

CSC384:Lecture7

- Lasttime

- `gametreesearch(note: notextreadings)`

- Today

- Intro to Planning; State and action representations
- planning: start on STRIPS planning

- Readings:

- Today: Ch. 8.1, 8.2 (STRIPS, **skip** EventCalculus), 8.3 (forward planning)
- Next week: 8.3 (STRIPS planning in depth, regression planning, briefly resolution-based planning)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

1

SearchinComplexSituations

- Search algorithms so far are graph-based:

- produces sequence of actions or moves that change the world in specific ways (start \rightarrow goal)
- has nothing to say about the actions other than that take you from one state to another
 - e.g., map location (courier), robot position
- states have no structure to speak of

- Realistic problems involve altering complex states of the world

- at each state, many different facts are true/false
- actions change the truth of different facts

CSC 384 Lecture Slides (c) 2002, C. Boutilier

2

Planning

- Planning problems are like search problems

- given an initial state of the world
- given some goal **conditions** (facts we would like to make true, in contrast to goal **states**)
- find a sequence of actions that will take you from the start state to some states satisfying the goal conditions (note: these **are** goal states, but specified differently)

- What's different?


- states are complex entities (possible worlds)
- actions tend to affect only certain facts

- End result: planning algorithms look a little different than search algorithms (usually)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

3

APlanningProblem



CSC 384 Lecture Slides (c) 2002, C. Boutilier

4

RobotExample:Domain

- Predicates:

- **loc(X)** - locationofrobotisX
- **adj(X,Y)** - locationsX,Yadjacent
- **rhk** - robothaskeys
- **cm** - coffeemade
- **chc** - craighascoffee(ormaybe
hc(P))
- **rhc** - robothascoffee(ormaybe
carrying(O))
- **lt** - labistdy

- Derived predicates possibly:

- ```
accessible(X,Y) ← adj(X,Y)&rhk.
accessible(X,Y) ← adj(X,Y)&unlck(X).
```

CSC 384 Lecture Slides (c) 2002, C. Boutilier

5

## StatesoftheWorld

- A state is an assignment of truth values to all (relevant) atoms

- e.g., think of it as an “interpretation” or *complete KB*
- it is a way the world could be

- Formally, given a finite number of predicates and domain objects, a **worldstate** is any assignment of truth values to *all ground atoms*

- e.g., True: loc(off),adj(cr,hall),...,rhc,rhk, etc...  
False: loc(hall),loc(mr),...,adj(off,hall ),...cm,  
chc, etc.

CSC 384 Lecture Slides (c) 2002, C. Boutilier

6

## Robot Example: Actions

- What actions might we consider?
  - `move(X,Y)`, `getkeys`, `dropkeys`, `makecoffee`, `grabcoffee` or `grab(X) ?`, `tidylab`, etc...
- Almost all actions have **preconditions**
  - can't apply an action at state S unless all preconditions are satisfied
  - `move(X,Y)` requires `accessible(X,Y)`
  - `givecoffee` at S requires `loc(off)`, `rhc`, etc.
- Actions change the state of the world
  - `move(X,Y)` makes `loc(X)` false, `loc(Y)` true
  - `givecoffee` makes `rhc` false, `chc` true
  - these actions affect no other facts!

CSC 384 Lecture Slides (c) 2002, C. Boutilier

7

## Neighbor Representation

- For any state S, action A, we could specify the state resulting from applying S at A; but...
- How many states?
  - if k ground atoms,  $2^k$  states (so far our trivial robot domain has about 40 ground atoms!)
- Action effects have a lot of structure
  - actions tend to affect only a few atoms, most unaffected
- But if we decided to do this (and ignore these issues), we could apply our favorite search algorithms to this planning problem!

