Principles of AI Planning

15. Strong nondeterministic planning

Albert-Ludwigs-Universität Freiburg

Bernhard Nebel and Robert Mattmüller

January 11, 2019

Strong planning

In this chapter, we will consider the simplest case of nondeterministic planning by restricting attention to strong plans.

Concepts

Strong plans

Recall the definition of strong plans:

Definition (strong plan)

Let S be the set of states of a planning task Π. Then a strong plan for Π is a function $\pi : S_\pi \to O$ for some subset $S_\pi \subseteq S$ such that

- $\pi(s)$ is applicable in s for all $s \in S_\pi$,
- $S_\pi(s_0) \subseteq S_\pi \cup S_\star$ (π is closed),
- $S_\pi(s') \cap S_\star \neq \emptyset$ for all $s' \in S_\pi(s_0)$ (π is proper), and
- there is no state $s' \in S_\pi(s_0)$ such that s' is reachable from s' following π in a strictly positive number of steps (π is acyclic).

Execution of a strong plan

1. Determine the current state s.
2. If s is a goal state then terminate.
3. Execute action $\pi(s)$.
4. Repeat from first step.

Strong plans

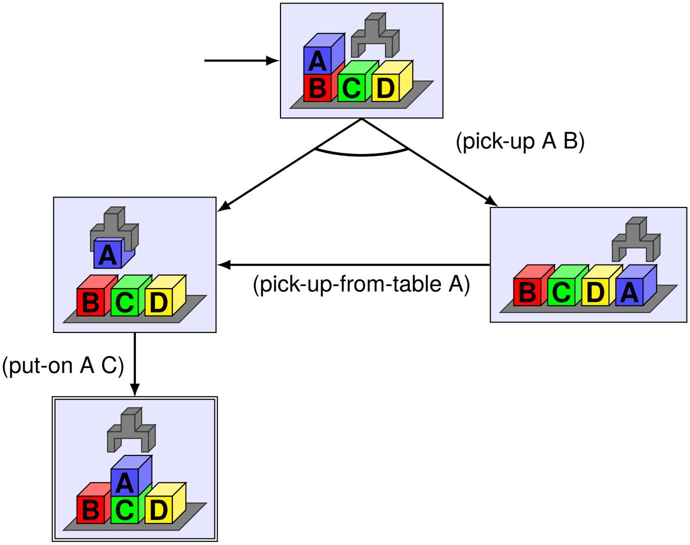

Image

The image of a set T of states with respect to an operator o is the set of those states that can be reached by executing o in a state in T.

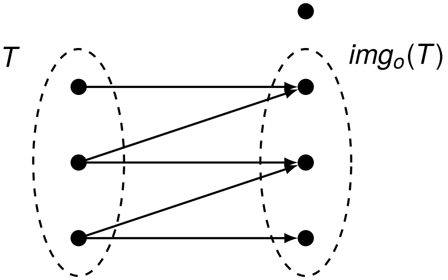

Images

Definition (image of a state)

$img_o(s) = \{s' \in S | s \xrightarrow{o} s'\} = app_o(s)$

Definition (image of a set of states)

$img_o(T) = \bigcup_{s \in T} img_o(s)$

Weak preimage

The weak preimage of a set T of states with respect to an operator o is the set of those states from which a state in T can be reached by executing o.

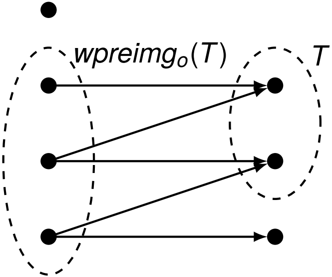

Weak preimages

Definition (weak preimage of a state)

$wpreimg_o(s') = \{s \in S | s \xrightarrow{o} s'\}$

Definition (weak preimage of a set of states)

$wpreimg_o(T) = \bigcup_{s \in T} wpreimg_o(s)$.

Strong preimage

The strong preimage of a set T of states with respect to an operator o is the set of those states from which a state in T is always reached when executing o.

