Artificial Intelligence (CS6380)

Agents and Environments

Agent

Agent : anything that acts within an **environment** to achieve a particular **goal**

- Agent **perceives** env.
- Makes decisions about actions.
- Actions change env.
- Repeat..

AI : algorithmic paradigm of making these decisions intelligent

Agent

Agent : anything that perceives the environment through its sensors and acts upon the environment using the actuators.

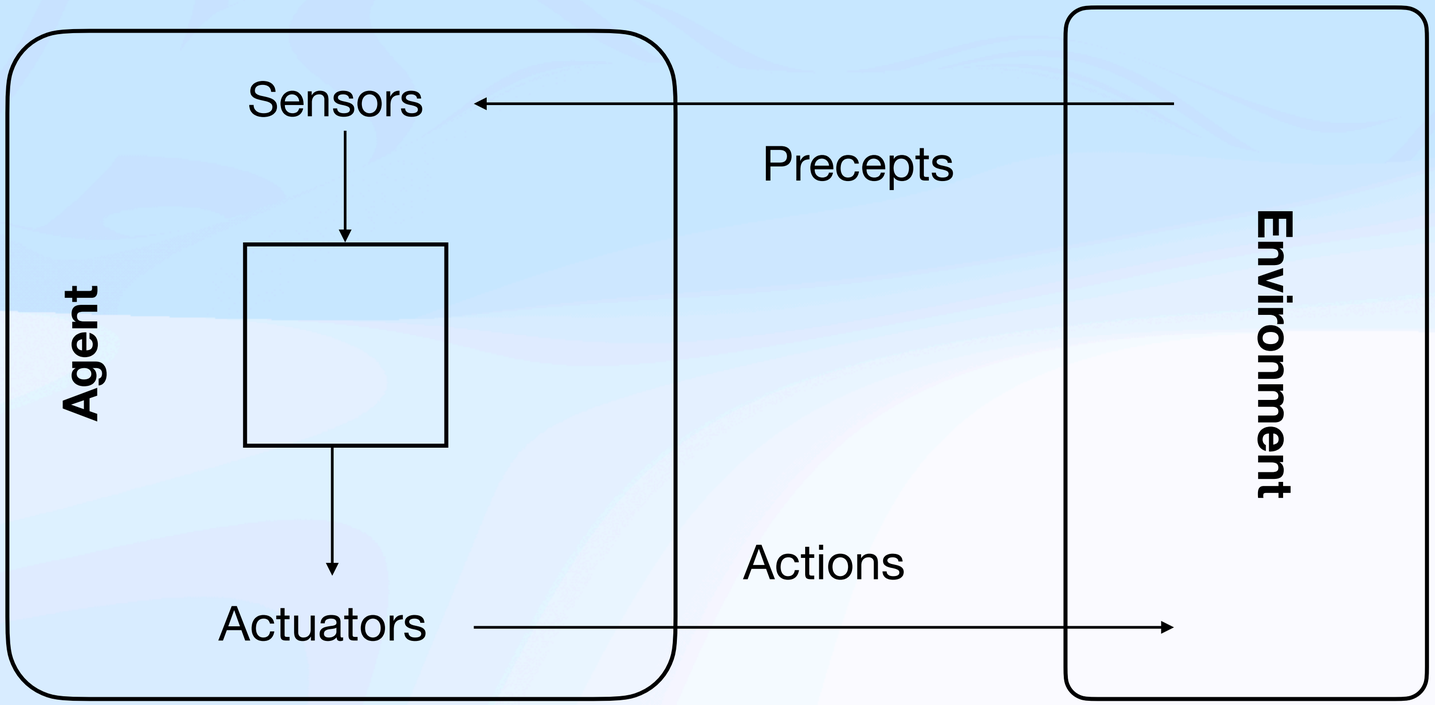

Human agent:

- **Sensors**: eyes, ears, ..
- **Actuators**: hands, legs, ..

Software agent:

- **Sensors**: receives files, human input (keyboard, mouse) ..
- **Actuators**: display on screen, sounds generated ..

Robotic agent:

- **Sensors**: cameras, infrared range finders, ..
- **Actuators**: motors, ..

Agent may have knowledge about the environment.

Agent: precepts and actions

Percept:

- Agents perceptual inputs at any instant
- Agents current belief of the world.
- Environments current state.

Percept sequence:

- Complete history of what the agent has perceived so far.

