution; after plans are generated ROBBIE executes them with a simulated robot using reactive planning (Firby, 1989). Execution evaluates the quality of the CBR-created plan, and permits ROBBIE to arrive at a solution despite flaws in the plan. In addition, the reactive planner handles facets of the simulated world that cannot be predicted by the CBR system, such as traffic lights and blocked streets.

The plans in memory are indexed by their starting and ending locations and other features ROBBIE learns through introspective reasoning. Cases describe navigational plans in terms of high-level steps such as "Move north on Oak to Birch" or "Turn east on Fir." These steps are then used as sequential goals for reactive execution, which breaks the problem down into individual actions of the simulated robot.

3 The Introspective Reasoning Task

ROBBIE's high-level task is to improve its reasoning process when reasoning failures are detected. ROBBIE's approach uses a declarative model describing the underlying system's reasoning processes (Fox & Leake, 1994; Birnbaum et al., 1991; Freed & Collins, 1994). The model contains assertions about the system's ideal reasoning behavior; the actual reasoning is compared to this ideal to detect discrepancies which are viewed as reasoning failures. Once a reasoning failure is detected the model is used to create an explanation of the failure and to suggest a repair (See Fox & Leake (1994) for a discussion of ROBBIE's failure detection). Currently, ROBBIE can re-index its memory to include new features.

The model's assertions describe states that would be ideally true of particular portions of the case-based reasoning process. Assertions range from details of ROBBIE's implementation to abstract descriptions of the CBR process. The assertions are organized by the component of the process to which an assertion refers, the level of specificity of the assertion, and by connections to related assertions. In Fox & Leake (1995b) we provide a detailed description of the model and the representation it uses.

ROBBIE can currently refine its reasoning methods by altering the retrieval criteria to include new features as indices. Such learning may not be triggered until after an imperfect retrieval as been adapted and executed. Before storing a new solution, ROBBIE retrieves the case in memory with the closest *solution*, according to the similarity of their plan steps. If a different case was retrieved originally, then the indices used in retrieval failed to select the best match. In repairing the failure ROBBIE must determine what, if any, relevant features were overlooked. ROBBIE currently considers a pre-defined set of feature types which are ranked by a fixed hierarchy of importance. ROBBIE looks for features which are common to the new solution and its closest match, and alters its indexing method to include those new features.

4 An Example

In this section we will present an implemented example of ROBBIE's introspective reasoning, shown in Figure 2 In this example, ROBBIE has in memory plan A, which describes how to get from L1 to L2, and plan B, which describes how to get from L2 to L3. The current task is to get from location L4 to location L2. Judging similarity by the geographic closeness of starting and ending locations, plan A appears to be the closest because it shares the same ending location (ROBBIE's retrieval criteria does not include reversals of known routes, so

