

Planning Agents

Sven Koenig, USC

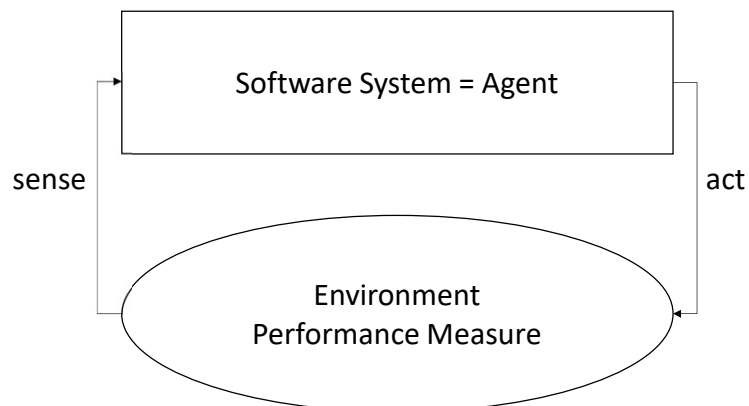
Russell and Norvig, 3rd Edition, Sections 2.4 and 3.1-3.2

These slides are new and can contain mistakes and typos.
Please report them to Sven (skoenig@usc.edu).

1

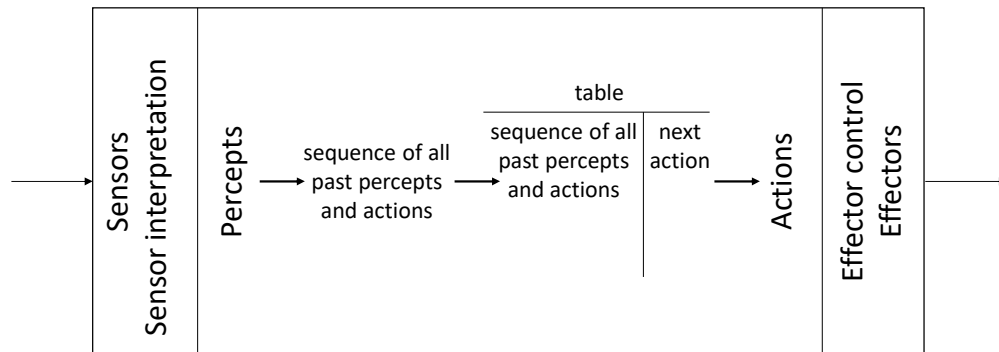
Search and Planning

- We now start with (deterministic) search and planning.



2

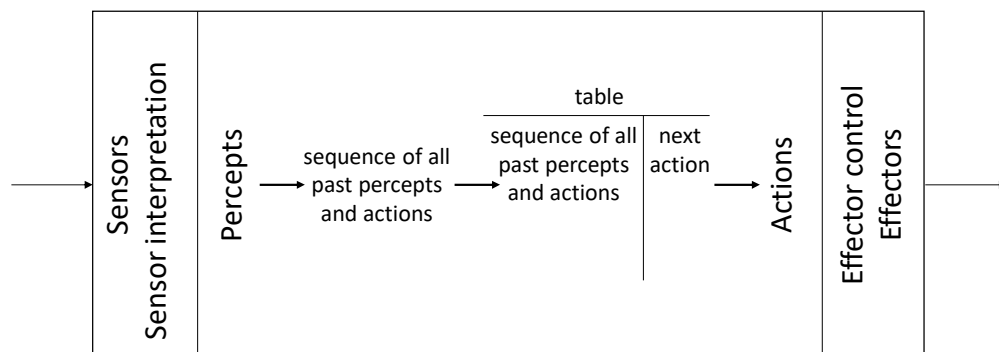
Architectures for Planning Agents



3


Architectures for Planning Agents

- Gold standard (but: results in a large table that is difficult to change)



4


Architectures for Planning Agents



5

Architectures for Planning Agents

- Reflex agent (“reactive planning”): often good for video games



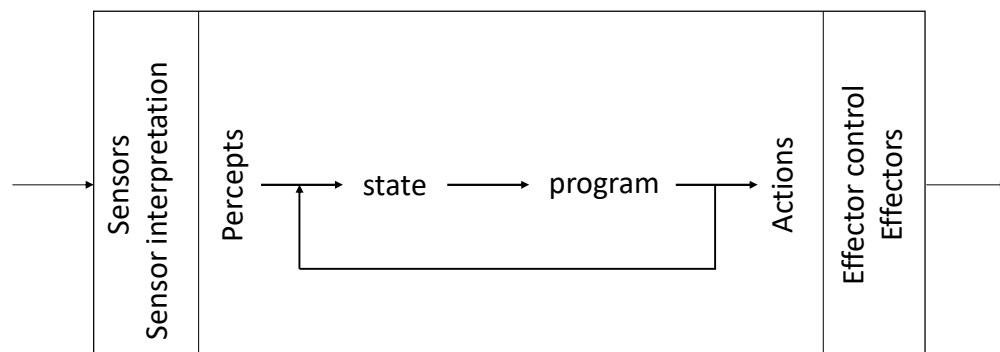
6

Architectures for Planning Agents

- An agent often does not need to remember the sequence of past percepts and actions to perform well according to its performance objective.
- A state characterizes the information that an agent needs to have about the past and present to pick actions in the future to perform well according to its performance objective.
- We are typically interested in minimal states.
- For example, a soda machine does not need to remember in which order coins were inserted and in which order coins and products were returned in the past. It only needs to remember the total amount of money inserted by the current customer.

