

Planning, Execution & Learning: Planning with POMDPs (II)

Reid Simmons
Manuela Veloso

Solving POMDPs

- Representational Choices
 - Exact V , exact b
 - Witness, policy-iteration, ...
 - Approximate V , exact b
 - Differentiable function approximators (higher-order polynomials, neural nets, ...)
 - Exact V , Approximate b
 - Grid-based
 - Sample-based (PBVI, HSVI)
- **Greedy Approaches** Based on Solving Underlying MDP

Approximating Value Function

- Use Function Approximator with “Better” Properties than Piece-Wise Linear
 - Continuous (differentiable), non-linear
 - Typically use on the order of one vector per action
- Comparisons
 - + Generally much more efficient
 - May poorly represent optimal solution (however, better function approximation usually implies better results)

SPOVA Algorithm (Parr, 1995)

- Approach
 - Use a small set of vectors to represent the value function
 - Approximate the value function by a smooth (differentiable) function

$$V(b) = \max_{v \in \Psi} (v \bullet b)$$


$$\approx \left\{ \sum_{v \in \Psi} v \bullet b \right\}^{1/k}$$
 - Use gradient descent to adjust components of the vectors

$$E(b) = V(b) - \beta \{ \max_{v_i} \{ R(a, b) + \gamma \sum_{b'} p(b' | a, b) V(b') \} \}$$

$$v_{i,t+1}(s) = v_{i,t}(s) + \alpha E(b) b(s) (v_i \bullet b)^{k-1} / V(b)^{k-1}$$

Approximating Belief Space

- Use Grid-Based Approximation
 - Discretize belief space: Place finite grid over belief simplex
 - Evaluate value function at grid points
 - Interpolate
- Regular Grid (Lovejoy)
 - + Simple method, easy interpolation
 - Exponential space needed
- Non-Regular Grid (Hauskrecht)
 - + More accurate – tries to follow value contours
 - Interpolation is difficult
- Variable-Resolution Grid (Zhou & Hansen)
 - + Fairly accurate – grid points added where distinctions are needed
 - + Interpolation is fairly easy – add virtual grid points




Point-Based Value Iteration (Pineau, 2003)

- **Approximate Belief Space**
 - Deals with only a subset of the belief points
 - Focus on the most relevant beliefs
 - Focus on the points that add the most information
- Tackles the **Curse of History**
 - Number of action/observation histories typically grows exponentially
 - Leads to exponential number of vectors in value function

PBVI Approach

1. Maintain One Alpha-Vector Per Belief
2. Perform Backups Over Only Those Vectors
3. Grow Belief Space Intelligently
 - Create only reachable belief points
 - Choose beliefs so as to create uniform density




Heuristic-Search Value Iteration (Smith, 2004)

- **Approximate Belief Space**
 - Deals with only a subset of the belief points
 - Focus on the most relevant beliefs
 - Focus on the most relevant actions and observations
- **Main Idea**
 - Value iteration is the dynamic programming form of a tree search
 - Go back to the tree and use heuristics to speed things up
 - But still use the special structure of the value function and plane backups


HSVI Approach

- Maintain Constraints on Value of Beliefs
 - Lower and upper bounds
 - Initialize upper bound to QMDP; Lower bound to "always a"
- Explore the "Horizon" Tree
 - Back up values to further constrain belief values
 - Lower bound is regular plane backups
 - Upper bound is set of points
 - solve linear program to interpolate
 - approximate upper bound (sawtooth)





HSVI Approach

- Need to decide:
 - When to terminate search?
 - Minimal gain
 - $width(V(b)) < \epsilon \gamma^t$
 - Which action to choose?
 - Highest upper bound:
 - $argmax_a Q(b, a)$
 - Which observation to choose?
 - Reduce excess uncertainty most
 - $argmax_o p(o | b, a) * (width(V(\tau(b, a, o))) - \epsilon \gamma^{t+1})$



HSVI Results

5x5 map, 7 rocks, 3200 states



Greedy Approaches to POMDP Planning

- Solve Underlying MDP
 - $\pi_{MDP}: S \rightarrow A$; $Q_{MDP}: S \times A \rightarrow \mathcal{R}$
- Choose Action Based on Current Belief State
 - "most likely" (Nourbakhsh, 1994) - $\pi_{MDP}(argmax_s(b(s)))$
 - "voting" (Simmons & Koenig, 1995) - $argmax_a(\sum_{s \in S} b(s) \delta(a, \pi_{MDP}(s)))$ where $\delta(a, b) = (1 \text{ if } a=b; 0 \text{ otherwise})$
 - "Q-MDP" (Cassandra, 1995) - $argmax_a(\sum_{s \in S} b(s) Q_{MDP}(s, a))$
- Essentially, try to act optimally as if the POMDP were to become observable after the next action
 - **Cannot plan to do actions just to gain information**

Greedy Approaches to POMDP Planning

- “Dual-Mode Control” (Cassandra 1996) –

Extension to Allow Information-Gathering Actions

- Compute entropy $H(b)$ of belief state
- If entropy is below a threshold, use a greedy method $Z(a, b)$ for choosing action
- If entropy is above a threshold, choose the action that reduces expected entropy the most

$$EE(a, b) = \sum_{b'} p(b' | a, b) H(b')$$

$$\pi(s) = \begin{cases} \operatorname{argmax}_a Z(a, b) & \text{if } H(b) < t \\ \operatorname{argmin}_a EE(a, b) & \text{otherwise} \end{cases}$$