CSC 384 Lecture Slides (c) 2002, C. Boutilier

8

## Planning Basics

- Almost all work in planning relies on some basic principles
  - use specific state representations that are easy to manipulate
  - use specific action representations that exploit the fact that actions affect only a few atoms generally
  - devise planning algorithms that exploit the "locality" of actions and the fact that we're concerned with only a few goal propositions
- We'll look at each of these in turn

CSC 384 Lecture Slides (c) 2002, C. Boutilier

9

## State Representation

- A world state is just like a KB, but we can't just write a bunch of facts/rules to assert what's true
  - we'll need to look at a bunch of worlds, not just one
  - and we can't retract/assert things in KB (easily)
- So we need a *state representation*
- Consider the state:
  - True: `loc(c)`, `adj(o,m)`, `adj(m,o)`, `adj(o,l)`, `adj(l,o)`, ..., `rhc`, `cm`, `lck(c)`, `lck(l)`, ...
  - False: `loc(o)`, `loc(l)`, `loc(m)`, `loc(h)`, `adj(c,o)`, etc..., `chc`, `rhc`, `lck(o)`, `lck(h)`, ...

CSC 384 Lecture Slides (c) 2002, C. Boutilier

10

## Explicit World Repres'n (EWR)

- EWR the simplest possible state representation
  - Keep two lists of ground atoms: *true list*, *false list*
  - or use negation and have one list:
  - e.g., `[loc(c), neg(loc(o)), neg(loc(m)), rhc, neg(chc)...]`
- Difficulties:
  - list has one entry for each ground atom in language
  - generally, there are far more negative literals than positive literals
  - e.g., `adj(X,Y)` false for many more pairs than true; `loc(X)` false for all but one location;
  - consider airline planner with relation `flight(City1,City2,FltNum)`: many false instances!!!

CSC 384 Lecture Slides (c) 2002, C. Boutilier

11

## Closed World Representation (CWR)

- CWR makes the closed world assumption
  - You specify every atom that is true (e.g., in a list)
  - If atom is not in list, it is assumed false
  - A state is a list of *positive ground facts*
  - e.g., `[loc(c), rhc, cm, adj(o,m), adj(m,o)...]`
- Give state S (list), when is literal L true at S?

```
holds(Atom,State) :- member(Atom,State).
```

```
holds(neg(Atom), State) :- not(member(Atom,State)).
```

CSC 384 Lecture Slides (c) 2002, C. Boutilier

12

## Meta-Interpreter

- The *holds* predicate: a (simple) *meta-interpreter*
  - the list *State* is our own little 'KB', built into a list
  - the *holds* predicate asks a "query" of this 'KB': is the (pos'tv or neg'tv) Fact true in *State* (our mini 'KB')
  - we allow *holds* to ask queries about negative literals (and we exploit Prolog's negation-as-failure mechanism to implement the CWA)
  - Thus we've implemented a (trivial) Prolog-like query interpreter in Prolog
- Note: Elements in *State* are "atoms" in our domain language, but are just *terms* in Prolog
  - allows us to treat a state ('KB') as an object in Prolog
  - why is this important?

CSC 384 Lecture Slides (c) 2002, C. Boutilier

13

## Derived Relations in CWR (CWR-D)

- Certain facts derivable from other facts
  - e.g., the *accessible(X,Y)* relation
- Don't want these explicitly represented in state
  - they are redundant (can be derived from others)
  - they complicate action rep'n (potential inconsistency)
  - e.g., robot drops keys (-rhk): what happens to accsbl?
- So we can separate *basic* relations in our domain from *derived* relations
  - only ground atoms of *basic* type allowed in state
- For each domain "predicate", must specify type

CSC 384 Lecture Slides (c) 2002, C. Boutilier

14

## Derived Relation: Example

- Basic relations specified as follows:

```
baseRel(rhk).
baseRel(loc(X)).
baseRel(lck(X)). etc...
```

- Derived relations include how they are derived:

```
derivedRel(accessible(X), [neg(lck(X))]).
derivedRel(accessible(X), [lck(X), rhk]).
```

Like Prolog rules:  
*accessible(X) :- not(lck(X))*  
*accessible(X) :- lck(X), rhk.*

CSC 384 Lecture Slides (c) 2002, C. Boutilier

15

## A Meta-Interpreter for CWR-D

```
holds(Fact, State) :- baseRel(Fact), member(Fact, State).
holds(neg(Fact), State) :- baseRel(Fact),
 not(member(Fact, State)). *
holds(Fact, State) :- derivedRel(Fact, BodyList),
 holdsAll(BodyList, State).
holds(neg(Fact), State) :- derivedRel(Fact, BodyList),
 not(holds(Fact, State)). *
```

### Notes:

- *holdsAll* just applies *holds* to every atom in the "body" list
- Can replace both \* with *holds(neg(F),S) :- not(holds(F,S))*.

