Principles of AI Planning

Prof. Dr. B. Nebel, Dr. R. Mattmüller
D. Speck, T. Schulte, M. Kantz
Winter Semester 2019/2020

University of Freiburg
Department of Computer Science

Exercise Sheet 6

Due: Friday, December 6th, 2019

Send your solution to `mario.kantz@gmail.com` (PDF only) or submit a hardcopy before the lecture. The exercise sheets may and should be worked on and handed in in groups of two or three students. Please indicate all names on your solution.

Exercise 6.1 (Delete relaxation, 2+2 points)

Consider the planning task $\Pi = \langle A, I, O, \gamma \rangle$ in positive normal form with

$$\begin{aligned}
A &= \{\mathit{haveCake}, \mathit{eatenCake}, \mathit{haveNoCake}\}, \\
I &= \{\mathit{have\text{-}cake} \mapsto 0, \mathit{eatenCake} \mapsto 0, \mathit{haveNoCake} \mapsto 1\} \\
O &= \{\mathit{eatCake}, \mathit{bakeCake}\}, \\
\mathit{eatCake} &= \langle \mathit{haveCake}, \neg\mathit{haveCake} \wedge \mathit{haveNoCake} \wedge \mathit{eatenCake} \rangle, \\
\mathit{bakeCake} &= \langle \mathit{haveNoCake}, \mathit{haveCake} \wedge \neg\mathit{haveNoCake} \rangle \text{ und} \\
\gamma &= \mathit{haveCake} \wedge \mathit{eatenCake}.
\end{aligned}$$

(a) Give the relaxation Π^+ of Π.
(b) Give a sequence π of operators (as short as possible) from O such that π is *not* a plan of Π, but π^+ is a plan of Π^+.

Exercise 6.2 (h^+ heuristic, 3+3 points)

A 15-puzzle planning task $\Pi = \langle A, I, O, \gamma \rangle$ is given as

$$\begin{aligned}
A &= \{\mathit{empty}(p_{i,j}) \mid 0 \le i, j \le 3\} \cup \{\mathit{at}(t_k, p_{i,j}) \mid 0 \le i, j \le 3, 0 \le k \le 14\}, \\
O &= \{\mathit{move}(t_m, p_{i,j}, p_{k,l}) \mid 0 \le i, j, k, l \le 3, 0 \le m \le 14, \\
&\quad (i = k \text{ and } |j - l| = 1) \text{ or } (j = l \text{ and } |i - k| = 1)\}, \\
\gamma &= \bigwedge_{0 \le m \le 14} \mathit{at}(t_m, p_{\lfloor m/4 \rfloor, m\%4})
\end{aligned}$$

Action $\mathit{move}(t_m, p_{i,j}, p_{k,l})$ moves tile t_m from position $p_{i,j}$ to position $p_{k,l}$:

$$\begin{aligned}
\mathit{move}(t_m, p_{i,j}, p_{k,l}) = \langle &\mathit{at}(t_m, p_{i,j}) \wedge \mathit{empty}(p_{k,l}), \\
&\mathit{at}(t_m, p_{k,l}) \wedge \mathit{empty}(p_{i,j}) \wedge \neg\mathit{at}(t_m, p_{i,j}) \wedge \neg\mathit{empty}(p_{k,l}) \rangle
\end{aligned}$$

A syntactically possible state is *legal* if each tile t_m is at some position p_{ij}, if no two tiles are at the same position and if the remaining position is the only one that is *empty*. The initial state is an arbitrary state that is legal.

One possible heuristic for the 15-puzzle is the Manhattan-distance heuristic $h^{Manhattan}$: It sums the Manhattan distances of all tiles from their current positions to their target positions, where the Manhattan distance between position $p_{i,j}$ and $p_{k,l}$ is given as $|i - k| + |j - l|$.
The h^+ heuristic estimates the distance of state s to the closest goal state as the length of the optimal plan in the relaxed planning task (with initial state s).

(a) Show that $h^+(s) \ge h^{Manhattan}(s)$ for each legal state s of a 15-puzzle planning task.
(b) Show that $h^+(s) > h^{Manhattan}(s)$ for at least one state s of a 15-puzzle planning task.