

L15: Planning with factored representations

AIMA4e: Chapter 11.1–11.3; 11.5

What you should know after this lecture


- STRIPS planning representation takes advantage of relations, factoring, locality, and sparsity to make transition model compact
- STRIPS models enable powerful domain-independent heuristics
- We can model partial information, and other extensions, in this formalism

Factored states and information

Factored, discrete states

Compact, sparse representation of T

Construct heuristics via relaxation



Making plans in complex domains

- We have seen how to frame planning for an agent as searching for a path through a state space.
- We have also seen how to describe states using a factored representation in terms of variables and values
- Can we combine them? Yes, with the following advantages:
 - Factoring state representations lets us compactly describe the goals and transition model
 - Factored structure enables a lot of relaxations that lead to powerful domain independent heuristics

“Classical planning” framework

- Make some structural assumptions about the domain
 - sparsity of effect: any action taken by an agent doesn't change many aspects of the environment state
 - locality of dependence: what effects an action will have depend only on a few aspects of the environment state
- Leads us to a special-purpose (but still domain independent) representation language for describing \mathcal{S} , \mathcal{A} , T , and G that
 - Is highly compact (and therefore learnable from few samples)
 - Can be used to plan efficiently
- Language is called STRIPS; standardized syntax and variations in PDDL (planning domain description language)

Planning domain description language

For now we are following syntax from AIMA—we'll show later what the “real” syntax is like.

Domain specification

- predicates: symbols, like *On* or *Airport*
- object variables: symbols, like *x*
- fluents: atoms, like *On(x, y)*
// These are the factors of our state representation
- operators: schematic, factored, description of T, like

Unload(obj, plane, loc)

- preconditions: *Aboard(obj, plane), At(plane, loc)*
- effect: *At(obj, loc), ¬Aboard(obj, plane)*

Planning domain description language

A ground fluent is a predicate applied to a tuple of constant symbols.

Problem specification

- constants: symbols, like *blockA* or *747_e35b2*
- initial state: set of ground fluents that are true in the initial state; assume all other ground fluents are false.
(This is often called the closed world assumption.)
- goal: conjunction (set) of ground fluents

Path-search problem given PDDL domain and problem

Mapping this back into the representation we used for path search problems

- \mathcal{S} :
 - Plug all combinations of constants into all predicates to get all ground fluents, like *Aboard(blockA, 747_e35b2)*
 - A state is an assignment of **True** or **False** to each ground fluent.
 - It is often most efficient to represent a state as the set of ground fluents that have the value **True**.
- \mathcal{A} : Plug all combinations of constants into all operators to get all ground operators. These are the possible actions.
- $G \subset \mathcal{S}$: All states in which all ground fluents in the goal are assigned to **True**
- s_0 : The initial state, set of ground fluents that are true initially

State transition function

Define $T(s, a)$ where

- s : set of true ground fluents
- a : ground operator instance

as follows:

- If $preconditions(a) \subseteq s$ then

$$T(s, a) = s - del(a) \cup add(a)$$

where $add(a)$ are positive fluents in $effects(a)$ and $del(a)$ are negated fluents in $effects(a)$

- Otherwise, the operator a is not applicable in state s , and we can think of it as having no effect, so

$$T(s, a) = s$$

Planning algorithms

Given a domain and problem description, how do we find a plan?

- Forward best-first search with heuristics that take advantage of the structured representation
- Regression (or backward chaining), works backwards from the goal, states in the search space are actually sets of fluents representing sub-goals (not environment states)
- Reduction to propositional satisfiability.

Why is this formalism useful?

- The domain description is independent of the particular universe of objects (constants)
- Similar in some ways to a graph neural network (you can think of nodes for fluents in the problem instance; the operator description specifies connectivity (which other fluents the new value of a fluent depends on) and parameters (what those fluent values actually are.)
- Generalizes broadly
- Takes advantage of sparsity
 - The effects of most actions don't depend on most factor values
 - Relatively few factors are affected by any action
- Provides leverage for defining effective domain-independent heuristics

Delete relaxation

- The thing that makes planning difficult is interference among the operators—executing an action might potentially undo some effect that you had already achieved or wanted to maintain from the initial state.
- A relaxation of the planning problem is to assume that this never happens, by ignoring the delete effects of an operator, so that our update is:

$$s' = s \cup add(a)$$

- In this relaxation, a fluent never become false once it becomes true! So, e.g. a robot can be in multiple locations. Weird, but convenient.
- An even more relaxed relaxation: allow all actions whose preconditions are satisfied to be executed in parallel!