CSC 384 Lecture Slides (c) 2002, C. Boutilier

16

## Static Facts

- Certain facts never change as actions performed
  - e.g., *adj(X,Y)* in our example is a *static fact*
- No need to carry these around in state list
- Static facts can be represented as derived relations with empty "bodies"

```
derivedRel(adj(o,m), []).
derivedRel(adj(o,l), []). etc...
```

or possibly:

```
derivedRel(adj(X,Y), []) :- neighbor(X,Y).
```

with *neighbor* specified in Prolog KB.

CSC 384 Lecture Slides (c) 2002, C. Boutilier

17

## The CWR-D Meta-Interpreter

- *holds(F,S)* for CWR-D looks a lot more like Prolog. This metainterpreter allows us to assert *facts* in the *state* (thus varying the "fact" part of the KB) and assert *rules* (via derived relations) that do not vary. The *holds* predicate essentially tries to find a proof for the query *F* using the KB *S* (and the rules).
- We can thus have many different "KBs" at once
- For more on meta-interpreters, see Ch.6 of the text (if you're interested).

CSC 384 Lecture Slides (c) 2002, C. Boutilier

18

## Action Representation

- Actions are **concrete** things an agent can do
  - e.g., `move(l,o)`, `tidylab`, `getkeys`, etc.
- Action schema**: a collection of "related" actions
  - e.g., `move(X,Y)` is a schema
  - ground instances of a schema are true actions
- As discussed, actions cause state changes
- A **neighbor** of state *S* is any state *S'* such that an action *A* can be applied at *S* and results in *S'*
- So we need to describe (and represent):
  - when *A* can be applied at *S*; i.e., *A*'s **preconditions**
  - the effect *A* has on state *S*; i.e., *A*'s **effects**

CSC 384 Lecture Slides (c) 2002, C. Boutilier

19

## Precondition Representation

- Precondition for action *A*: a logical condition that must hold at *S* for *A* to be applicable
  - precondition representation: a list of ground literals
  - for an action schema, precondition literals may include variables mentioned in the action schema

| Action (schema)          | Preconditions                                |
|--------------------------|----------------------------------------------|
| <code>move(X,Y)</code>   | <code>loc(X), accessible(Y), adj(X,Y)</code> |
| <code>tidylab</code>     | <code>loc(l), neg(tidy(l))</code>            |
| or: <code>tidy(X)</code> | <code>loc(X), neg(tidy(X))</code>            |
| <code>getkeys</code>     | <code>loc(o), neg(rhk) ...</code>            |

- Assert **precond** predicate (one fact per action)
  - e.g., `precond(move(X,Y), [loc(X), acc(Y), adj(X,Y)])`.

CSC 384 Lecture Slides (c) 2002, C. Boutilier

20

## Effect Representation

- Actions affect only a few world facts, so we're best off describing only **what changes** when an action is performed
  - assume unmentioned facts are unaffected (like CWA)
- Effects are either **positive** or **negative**
  - action makes an atom true or false
- In CWR or CWR-D state rep'n:
  - we should **add** positive effects to the state list
  - we should **delete** negative effects from the state list

CSC 384 Lecture Slides (c) 2002, C. Boutilier

21

## STRIPS Action Representation

- STRIPS representation uses add and delete lists (and precondition lists) to represent an actions
  - `addlist(givecoffee, [ chc, craighappy ])`.
  - `deletelist(givecoffee, [ rhc ])`.
  - `addlist(move(X,Y), [ loc(Y) ])`.
  - `deletelist(move(X,Y), [ loc(X) ])`.
- Important: if you use CWR-D relation, add and delete lists should not include derived relations
  - changes to derived facts due to changes in base facts
- STRIPS: Stanford Res. Inst. (SRI) Planning System (1970-72)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

