Planning and Optimization

G4. Heuristics for Probabilistic Planning

Malte Helmert and Thomas Keller

Universität Basel

December 9, 2019

Content of this Course

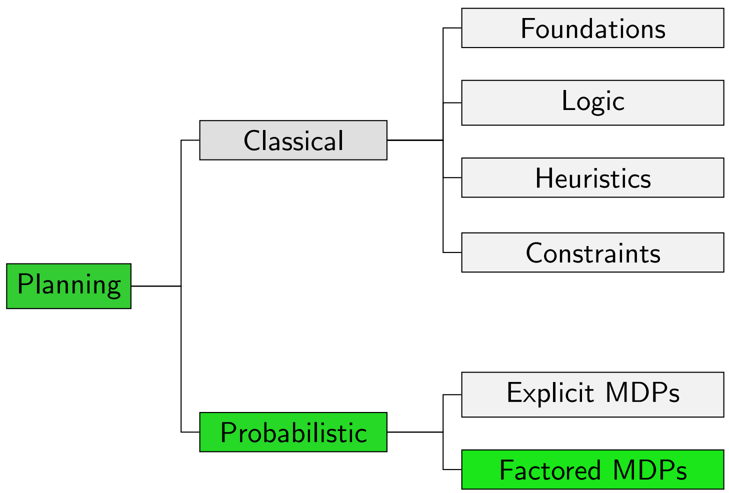

Content of this Course: Factored MDPs

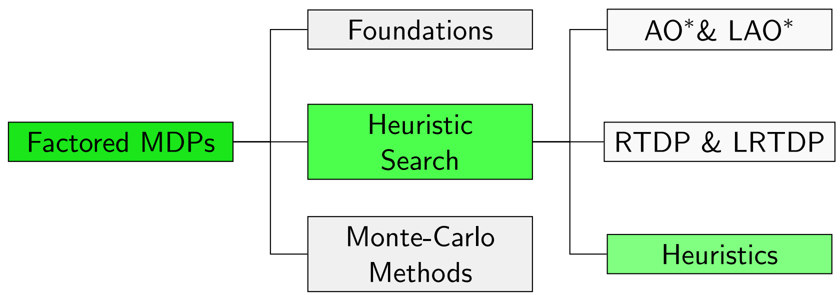

G4.1 Motivation

Heuristics

Heuristics for probabilistic planning

- most heuristics are based on domain-specific knowledge
- for a long time, determinization-based heuristics were the only notable domain-independent solution
- Recent progress: occupation measures (only covered in exercises)

G4.2 Determinization

What is a Determinization?

- Replace SSP operators with deterministic operators
- Results in classical planning task
- For SSPs, this is a classical planning task
- SSP and its determinization are related but not equivalent

How to Come up with a Determinization?

Typically, two types of determinization are distinguished:

- All-outcomes determinization
  - Create one deterministic transition to each outcome
- Single-outcome determinization
  - Pick one outcome of each probabilistic transition ...
  - ... and turn it into a deterministic transition
  - often, the most likely outcome is preserved
  - or one outcome is sampled according to its probability

All-outcomes Determinizion

Definition (All-outcomes Determinization)

Let $\mathcal{T} = \langle S, L, c, T, s_0, S_\star \rangle$ be an SSP. The all-outcomes determinization of $\mathcal{T}$ is the (deterministic) transition system $\mathcal{T}^d = \langle S, L^d, c, T^d, s_0, S_\star \rangle$ with

$$L^d = \{\langle s, \ell, s' \rangle \mid s, s' \in S, \ell \in L \text{ and } T(s, \ell, s') > 0\}$$
$$T^d = \{\langle s, \ell^d, s' \rangle \mid \ell^d = \langle s, \ell, s' \rangle \in L^d\}.$$

