AI Planning
Lecture 10

Rostislav Horčík

Czech Technical University in Prague
Faculty of Electrical Engineering
xhorcik@fel.cvut.cz

Factored representation

Probabilistic FDR

Definition

Probabilistic FDR (PFDR) task is $\Pi = \langle V, A, s_0, g \rangle$ where

- V is a finite set of variables, each $v \in V$ endowed with a finite domain $\text{dom}(v)$,
- s_0 is an initial state,
- g is a partial state called goal.
- A is a finite set of actions together with a cost function $\text{cost}\colon A \to \mathbb{R}^+$.

Action $a = \langle \text{pre}_a, \text{eff}_a \rangle$ where

- pre_a is a partial state called precondition
- eff_a is a set of partial states called effects. There is a probability distribution P_a over eff_a.

Induced SSP

An action $a \in A$ is applicable in a state s if $\text{pre}_a \subseteq s$.

For a state s, action a and effects $e \in \text{eff}_a$, we define:

$$e(s) = \begin{cases} e(v) & \text{if } e(v) \text{ is defined,} \\ s(v) & \text{otherwise.} \end{cases}$$

Each probabilistic FDR task $\Pi = \langle V, A, s_0, g \rangle$ induces an SSP $\Sigma_\Pi = \langle S, A, T, s_0, G \rangle$ where

- S is the set of all states,
- $G = \{s \in S \mid g \subseteq s\}$,
- $T(s, a) = \{e(s) \mid e \in \text{eff}_a, \text{pre}_a \subseteq s, a \in A\}$ and $P(e(s) \mid s, a) = P_a(e)$.

Heuristics

All-outcome determinization

Definition

Let $\Sigma = \langle S, A, T, s_0, G \rangle$ be an SSP. Its all-outcome determinization is an LTS $\Sigma_d = \langle S, A_d, T_d, s_0, G \rangle$ where

- $A_d = A \times \{1, \dots, m\}$, $m = \max_{s \in S, a \in A} |T(s, a)|$,
- $T_d = \{\langle s, \langle a, i \rangle, t_i \rangle \mid T(s, a) = \{t_1, \dots, t_n\}\}$.

Given a PFDR task $\Pi = \langle V, A, s_0, g \rangle$, we can make an FDR task $\Pi_d = \langle V, A_d, s_0, g \rangle$ such that $A_d = \{\langle \mathrm{pre}_a, e \rangle \mid e \in \mathrm{eff}_a, a \in A\}$.

LTS induced by Π_d is the all-outcome determinization of the SSP induced by Π.

Admissible heuristics from classical planning

Lemma
Let Σ be an SSP. Any admissible heuristic h for its all-outcome determinization Σ_d is an admissible heuristic for Σ.

Syntactic projection

Definition

Let $\Pi = \langle V, A, s_0, g\rangle$ be a PFDR task, $P \subseteq V$ a pattern, and $\pi_P(s) = \{\langle v, d\rangle \in s \mid v \in P\}$. The syntactic projection of Π to P is a PFDR task $\Pi_P = \langle P, A_P, \pi_P(s_0), \pi_P(g)\rangle$ where

- $A_P = \{a_P \mid a \in A\}$,
- $a_P = \langle \pi_P(\mathrm{pre}_a), \pi_P(\mathrm{eff}_a)\rangle$,
- $\pi_P(\mathrm{eff}_a) = \{\pi_P(e) \mid e \in \mathrm{eff}_a\}$, and
- $$P_{\pi_P(a)}(\pi_P(e)) = \sum_{\substack{e' \in \mathrm{eff}_a \\ \pi_P(e') = \pi_P(e)}} P_a(e').$$

Finally, $\mathrm{cost}(a_P) = \mathrm{cost}(a)$.

Pattern databases

Definition

Let $\Pi = \langle V, A, s_0, g \rangle$ be a PFDR task and $P \subseteq V$ a pattern. Further, let V_P^* be the optimal value function for the SSP Σ_{Π_P} induced by the syntactic projection Π_P. The PDB heuristic for P is defined by $h_P(s) = V_P^*(\pi_P(s))$ for any state s.

Theorem

The heuristic $h_P(s) = V_P^*(\pi_P(s))$ *is goal-aware and consistent, thus admissible.*