

Multi-Model Markov Decision Processes


IIE Transactions, 2021

Lauren N. Steimle, David L. Kaufman, and Brian T. Denton

Markov Decision Processes

Sequential decision-making problems

- ▶ Properties of the problem
 - ▶ Sequential decisions
 - ▶ Uncertain probabilistic outcomes
- ▶ What if we do not know the probabilities?
 - ▶ Multiple models



Motivation

- ▶ Medical setting
- ▶ Multiple disagreeing "experts"
- ▶ Weighted value problem (WVP): Find a policy that maximizes the expected reward given a distribution over the possible models
- ▶ Markovian policies

Contributions

- ▶ Complexity of WVP (not-novel)
- ▶ Type of optimal decisions/policies
- ▶ Exact method to solve WVP
- ▶ Approximate method to solve WVP
 - ▶ Error bound for 2 models

Possible Second Paper Directions

- ▶ Other objectives
 - ▶ Qualitative
 - ▶ Multiple policies
- ▶ Complexity
- ▶ Algorithmic
- ▶ Partial observability