Agent function:

- Input: percept sequence
- Output: action

Agent function may depend on

- Entire percept sequence
- Only current state

Vacuum cleaner world

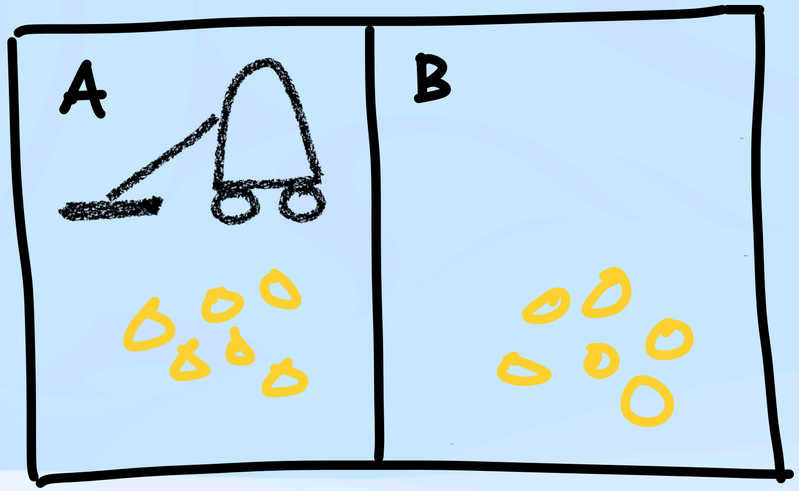

Percept sequence	Action
[A, clean]	Right
[A, dirty]	Suck
[B, clean]	Left
[A, clean], [A, clean]	Right

Agent: vacuum cleaner
Env: 2 cells A and B (can be clean or dirty)
Actions: left, right, suck

Legal vs illegal actions

- **Moving right in cell B**
- **Suck when clean**

What should the agent do when it perceives that cell A is clean?

Good behaviour: Rationality

Rational agent is supposed to do the right thing.

What is the right thing?

Consequentialism:

- Actions cause a sequence of states.
- Is this sequence desirable?
  - [A, dirty] right [B, dirty] left

Performance measure:

- Evaluates a given sequence of env. states.
- Given by the designer of the agent.

Vacuum cleaner:

- Amount of dirt sucked in a time window?
- One point for clean floor, penalty for elec. used, noise generated

Rational agent

Rational agent is supposed to do the **right thing.**

Rationality depends on:

- Agents percept sequence (S)
- Agents prior knowledge about env. (E)
- Actions that the agent can perform (A)
- Performance measure (P)

Vacuum cleaner world

- Loc. senor and dirt sensor
- Geography of loc. is known. What does left do?
- Right, left, suck
- One point for each clean square over fixed time steps

For each percept sequence, rational agent selects an action that is *expected to maximize* its performance measure, given the evidence provided by the percept sequence and agents built in knowledge.

PEAS model for problems

Rational agent contd.

For each percept sequence, rational agent selects an action that is *expected to maximize* its performance measure, given the evidence provided by the percept sequence and agents built in knowledge.

Rationality vs. perfection

- ***What if one of the cells catches fire*?**

Information gathering.

- ***What if there are more cells*?**

Learning

- ***Which cells get dirty, how frequently*?**

Env. and agents : types

Types of Env.

- Fully vs partially observable
- Single vs multiagent
- Deterministic vs stochastic
- Static vs dynamic
- Discrete vs continuous
- Episodic vs sequential

Types of Agents

- Table driven agent
- Simple reflex agent
- Model based reflex agent
- Goal based agent
- Utility based agent
- Learning agent

Table driven agent

Function: Table-Driven-Agent (input-percept)

Persistent: precepts, a sequence initially empty
table of actions, indexed by percept sequence
fully specified.

Append input-percept to the end of percepts.
Action = Lookup (percepts, table)
Return Action

$\mathcal{P}$: Set of precepts

$\mathcal{T}$: Total number of times agent receives the percepts.

What is the size of the table?

Key challenge: write a small program instead of a vast table that produces rational