All-outcomes Determinization: (Unit-cost) Example

Generate one action for each outcome

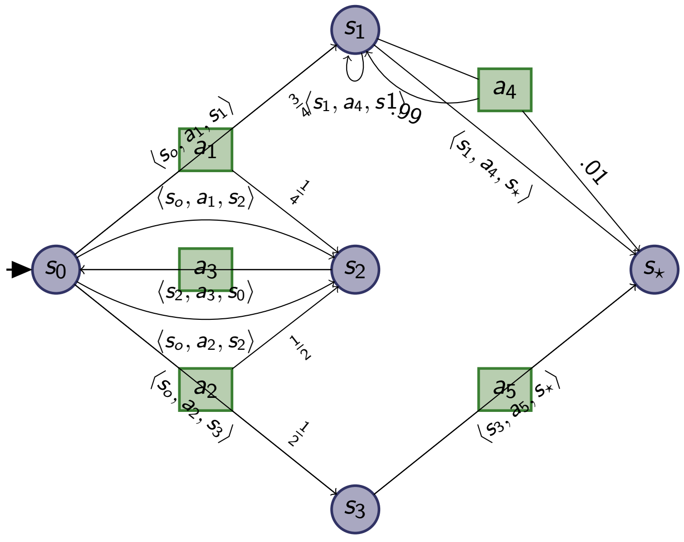

Single-outcome Determinization

Definition (Single-outcome Determinization)

Let $\mathcal{T} = \langle S, L, c, T, s_0, S_\star \rangle$ be an SSP and let

$$L^{aod} = \{\langle s, \ell, s' \rangle \mid s, s' \in S, \ell \in L \text{ and } T(s, \ell, s') > 0\}.$$

A (deterministic) transition system $\mathcal{T}^d = \langle S, L^d, c, T^d, s_0, S_\star \rangle$ is a single-outcome determinization of $\mathcal{T}$ if

- $L^d \subseteq L^{aod}$ s.t. for all $s \in S$ and $\ell \in L(s)$ there is exactly one $\langle s, \ell, s' \rangle \in L^d$ and
- $T^d = \{\langle s, \ell^d, s' \rangle \mid \ell^d = \langle s, \ell, s' \rangle \in L^d\}$

Single-outcome Determinization: (Unit-cost) Example

Remove transition to all outcomes but one

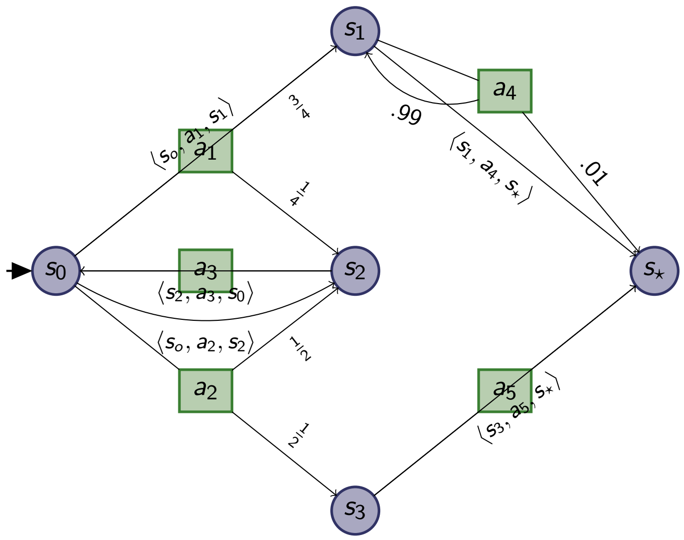

G4.3 Properties

Single-outcome Determinization: Properties

- single-outcome determinizations are not well-suited to be used as a heuristic:
  - can be inadmissible
  - and even unsafe

Single-outcome Determinization: Inadmissibility

Remove transition to all outcomes but one

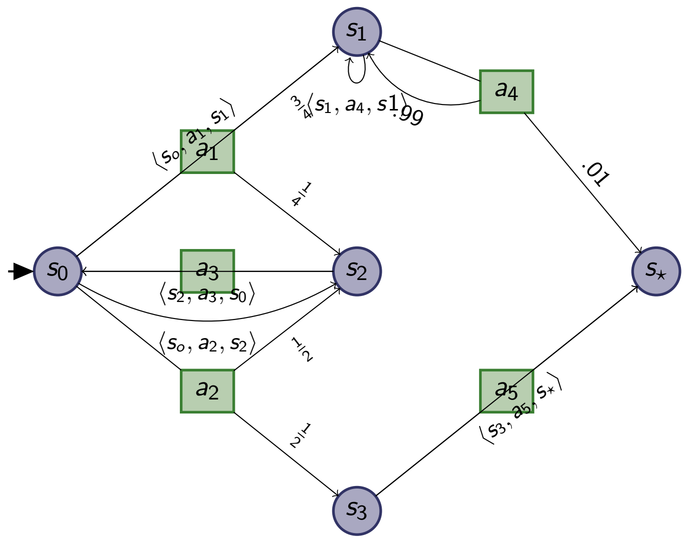

Single-outcome Determinization: Properties

- single-outcome determinizations are not well-suited to be used as a heuristic:
  - can be inadmissible
  - and even unsafe
- often part of domain-specific solutions (e.g., by ignoring that some action may "fail")
- and as part of algorithms that average over several samples (Chapter G5)