22

## Domain Descriptions

- A **domain description**: a specification of the actions (and language) required in your specific planning domain
  - for each action: precondition, addlist, deletelist statements should be asserted in KB
  - if CWR-D, state which relations are basic, derived
  - might have a list of all action names (schemas) in KB

CSC 384 Lecture Slides (c) 2002, C. Boutilier

23

## Neighbors (Result of Action)

- Exercises
  - define predicate `results(Action,State,NewState)`
    - given *State* in CWR-D, *Action*, what is resulting state
    - if *Action* preconditions not satisfied at *State*, fails
  - define neighbor predicate for STRIPS/CWR-D
    - `nb(State,NBList) : S'` is a neighbor of *S* if exists an *A'* s.t. preconds of *A'* satisfied at *S*, and `results(A',S,S')`
    - assume `actionList([ move(X,Y), givecof, ...])` in KB

CSC 384 Lecture Slides (c) 2002, C. Boutilier

24

## Planning Problems

- Given a domain description
  - preconditions, addlist, deletelist, actionlist, derivedRel, baseRel, all specified in KB
- Given an initial state  $s_0$  in CWR or CWR-D
  - initial state is just a list of positive ground atoms
- Given a goal set  $G$ 
  - $G$  is a list of (post'v or negt'v) literals to be made true
- Find a sequence of actions (a **plan**)  $a_1, a_2, \dots, a_n$  s.t. applying that sequence to  $s_0$  leads to a state  $s_n$  satisfying all goal literals
  - $results(a_1, s_0, s_1), results(a_2, s_1, s_2), \dots, results(a_n, s_{n-1}, s_n)$  and  $holdsall(G, s_n)$  are all true

CSC 384 Lecture Slides (c) 2002, C. Boutilier

25

## Example Planning Problem

- Initial State (CWR-D):  $[loc(o), lck(c)]$ 
  - unmentioned:  $neg(cm), neg(rhk), neg(chc), \dots$
  - unmentioned:  $adj(o, l), adj(o, c), neg(adj(l, c)), accessible(l), neg(accessible(c)), \dots$
- Goal:  $[loc(o), chc]$ 
  - note: this is **not a state** in CWR-D, it is just a set of conditions we want true

CSC 384 Lecture Slides (c) 2002, C. Boutilier

26

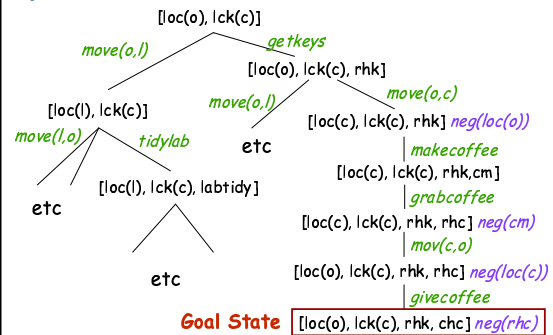
## Planning as Search

- Given this formulation of neighbors, goals, it is easy to see how planning can be solved using any standard search algorithm
  - initial state  $s_0$  is root of search tree
  - for any  $s$  in the tree, its children (neighbors) are determined by  $nb$  (states reachable using one action)
  - $is\_goal(S)$  determined by  $holdsAll(G, S)$
- Your favorite search algorithm builds a sequence of states from start to (some) goal state
  - note: you don't want state sequence, but an *action* sequence (so modify  $nb$  predicate, search algorithm to extract actions you took to reach these states)