Reflex agent

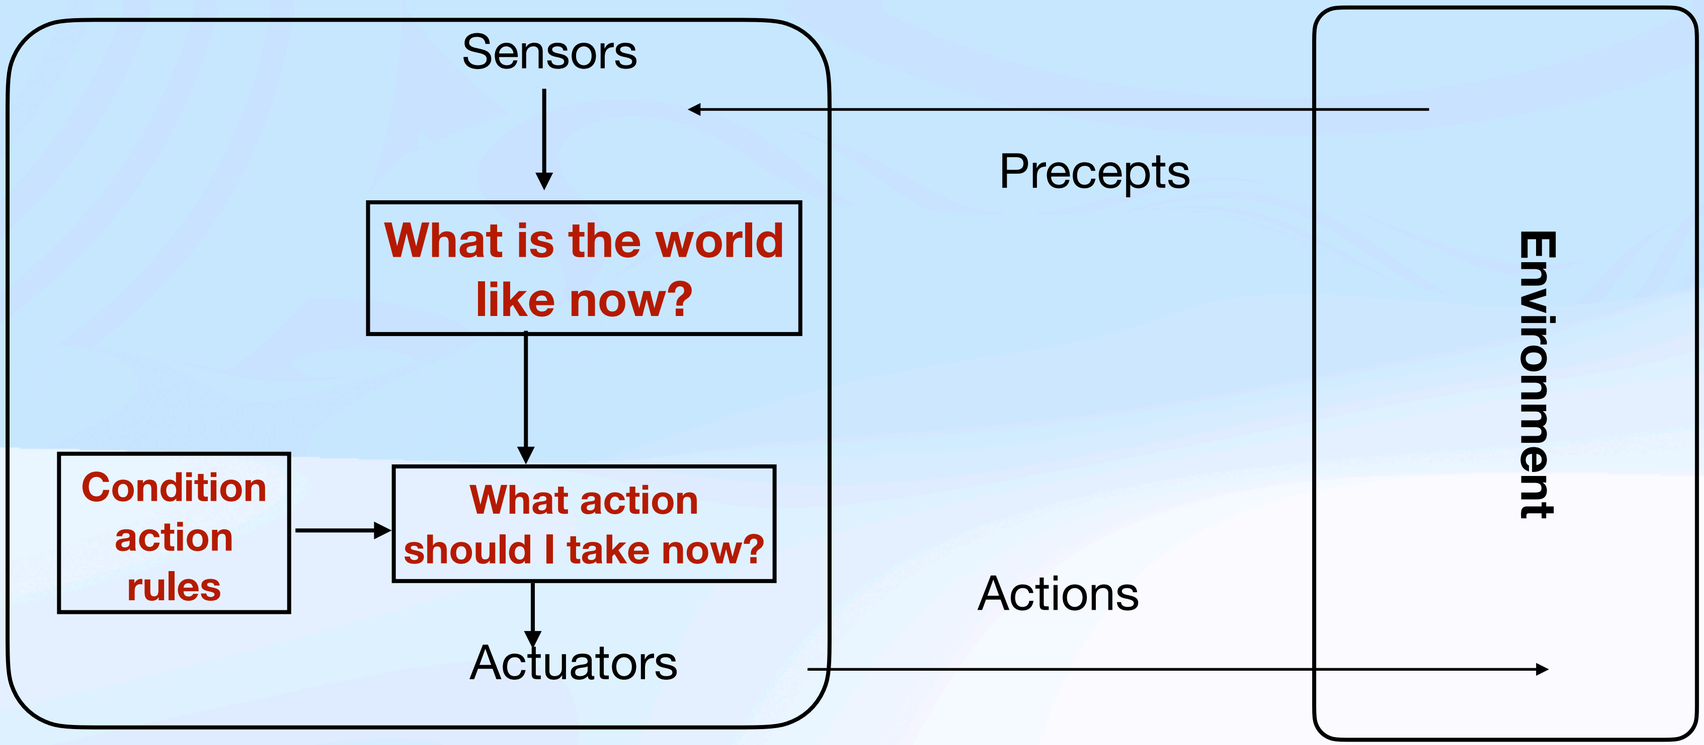

	Loc A	Loc B
Dirty	Suck	Suck
Clean	Right	Left

What if agent only one percept?