The relaxed planning graph (RPG) is computed by computing a sequence of relaxed, parallel state updates.

Example relaxed plan graph

Left

- $P = \text{In}(\text{Robot}, R_2)$
- $D = \text{In}(\text{Robot}, R_2)$
- $A = \text{In}(\text{Robot}, R_1)$

Suck(R_1)

- $P = \text{In}(\text{Robot}, R_1)$
- $D = \emptyset$ [empty set]
- $A = \text{Clean}(R_1)$

Right

- $P = \text{In}(\text{Robot}, R_1)$
- $D = \text{In}(\text{Robot}, R_1)$
- $A = \text{In}(\text{Robot}, R_2)$

Suck(R_2)

- $P = \text{In}(\text{Robot}, R_2)$
- $D = \emptyset$ [empty set]
- $A = \text{Clean}(R_2)$





Image source: J.C.Latombe, Stanford

Compute relaxed planning graph

COMPUTE-RPG(s_0, \mathcal{A}, G)

```
1  //  $s_0$  and  $G$  are both sets of ground fluents
2  //  $\mathcal{A}$  is a set of ground operator descriptions
3   $F_0 = s_0; t = 0$ 
4  while  $G \not\subseteq F_t$ 
5       $A_t = \{a \in \mathcal{A} \mid pre(a) \subseteq F_t\}$            // Do all applicable actions!
6       $F_{t+1} = F_t \cup \bigcup_{a \in A_t} add(a)$          // Add all add effects!
7      if  $F_{t+1} = F_t$ : return                        // Goal is infeasible ☹️
8       $t = t + 1$ 
9  return  $F_0, \dots, F_t, A_0, \dots, A_{t-1},$ 
```

Heuristics based on RPG

- Add up the levels at which each goal fluent appear: not admissible

$$H_{\text{add}}(s, G) = \sum_{f \in G} \underset{t}{\operatorname{argmin}} f \in F_t$$

- Max of the levels at which each goal fluent appear: admissible but weak

$$H_{\text{max}}(s, G) = \max_{f \in G} \underset{t}{\operatorname{argmin}} f \in F_t$$

- Optimal solution to the delete-relaxation problem (without parallel actions): still NP-hard!
- H_{ff} : Approximate solution to the delete-relaxation problem, searching backward in the RPG for a relaxed plan

Computing H_{ff}

$H_{ff}(s, G, RPG)$

```
1   $M = \max_{f \in G} RPG.level(f); \text{plan} = \{\}$ 
2  for  $t \in 0 \dots M$ :      // Fluents we need to make true at each level
3       $G_t = \{f \in G \mid RPG.level(f) = t\}$ 
4  for  $t = M \dots 1$ :
5      for  $f \in G_t$ :      // Find any applicable  $a$  with result  $f$ 
6           $a = \{a \mid RPG.level(a) = t - 1, f \in add(a)\}[0]$ 
7           $\text{plan} = \text{plan} \cup \{a\}$       // Add action  $a$  to plan
8          for  $p \in pre(a)$ 
9               $G_{RPG.level(p)} = G_{RPG.level(p)} \cup \{p\}$ 
10 return  $|\text{plan}|$ 
```

$RPG.level(f) = \min_t f \in F_t$

$RPG.level(a) = \min_t a \in A_t$

Extensions

There are lots of extensions to classical planning!

- Conformant planning: have, for each predicate P , BP and $BNotP$.
- Temporal planning: discrete time steps, actions take time
- Cost-sensitive planning: add action costs
- Conditional planning: add observe actions

Actual PDDL syntax example

LISP and prefix syntax used to be a thing!

```
(:action unload
  :parameters (?obj ?plane ?loc)
  :precondition (and
    (package ?obj)
    (plane ?plane)
    (location ?loc)
    (at ?plane ?loc)
    (aboard ?obj ?plane))
  :effect (and
    (not (aboard ?obj ?plane))
    (at ?obj ?loc)))
```

Next time

- Action planning in continuous state and action spaces