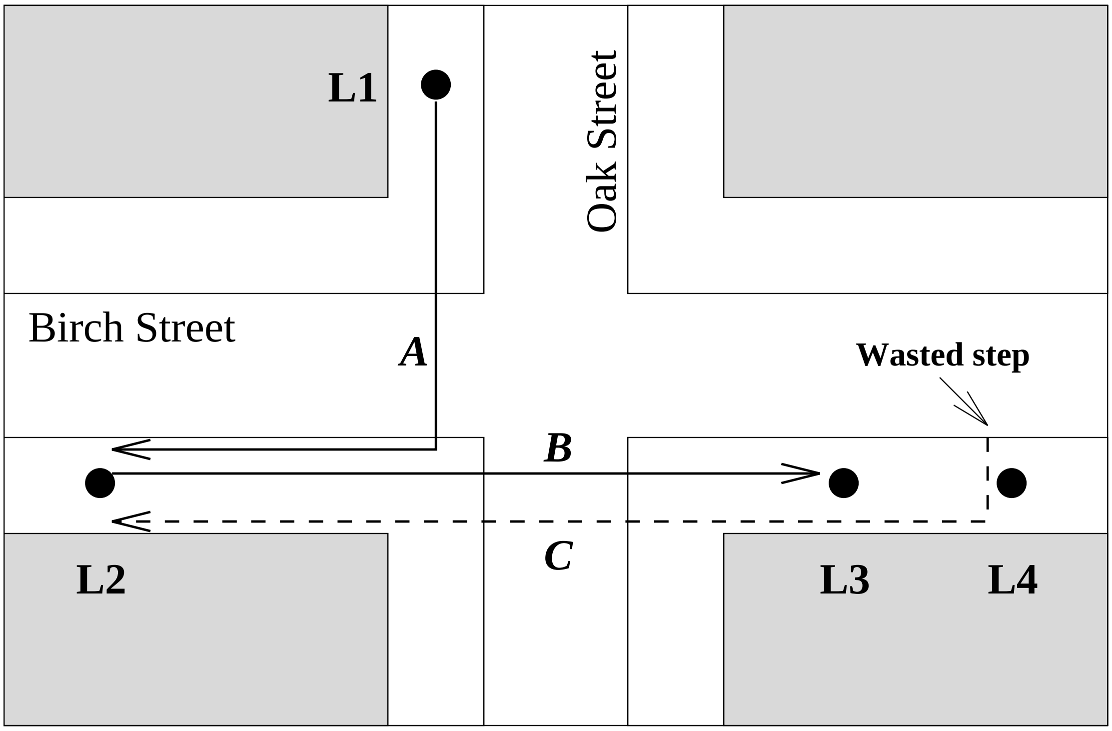

Figure 2: Example of ROBBIE's Introspective Learning

plan B does not look similar at all). Plan A is selected and adapted to create plan C (dashed line in Figure 2). During this process, the introspective reasoner monitors the system's behavior but detects nothing wrong. When the plan is executed, however, the wasted plan steps will be eliminated, producing a plan which moves straight west. When the resulting plan is stored in memory, a reasoning failure is detected: the final solution's steps are most similar to plan B and not the retrieved plan.

In explaining this failure, ROBBIE will reconsider relevant assertions connected to the failed assertion by links of the model. One such assertion is that the retrieval index will include all the features needed to retrieve the closest case. In re-considering this assertion the system discovers that the detected failure can be accounted for by a failure to consider a feature of the two cases: that each involves moving straight along an east-west street. This assertion suggests a repair: add "moves straight east-west" to the features used in indexing cases, and re-index the cases in memory to include the new feature.

5 Evaluating ROBBIE

We have run extensive tests on ROBBIE to evaluate its performance with and without introspective re-indexing of memory. ROBBIE's performance was measured over many goals presented in sequence. To avoid biases due to order of presentation, twenty six sequences were created ranging from well-ordered to random arrangements of goals.[2] For all but one

[2]ROBBIE was presented with a total of 41,600 goals: 26 sequences, each 40 goals long; and 20 runs with introspection, 20 without.

sequence, ROBBIE succeeded more often with introspective learning than with case learning alone. At the same time, the retrieval process became more efficient through fine-tuning of retrieval indices. The one anomalous case can be explained as a combination of poor case-based performance with over-restriction of the retrieval process by the addition of new indices. (See Fox & Leake (1995a) for an in-depth discussion of these results.)

6 Conclusions

Using introspective reasoning about a system's reasoning methods in order to detect and repair flaws in those methods has been proposed as a means of permitting a system to adapt better to new circumstances and to refine its reasoning to match the actual situations it faces. Combining introspective reasoning with an underlying learning system presents the prospect of learning both about the domain and about the reasoning task itself.

We propose a framework which combines a case-based learning system with an introspective reasoning component that allows refinement of the system's own reasoning process. The introspective reasoner uses a model of expected ideal reasoning performance to detect discrepancies in actual performance, and to diagnose and repair such expectation failures. ROBBIE implements this framework for a planner that learns new plans as it creates them and introspectively learns which features are required for retrieving relevant cases from memory.

Through empirical tests we have determined that the addition of introspective reasoning improves performance for ROBBIE both in terms of its rate of success and in terms of the efficiency of its retrieval process. In future research, ROBBIE will be extended to include more introspective repairs, and further experiments should be performed with a larger initial case base and more complex maps.

References

Birnbaum, L., Collins, G., Brand, M., Freed, M., Krulwich, B., & Pryor, L. (1991). A model-based approach to the construction of adaptive case-based planning systems. In Bareiss, R. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 215–224 San Mateo. DARPA, Morgan Kaufmann, Inc.

Firby, R. J. (1989). *Adaptive Execution in Complex Dynamic Worlds.* Ph.D. thesis, Yale University, Computer Science Department. Technical Report 672.

Fox, S. & Leake, D. (1994). Using introspective reasoning to guide index refinement in case-based reasoning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 324–329 Atlanta, GA. Lawrence Erlbaum Associates.

Fox, S. & Leake, D. (1995a). An introspective reasoning method for index refinement. In *Proceedings of 14th international Joint Conference on Artificial Intelligence*. IJCAI.

Fox, S. & Leake, D. (1995b). Modeling case-based planning for repairing reasoning failures. In *Proceedings of the 1995 AAAI Spring Symposium on Representing Mental States and Mechanisms* Stanford, CA. AAAI. (ftp.cs.indiana.edu:/pub/leake/p-95-02.ps.Z).

Freed, M. & Collins, G. (1994). Adapting routines to improve task coordination. In *Proceedings of the 1994 Conference on AI Planning Systems*, pp. 255–259.

Hammond, C. (1989). *Case-Based Planning: Viewing Planning as a Memory Task*. Academic Press, San Diego.

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

Oehlmann, R., Edwards, P., & Sleeman, D. (1994). Changing the viewpoint: re-indexing by introspective questioning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 675–680. Lawrence Erlbaum Associates.

Ram, A. & Cox, M. (1994). Introspective reasoning using meta-explanations for multistrategy learning. In Michalski, R. & Tecuci, G. (Eds.), *Machine Learning: A multistrategy approach Vol. IV*, pp. 349–377. Morgan Kaufmann.

Stroulia, E. & Goel, A. (1994). Task structures: what to learn?. In desJardins, M. & Ram, A. (Eds.), *Proceedings of the 1994 AAAI Spring Symposium on Goal-driven Learning*, pp. 112–121. AAAI Press.