Decision-Theoretic Plan Failure Debugging and Repair

Lisa J. Burnell[1]

Department of Computer Science Engineering
The University of Texas at Arlington, Box 19015, Arlington, Texas 76019
burnell@cse.uta.edu

A number of strategies exist for the recovery from execution-time plan failures. One manner in which these strategies differ is the degree of dependence on the reliability and availability of the planner's knowledge. The best strategy, however, may be dependent on a number of considerations, including the type of plan failure, the criticality of the failure, the availability of resources, and the reliability and availability of the knowledge involved in a given plan failure instance. We are examining a decision-theoretic approach to diagnose plan failures and to dynamically select from multiple failure recovery strategies when an execution-time plan failure occurs.

Existing failure recovery strategies generally classify, with assumed or proven certainty, the type of error that occurred during plan execution and then select a fixed strategy to recover from that error. Assumptions regarding the accuracy and completeness of the planner's domain model vary. On one end are approaches that use deterministic heuristics or purely syntactic analyses to debug and repair. These approaches are efficient and require limited knowledge, but generally are limited in the level of diagnosis and repair they can perform. On the other end are logic-based approaches, that are robust, but are knowledge and resource intensive. Such approaches are not always feasible.

Incomplete or uncertain knowledge of previous planning actions, as in multi-agent planners, precludes a complete logical analysis. Even when possible, the costs of collecting and reasoning with complete information may be intractable, especially given time pressures and other resource constraints. The goal of our work is the development of an approach that can intelligently select and apply failure recovery strategies that are appropriate to the situation and that can cope with uncertainty.

There are three primary components of our research: diagnosis of plan failures, plan repair and planner modification. When a failure is detected, we use a probabilistic method for determining the error class and the source(s) of the error. This method has already proven effective in debugging programs (Burnell & Horvitz 1993), and is being adapted for debugging generated plans. Plan repair strategies are selected using a decision theoretic approach similar to (Howe & Cohen 1991), with the added feature of dealing with potential uncertainties in the error classification and resource constraints. Finally, machine learning methods are employed, when warranted, to correct and refine the planner itself.

Our approach uses probabilistic models, represented as belief networks, to construct an ordering over classes of errors, to identify the likely source(s) of the error and to recommend an appropriate repair strategy. The belief networks model the uncertain relationships about the nature and structure of planning actions and the likelihood of types of errors. Value of information calculations recommend which computationally complex logical analyses are worth undertaking collect evidence. Also modeled is the decision problem of selecting a preferred repair strategy, which may include planner modification, based on the likely error class, failure criticality and resource availability, as in (Horvitz 1988). For example, in a multi-agent planner, repair strategies may include local selection of a reactive failure-recovery action or requesting replanning from a more sophisticated planner.

[1] This work is supported in part by the Advanced Research Projects Agency, State of Texas Advanced Technology Projects and the National Science Foundation.

References

Burnell, L. J. and Horvitz, E. J. 1993. A Synthesis of Logical and Probabilistic Reasoning for Program Understanding and Debugging. Proc. of the 9th Conf. on Uncertainty in AI, 285-291. San Mateo, CA: Morgan Kaufmann.

Horvitz, E.J. 1988. Reasoning under varying and uncertain resource constraints. Proc. of the 7th National Conf. on Artificial Intelligence, 111-116. San Mateo, CA:Morgan Kaufmann.

Howe, A. and Cohen, P. 1991. Failure Recovery: A model and Experiments. Proc. of the 9th National Conf. on Artificial Intelligence, 801-808. Menlo Park, CA: AAAI.