


Planning, Execution & Learning: Planning Under Uncertainty 2

Reid Simmons
Manuela Veloso

Replanning to Handle Uncertainty

- Basic Idea:
 - Plan using a deterministic model
 - Start executing until outcome diverges from expected state
 - Replan



Determinization

- Generate Deterministic Domain from Probabilistic Effects
- Single-Outcome
 - Select one outcome (e.g., most likely; most add effects)
 - Very efficient planning
 - Cannot handle situations where alternate effects are needed
- All-Outcomes
 - Generate one action for *each* combination of effects
 - Exponential blow-up in number of actions
 - May plan to include effects that are unlikely to occur

Determinization in Hindsight

- Dynamic Generation of Deterministic Actions
 - At each time step in the plan, select effects according to probability distribution
- Compute Expectation over Multiple Plans
 - $V^*(s, T) = \max_{\pi} E[R(s, F, \pi)]$
 - $Q^*(s, a, T) = R(S) + E[V^*(s', T-1)]$
 - $V_{hs}(s, T) = E[\max_{\pi} R(s, F, \pi)]$ ← **Overestimate**
 - $Q_{hs}(s, a, T) = R(S) + E[V_{hs}(s', T-1)]$
- Choose Best First Action
 - For completeness, need random tie breaking

FF-Replan


- FF-Replan
 - Originally developed as a “baseline” for planning under uncertainty
 - Highly effective in planning-under-uncertainty competitions
 - Led to development of domains that are “probabilistically interesting”

Preference Models




Decision Theory

- Descriptive Theory: *How to make choices among alternatives*
 - Does not describe how to *generate* the alternatives
- Utility Models
 - Capture preferences for rewards and resource consumption
 - Capture risk attitude (risk-neutral, risk-averse, risk-seeking)
 - *Expected Reward* vs. *Expected Utility*



Plan Dominance

- What?
 - Eliminate plan from consideration if provably worse than another plan under consideration
- How?
 - For **complete** plans: P_1 dominates P_2 if $EU(P_1) > EU(P_2)$
 - For **partial** plans:
 - Maintain range of utilities: $[EU, \bar{EU}]$
 - P_1 dominates P_2 if $EU(P_1) > EU(P_2)$

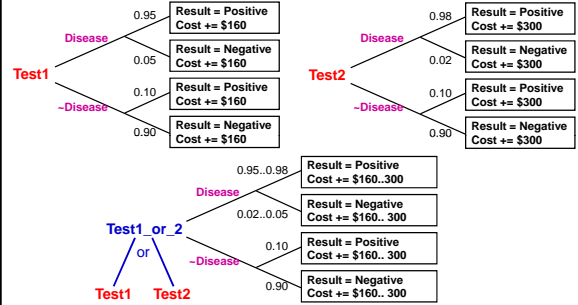


- Which?
 - *Provably optimal* to select plan with greatest upper bound (Goodwin, 1996)

DRIPS (Haddawy & Hanks, 1994)

- Decision-Theoretic Hierarchical Refinement (HTN) Planner
 - Probabilistic; Maximizes expected utility
- Action Representation:
 - *Concrete* Actions
 - Primitives, fully refined
 - *Abstract* Actions
 - *Sequence* of actions (e.g., “load and drive”)
 - *Alternative* actions (e.g., “drive slowly”, “drive fast”)
- Planner
 - Works by refining abstract actions
 - Eliminate dominated plans

DRIPS Example (I)



DRIPS Example (II)

