

A Decision-Theoretic Academic Advisor: Extended Abstract

Joshua T. Guerin

Department of Computer Science
University of Kentucky
jtguer2@uky.edu

1 Introduction

This paper describes work related to stochastic modeling and decision-theoretic (DT) planning methods applicable to the real-world domain of academic advising.

The uncertainty associated with stochastic planning is often approached as a problem of finding a complete mapping from states to actions (*a policy*). A need to plan for all contingent futures makes (DT) planning an inherently difficult problem in large, real-world domains. This is different from classical planning where the effects of actions are known with certainty, and a step-by-step *plan* is sufficient for reaching a goal state.

Recently, replanning and heuristic planning techniques which operate over *determinized* (deterministic) versions of stochastic planning problems have seen much success in competitions and in planning literature. This approach typically involves direct use of deterministic solutions in place of a complete policy or as a heuristic for finding a partial or complete policy. These successes have caused some to call into question the types of problems which are suited to DT planning techniques and to examine the role of approximate planning in stochastic domains [Little and Thiébaux, 2007; Sanner, 2008].

Deterministic planning techniques have seen particular success in domains which possess inherently deterministic qualities. Deterministic methods also provide a means for planning in domains which are too large (in terms of size and branching factor) for fast stochastic planning methods. However, current determinization approaches fail on domains with properties such as avoidable dead-end states and multiple goal-reaching paths. Can the success of deterministic planning techniques be extended into domains which exhibit both categories of properties? The domain of academic advising exhibits properties which confound both stochastic planners (size, branching factor) and deterministic planners (avoidable dead-end states, multiple goal-reaching trajectories).

Applying stochastic modeling and DT planning to academic advising domains may also provide real-world benefit to students and advisors. The effect and importance of advising services in post-secondary education has been a major topic of discussion in education literature for several decades. Available, quality advising services have been linked to a number of facets of overall student success, satisfaction, and perceived quality and utility of education [King, 1993;

Metzner, 1989; Titley and Titley, 1982; Lowe and Toney, 2001]. Despite the importance of academic advising to student success, academic advising services are often found to be inconsistent, and are often not treated as a high priority within the educational community [Titley and Titley, 1982; Lowe and Toney, 2001].

DT planning in this research is in the context of a Markov decision process (MDP) planning model. MDPs are a mathematical formalism used to model a controlled stochastic processes. An MDP is represented by a 4-tuple: $\langle \mathcal{S}, \mathcal{A}, T(s, a, s'), R(s) \rangle$ where $\mathcal{S}$ and $\mathcal{A}$ are a (finite) set of states and actions, $T(s, a, s') \rightarrow [0, 1]$ is a transition function denoting the probability of arriving at state s' after taking action a from state s , and a real-valued utility function $R(s) \rightarrow \mathbb{R}$ denoting the a reward value for occupying state s .

2 Modeling

In our current work, we learn a stochastic predictive model for student grades from a large set of student transcripts. Students' reward functions are personalized, and we use interviews with students and advisors to determine a set of utility/reward variables. We mine transcripts, teaching evaluations, and other sources in order to determine possible assignments to state and reward variables.

We represent transition functions as *dynamic Bayesian network* (DBN) models. A DBN is a directed acyclic graph, where edges represent temporal dependencies between variables and nodes are domain variables. Nodes in a DBN have associated *conditional probability tables* which give probabilities over possible assignments given the assignment of *parent variables* (variables with outgoing edges).

Actual transcript data is very sparse, even with a large set of transcripts. In order to construct DBN models we must generate probability distributions over outcomes for all possible assignments to parent variables. A DBN for a single course with 5 parents, each of which has 6 possible values (A-D, failure, and not.taken) this would require $6^5 = 7,776$ probability distributions. Popular or highly required courses may be taken by hundreds or even thousands of students within the span of several years, but even this is insufficient to derive realistic probability distributions through direct analysis.

For model generation we have introduced a method of generating probability distributions, based on the predictive power of item-based collaborative filtering [Guerin *et al.*,

2010]. This algorithm uses a generic collaborative filtering framework which treats grade prediction as a problem of predicting ratings over items given a user's rating over similar items, and the ratings trends of past users. Network structure (parent nodes for each variable) is learned by computing similarity between rated items (courses). Probability distributions are generated by determining the deviation of actual ratings from predicted ratings over sets of courses from each course level. Deviations are recorded in an $n \times n$ matrix where n is the number of possible assignments to variables; rows of the matrix consist of values predicted by collaborative filtering and columns of the matrix consist of actual values. Rows from this matrix serve as probability distributions in associated CPTs; row i being the probability distribution whenever the CF algorithm predicts a value of i .