Min-min Heuristic

Definition (Min-min Heuristic)

Let $\mathcal{T}$ be an SSP and let $\mathcal{T}^d$ be the all-outcomes determinization of $\mathcal{T}$. The min-min heuristic h_{min} maps each state $s \in S$ to the cost of the cheapest path from s to a goal state in $\mathcal{T}^d$.

Min-min Heuristic: Admissibility

Theorem (Admissibility)

The min-min heuristic is admissible.

Proof Sketch.

$$\begin{aligned} V^\star(s) &\overset{(1)}{=} \min_{\ell \in L(s)} c(\ell) + \sum_{s' \in S} T(s, \ell, s') \cdot V^\star(s') \\ &\overset{(2)}{=} \min_{\ell \in L(s)} c(\ell) + \sum_{s' \in \text{succ}(s,\ell)} T(s, \ell, s') \cdot V^\star(s') \\ &\overset{(3)}{\geq} \min_{\ell \in L(s)} c(\ell) + \min_{s' \in \text{succ}(s,\ell)} V^\star(s') = h_{\text{min}}(s) \end{aligned}$$

(1) is the Bellman equation, (2) holds because $T(s, \ell, s') = 0$ for all $s' \notin \text{succ}(s, \ell)$ and (3) because the weight from more expensive outcomes is shifted to the cheapest one.

Min-min Heuristic: Properties

- The min-min heuristic is also called the optimistic heuristic ...
- ... because the planner may choose its preferred outcome
- Min-min heuristic can be well informed ...
- ... but can also be utterly optimistic

All-outcomes Determinization: Quality

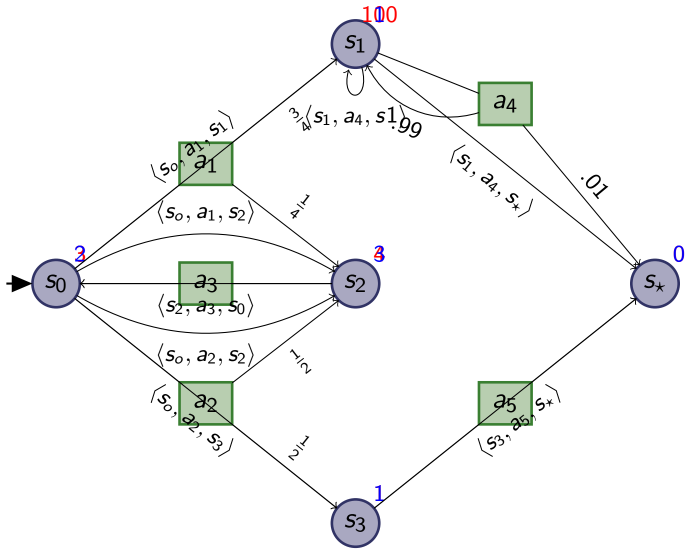

Annotation: probabilistic state-values (in red)deterministic state-values (in blue)

Min-min Heuristic: Properties

- The min-min heuristic is also called the optimistic heuristic ...
- ... because the planner may choose its preferred outcome
- Min-min heuristic can be well informed ...
- ... but can also be utterly optimistic
- Min-min heuristic often solvable in practice even if SSP is not
- If still unsolvable: compute classical heuristic of all-outcomes determinization

G4.4 Summary

Summary

- Almost all heuristics in probabilistic planning are either domain-specific or based on a determinization
- A single-outcome determinization removes all outcomes from a transition except for one
- The all-outcomes determinization creates a deterministic transition for each outcome
- The min-min heuristic computes the shortest path in the all-outcomes determinization
- The min-min heuristic is admissible