# Combining Strategies for Planning in Huge Uncertain Domains

Kimberly G. Waters

Artificial Intelligence Laboratory
Massachusetts Institue Of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

**The Problem:** We would like to create a robust system for planning in huge uncertain domains. Markov Decision Processes (MDPs) are a standard way of devising detailed plans in stochastic domains. They can provide a complete policy so that regardless of what occurs, you know how best to react to the situation. However, in large domains it becomes computationally intractable to compute this full policy in advance. We are interested in finding methods of creating full-bodied plans which do not require the consideration of every state in the domain as a possible successor to the current state nor the explicit definition of tables of rewards and transition probabilities.

**Motivation:** Our ultimate goal is to be able to produce robots that are flexible enough to operate in very complex environments for long periods of time. They would have to handle many different, somtimes conflicting, goals. Currently no technique is able to address a problem of this magnitude.

**Previous Work:** Research has been conducted which explores the use of different techniches to simplify the complexity of planning. These methods construct a simpler approximation of the problem making the policy building process easier. Our research is built upon the ideas of limiting the number of outcomes considered and limiting the resolution at which those states are described[1, 2, 3].

**Approach:** First we implemented the Value Iteration algorithm for planning. This algorithm computes the optimal solution and can be used to compare our results against for relatively small domains. Though it does give an optimal policy, the trade off is that for huge domains it is impossible to compute.

Our approach is to combine the planning strategies of three currently available algorithms. The first is a Sparse Sampling algorithm[2]. It works by using a generative model (simulator) to sample possible next states based on the probability they will occur. It is efficient in that it only considers the most probable next states and thus has no dependancy on the total number of states in the MDP.

The next method that we consider is the use of factored representations[3]. Factored representations are beneficial in our situation because they limit the resolution at which an MDP's states are described. This is analogous to combining similar states into one, further reducing the number of states which have to be considered.

Finally, we integrate a restricted planning domain[1]. This helps by limiting the attention of our algorithm to an envelope of states surrounding a nominal path from the start to the goal state. The envelope is simply a subset of the total state space.

**Impact:** Upon completion, this work will provide an algorithm for more efficient computation of policies in large stochastic domains.

**Future Work:** The next step in this project is to complete the implementation of the sparse sampling algorithm. After this has been done we will:

- incorporate factored representations into the current version of the algorithm,
- limit the number of states considered as possible next states for the augmented algorithm,
- analyze the effiency and correctness of this planning strategy compared to the Value Iteration method as well as the individual strategies upon which it was based.

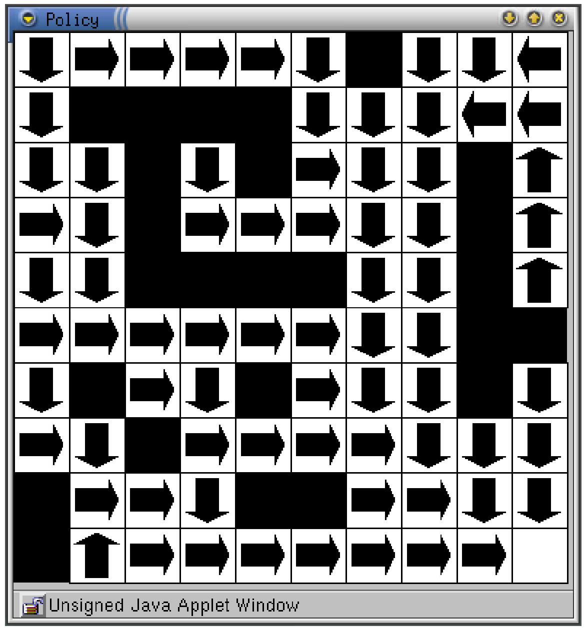


Figure 1: The optimal policy for a very simple maze domain with walls, pits, a starting state in the upper left corner, and a goal state in the lower right corner.

**Research Support:** This work was supported by a Walter A. Rosenblith fellowship.

**References:**

[1] T. Dean, L. Kaelbling, J. Kirman, and A. Nicholson. Planning Under Time Constraints in Stochastic Domains. *Artificial Intelligence*, 76, 1995.

[2] M. Kearns, Y. Mansour, and A. Ng. A Sparse Sampling Algorithm for Near-Optimal Planning in Large Markov Decision Processes. *Proceedings of the Sixteenth International Conference on Artificial Intelligence (IJCAI-99)*, Morgan Kaufmann, Los Altos/Palo Alto/San Francisco, 1324-1331, 1999.

[3] D. McAllester and S. Singh. Approxiamte Planning for Factored POMDPs using Belief State Simplification. *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI-99)*, 409-416, Morgan Kaufmann Publishers, 1999.