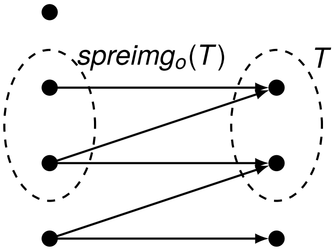

Strong preimages

Definition (strong preimage of a set of states)

$$spreimg_o(T) = \{s \in S \mid \exists s' \in T : s \xrightarrow{o} s' \wedge img_o(s) \subseteq T\}$$

Algorithms

Algorithms for strong planning

1. Dynamic programming (backward)
   Compute operator/distance/value for a state based on the operators/distances/values of its all successor states.
   1. Zero actions needed for goal states.
   2. If states with i actions to goals are known, states with $\leq i+1$ actions to goals can be easily identified.

   Automatic reuse of plan suffixes already found.
2. Heuristic search (forward)
   Strong planning can be viewed as AND/OR graph search.

   OR nodes: Choice between operators

   AND nodes: Choice between effects

   Heuristic AND/OR search algorithms:
   AO*, Proof Number Search, . . .

Dynamic programming

Planning by dynamic programming

If for all successors of state s with respect to operator o a plan exists, assign operator o to s.

- Base case $i = 0$: In goal states there is nothing to do.
- Inductive case $i \geq 1$: If $\pi(s)$ is still undefined and there is $o \in O$ such that for all $s' \in img_o(s)$, the state s' is a goal state or $\pi(s')$ was assigned in an earlier iteration, then assign $\pi(s) = o$.

Backward distances

If s is assigned a value on iteration $i \geq 1$, then the backward distance of s is i. The dynamic programming algorithm essentially computes the backward distances of states.

Example

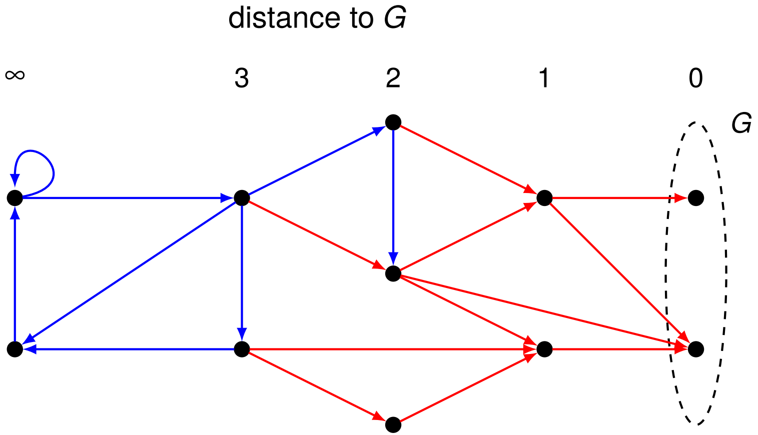

Backward distances

Definition (backward distance sets)

Let G be a set of states and O a set of operators.
The backward distance sets D_i^{bwd} for G and O consist of those states for which there is a guarantee of reaching a state in G with at most i operator applications using operators in O:

$$D_0^{bwd} := G$$
$$D_i^{bwd} := D_{i-1}^{bwd} \cup \bigcup_{o \in O} spreimg_o(D_{i-1}^{bwd}) \text{ for all } i \geq 1$$

Backward distances

Definition (backward distance)

Let G be a set of states and O a set of operators, and let $D_0^{bwd}, D_1^{bwd}, \ldots$ be the backward distance sets for G and O. Then the backward distance of a state s for G and O is

$$\delta_G^{bwd}(s) = \min\{i \in \mathbb{N} \mid s \in D_i^{bwd}\}$$

(where $\min \emptyset = \infty$).

Strong plans based on distances

Let $\Pi = \langle V, I, O, \gamma \rangle$ be a nondeterministic planning task with state set S and goal states $S_\star$.

Extraction of a strong plan from distance sets

1. Let $S' \subseteq S$ be those states having a finite backward distance for $G = S_\star$ and O.
2. Let $s \in S'$ be a state with distance $i = \delta_G^{bwd}(s) \geq 1$.
3. Assign to $\pi(s)$ any operator $o \in O$ such that $img_o(s) \subseteq D_{i-1}^{bwd}$. Hence o decreases the backward distance by at least one.

Then π is a strong plan for $\mathscr{T}$ iff $I \in S'$.

Question: What is the worst-case runtime of the algorithm?

Summary

Summary

- We have considered the special case of nondeterministic planning where
  - planning tasks are fully observable and
  - we are interested in strong plans.
- We have introduced important concepts also relevant to other variants of nondeterministic planning such as
  - images and
  - weak and strong preimages.
- We have discussed one basic classes of algorithms: backward induction by dynamic programming.