CSC 384 Lecture Slides (c) 2002, C. Boutilier

27

## Partial Search Tree in Example



CSC 384 Lecture Slides (c) 2002, C. Boutilier

28

## Difficulties with Generic Search

- Search tree is generally quite large
  - our branching factor roughly 3; solution depth is 7
  - BFS: about 3300 node expansions
  - DFS is easily lead astray in this problem
- Goal provides no guidance
  - path selection uninfluenced by goal (e.g., why even consider going to lab as the first step?)
  - suggests we need heuristics (people do work on this)
- Can we get by without heuristics?
  - heuristics generally handcrafted for *specific* domains
  - want a *general purpose planner* that is directed toward the goal somehow

CSC 384 Lecture Slides (c) 2002, C. Boutilier

29

## Goal-Influenced Planning

- In our example, intuitively the action *tidylab* is irrelevant to the literals in the goal list
  - it doesn't achieve a goal literal, it doesn't achieve the precondition of any action that achieves a goal literal, etc...
- Most planning algorithms work backward from the goal list. The basic idea/intuition:
  - find an action  $a_n$  that achieves a goal literal
  - then find an action  $a_{n-1}$  that achieves one of  $a_n$ 's preconditions, etc...
- Focuses attention on goal literals or things that need to be made true to achieve goal literals
  - actions that have some (indirect) impact on the goal

CSC 384 Lecture Slides (c) 2002, C. Boutilier

30

## STRIPS Planner

- STRIPS one of earliest AI planning algorithms
  - note: distinguish STRIPS action repn (a good idea) from STRIPS planning algorithm (not so great)
  - focuses on actions relevant to your goals (so more informed than blind search)
- Basically a divide-and-conquer approach to "solving a goal list"
  - tries to find independent plans for individual subgoals and then pieces these plans together
  - recursively tries to achieve necessary preconditions

CSC 384 Lecture Slides (c) 2002, C. Boutilier

31

## STRIPS: Intuitive Sketch

- Given a goal list  $[g_1, g_2]$ , initial state  $s_0$ 
  - Select a goal to achieve, let's say  $g_1$
  - Find a sequence of actions that achieves  $g_1$  from  $s_0$ , let's say  $a_1, a_2, a_3$
  - Compute state  $s_1$  that results from applying  $a_1, a_2, a_3$  to  $s_0$ :  $s_1 = \text{result}(a_3, \text{result}(a_2, \text{result}(a_1, s_0)))$
  - Select a remaining goal (here only  $g_2$  remains)
  - Find a sequence that achieves  $g_2$  from  $s_1$ , let's say  $b_1, b_2, b_3$
  - Return solution: Plan  $P = a_1, a_2, a_3, b_1, b_2, b_3$

CSC 384 Lecture Slides (c) 2002, C. Boutilier

32

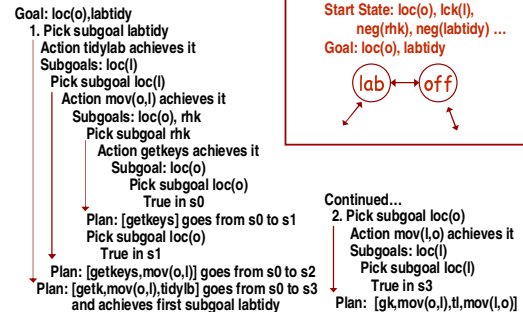
## How to Solve a Subgoal

- How to find sequence to achieve  $g_1$  from  $s_0$ ?
  - find an action (say,  $a_3$ ) that achieves  $g_1$  (so  $g_1$  must be an effect of  $a_3$ )
  - to ensure we can execute  $a_3$ , we make all of its preconditions new goals, or *subgoals* (say  $p_1, p_2$ )
  - recursively call STRIPS on subgoals  $[p_1, p_2]$  with start state  $s_0$ , to get sequence  $a_1, a_2$  that achieves them
  - sticking  $a_3$  on the end of  $a_1, a_2$  ensures achievement of  $g_1$
- Process terminates when all subgoals true at  $s_0$

CSC 384 Lecture Slides (c) 2002, C. Boutilier

33

## A Simple Example of STRIPS



CSC 384 Lecture Slides (c) 2002, C. Boutilier

34