In order to construct a model of utility in the academic domain we are using course evaluations to determine user preferences based on perceived utility. Much work has been done studying the biases in evaluations, but conclusions vary and studies of computer science evaluations are virtually non-existent. Linear models may serve as a way to eliminate these biases, while retaining information about perceived utility. We are examining the use of model residuals as a more neutral metric for approximating student utility [Guerin and Michler, 2011]. The linear model approximates ratings to questions given the values of several potential sources of bias (expected grade, student classification, course requiredness, etc.). Model residuals show how high or low actual ratings are, compared to this expectation, which indicate how ratings in a particular course compare to similar courses.

3 Future Work

Academic advising is an important real-world planning problem that yields significant technical challenges. The next challenge that we will tackle is that of planning in such a large, multi-valued, complex domain. Our intent is to leverage modern fast-planning techniques. Deterministic planners such as RFF and FF-Replan have demonstrated that fast planning in many competition domains can be accomplished through little more than solving a determinized version of a stochastic planning problem, stepping through a resulting deterministic plan, and replanning when necessary [Teichteil-Königsbuch *et al.*, 2010; Yoon *et al.*, 2007]. The two most prominent determinization schemes are all-outcomes (AO), where each outcome of a stochastic action is mapped to a separate deterministic action, and single-outcomes (SO), where (in most cases) the *most likely* outcome of each action is treated as a deterministic outcome.

These determinization schemes offer trade-offs in terms of time, space, and quality. Determinization using SO produces smaller deterministic problems for which a solution may be found quickly, and resulting plans have a high probability of success. However, SO determinization may ignore possible outcomes which are necessary for finding a solution (in such domains AO may be a better choice) [Teichteil-Königsbuch *et al.*, 2010]. AO has also been shown to produce a more informed heuristic than SO, but at the cost of both time and space [Kolobov *et al.*, 2010].

We are investigating two issues related to determinized planning methods. These are: (1) examining a middle ground between AO and SO determinization by examining what happens when *many-outcomes* (MO) are considered during determinization and (2) identifying how close to optimal plans generated by these determinizing planners are.

AO and SO determinization use a threshold of one or all outcomes in generating deterministic planning problems. In implementing many-outcomes determinization we are examining what happens when this threshold is changed depending on the planning problem. Examples of this are an “ n most probable” outcomes determinization or a “probability thresholded outcomes” determinization where some subset of outcomes are mapped to new deterministic actions.

References

- [Guerin and Michler, 2011] J.T. Guerin and D. Michler. Analysis of undergraduate teaching evaluations in computer science. In *SIGCSE 2011*. Sheridan Printing, 2011.
- [Guerin *et al.*, 2010] J.T. Guerin, R. Crawford, and J. Goldsmith. Constructing dynamic bayes nets using recommendation techniques from collaborative filtering. Technical Report TR-515-10, University of Kentucky, December 2010.
- [King, 1993] M.C. King. Academic advising, retention, and transfer. *New Directions for Community Colleges*, 1993(82):21–31, 1993.
- [Kolobov *et al.*, 2010] A. Kolobov, Mausam, and D.S. Weld. Classical planning in MDP heuristics: with a little help from generalization. In *ICAPS 2010*, 2010.
- [Little and Thiébaux, 2007] I. Little and S. Thiébaux. Probabilistic planning vs. replanning. In *In ICAPS Workshop on IPC: Past, Present and Future*, 2007.
- [Lowe and Toney, 2001] A. Lowe and M. Toney. Academic advising: Views of the givers and takers. *Journal of College Student Retention*, 2(2):93–108, 2001.
- [Metzner, 1989] B.S. Metzner. Perceived quality of academic advising: The effect on freshman attrition. *American Educational Research Journal*, 26(3):422–442, 1989.
- [Sanner, 2008] S. Sanner. How to spice up your planning under uncertainty research life. In *Workshop: A Reality Check for Planning and Scheduling Under Uncertainty (ICAPS-08)*, 2008.
- [Teichteil-Königsbuch *et al.*, 2010] F. Teichteil-Königsbuch, U. Kuter, and G. Infantes. Incremental plan aggregation for generating policies in MDPs. In *AAMAS 2010*, pages 1231–1238, 2010.
- [Titley and Titley, 1982] R.W. Titley and B.S. Titley. Academic advising: The neglected dimension in designs for undergraduate education. *Teaching of Psychology*, 9(1):45–49, February 1982.
- [Yoon *et al.*, 2007] S. Yoon, A. Fern, and R. Givan. FF-Replan: A Baseline for Probabilistic Planning. In *ICAPS 2007*, pages 352–359. AAAI, 2007.