

Joshua T. Guerin, Logic and AI Lab – University of Kentucky



You still need to take one course from list “G” and two courses which fulfill general requirement “N”.

No more morning classes... and how about we avoid theory this semester... I’m sick of writing proofs!

### The problem:

- Decision factors: requirements and the student’s goals (graduation, career, grad school, etc.) and preferences.
- Students and advisors may approach the problem from different perspectives.
- Success can only be predicted probabilistically.
- Actions have short and long-term effects—next semester’s schedule is only part of the picture.
- An academic domain is very large (trillions of possible transcripts).

### Decision-theoretic planning uses:

- probabilities of success and failure
  - student preferences and goals
- To optimize expected value to the student over time.

### A Markov Decision Process model of advising:

*States:* transcripts  
*Actions:* take/retake classes  
*Utility:* elicited from student, available data<sup>1</sup>  
*Transitions:* Learned from transcripts<sup>2</sup>


*Policy:* a complete mapping from states to actions.

A policy provides the optimal action regardless of the circumstances.

### Problems:

- Decision-theoretic planning has a high computational cost
- We need a way to convey complex probabilistic information to users.

- Optimal long-term planning can done using MDP model.
- Planning with an MDP yields a *policy*.



A Policy

Project Members: Crawford, Dodson, Goldsmith, Guerin, Mattei, Mazur, Michler

<sup>1</sup> “Analysis of Undergraduate Teaching Evaluations in Computer Science,” Joshua T. Guerin, Daniel Michler, The Proceedings of the ACM Special Interest Group in Computer Science Education (SIGCSE 2011).  
<sup>2</sup> “Tech Report: Constructing Dynamic Bayes Nets Using Recommendation Techniques From Collaborative Filtering,” Joshua T. Guerin, Robert Crawford, and Judy Goldsmith. TR 515-10