7

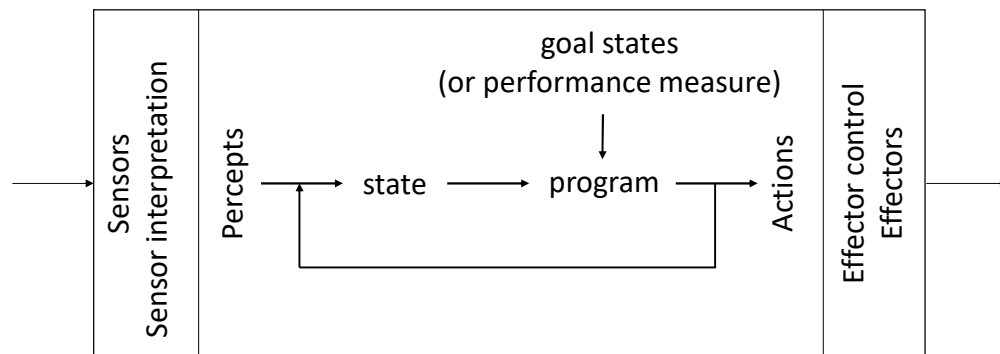
Architectures for Planning Agents



8

Architectures for Planning Agents

- Planning agent



9

Examples

- What are the states, actions and action costs?
- Eight puzzle

1	3	2
5	6	
7	8	4

start (= current) configuration

1	2	3
4	5	6
7	8	

goal configuration

10

Examples

- What are the states, actions and action costs?
- Missionaries and cannibals problem

Three missionaries and three cannibals are on the left side of a river, along with a boat that can hold one or two people. Find the quickest way to get everyone to the other side, without ever leaving a group of missionaries in one place outnumbered by the cannibals in that place.

11

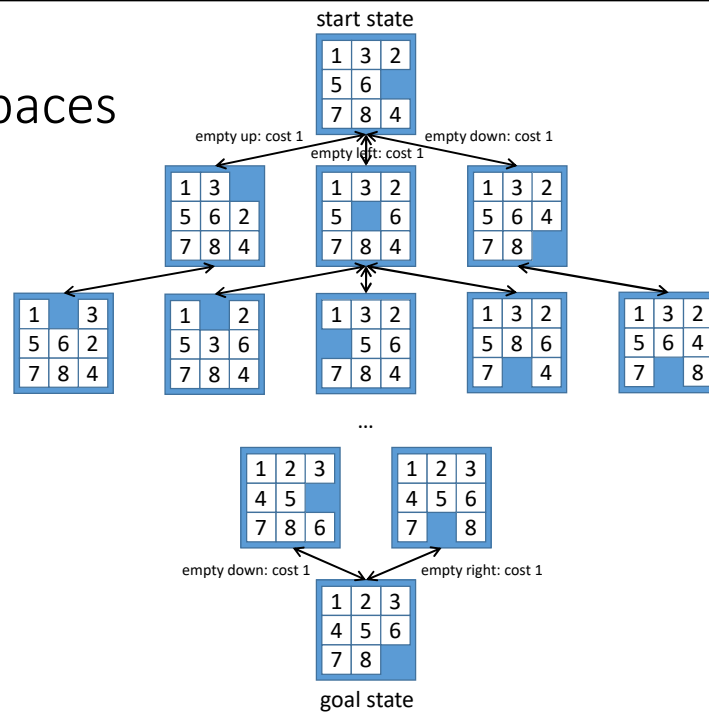
Examples

- What are the states, actions and action costs?
- Traveling salesperson problem

Visit all given cities in the plane with a shortest tour (= with the smallest round-trip distance).

12

State Spaces




13

State Spaces

• Example application of hillclimbing: Boolean satisfiability

- Find an interpretation that makes a given propositional sentence true.
- Transform the propositional sentence into conjunctive normal form, assign random truth values to all propositional symbols, then repeatedly switch the truth value of some symbol to decrease the number of clauses that evaluate to false.
- $S \equiv (P \text{ OR } Q) \text{ AND } (\text{NOT } P \text{ OR } \text{NOT } R) \text{ AND } (P \text{ OR } \text{NOT } Q \text{ OR } R)$

start state = assignment of random truth values to all propositional symbols



- Costs do not matter since we are not interested in finding a minimum-cost path.
- There is more than one goal state, e.g. P, Q, NOT R and P, NOT Q, NOT R.
- There is a goal test, namely whether S is true.

14

State Spaces

- Graph
- Vertex
- Edge
- Edge cost
- Start vertex
- Goal vertex
- State space
- State
- Action = operator = successor function $\text{succ}(s,s') \in \text{States}$
- Action cost = operator cost
- Start state
- Goal state or goal test $\text{goal}(s) \in \{\text{true}, \text{false}\}$
- Solution is a (minimum-cost) path from the start vertex to any goal vertex
- Solution is a (minimum-cost) action sequence = operator sequence from the start state to any goal state