Model based Reflex agent

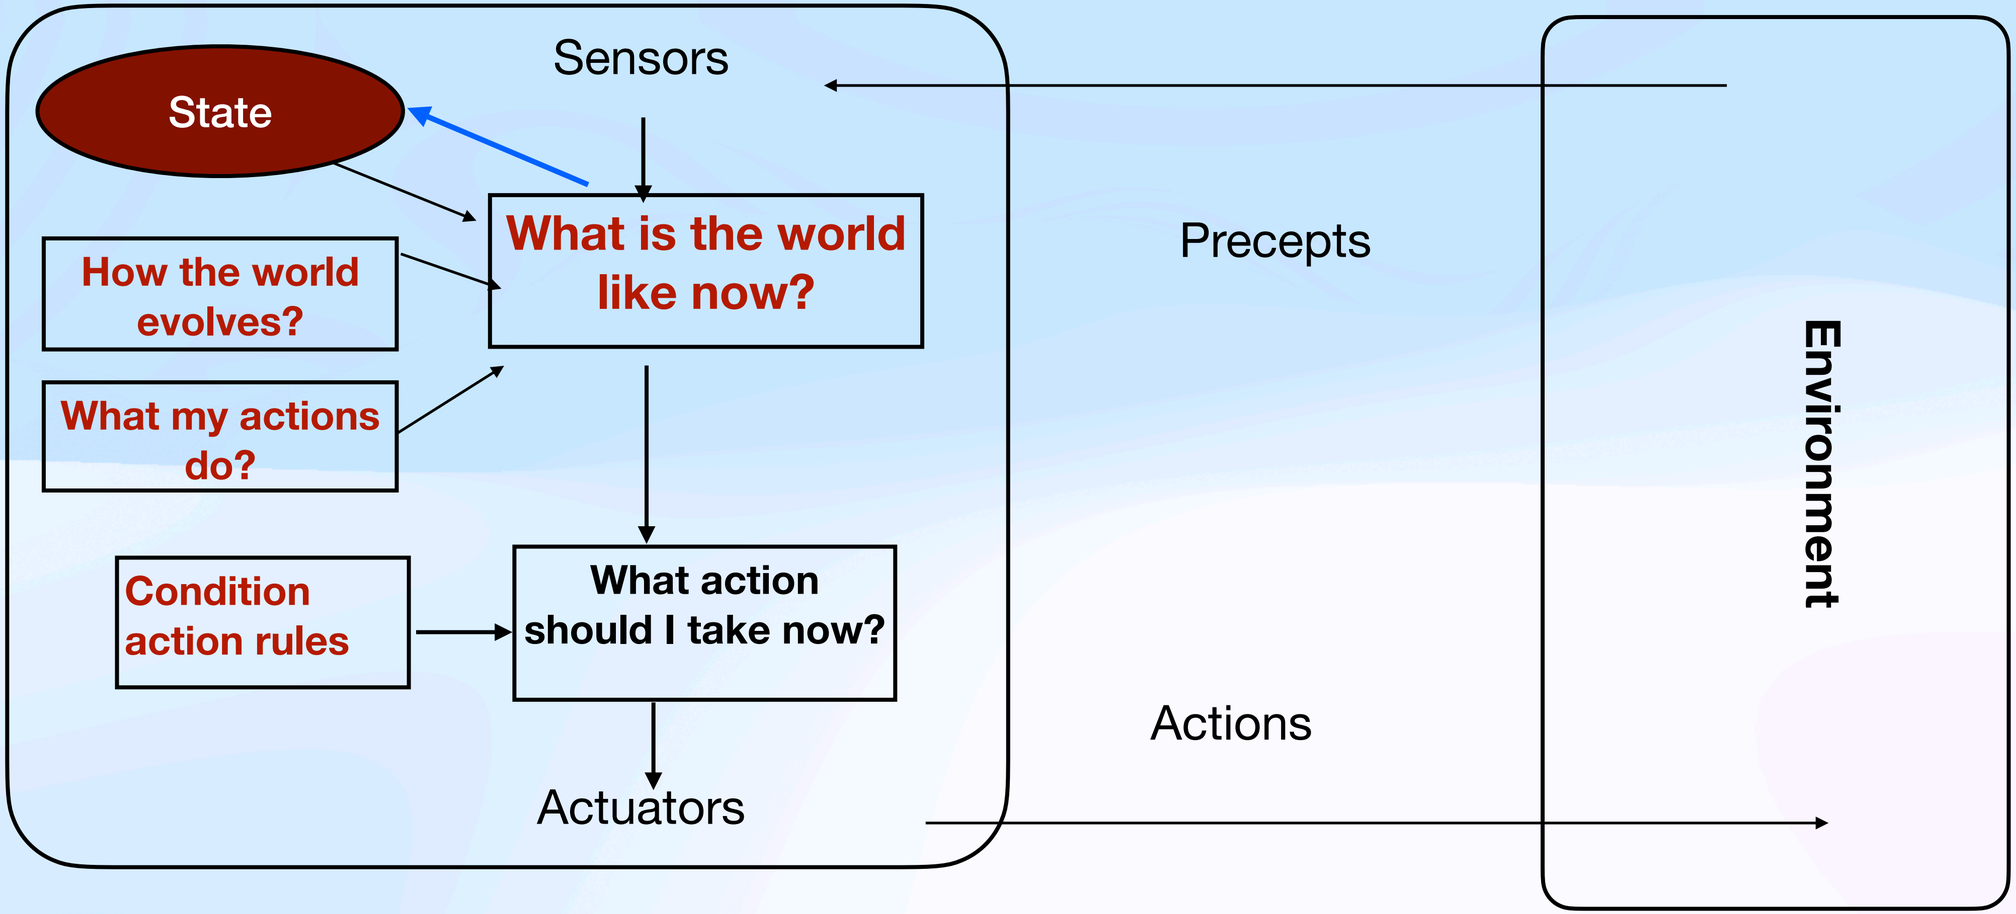

Model based goal-based agent

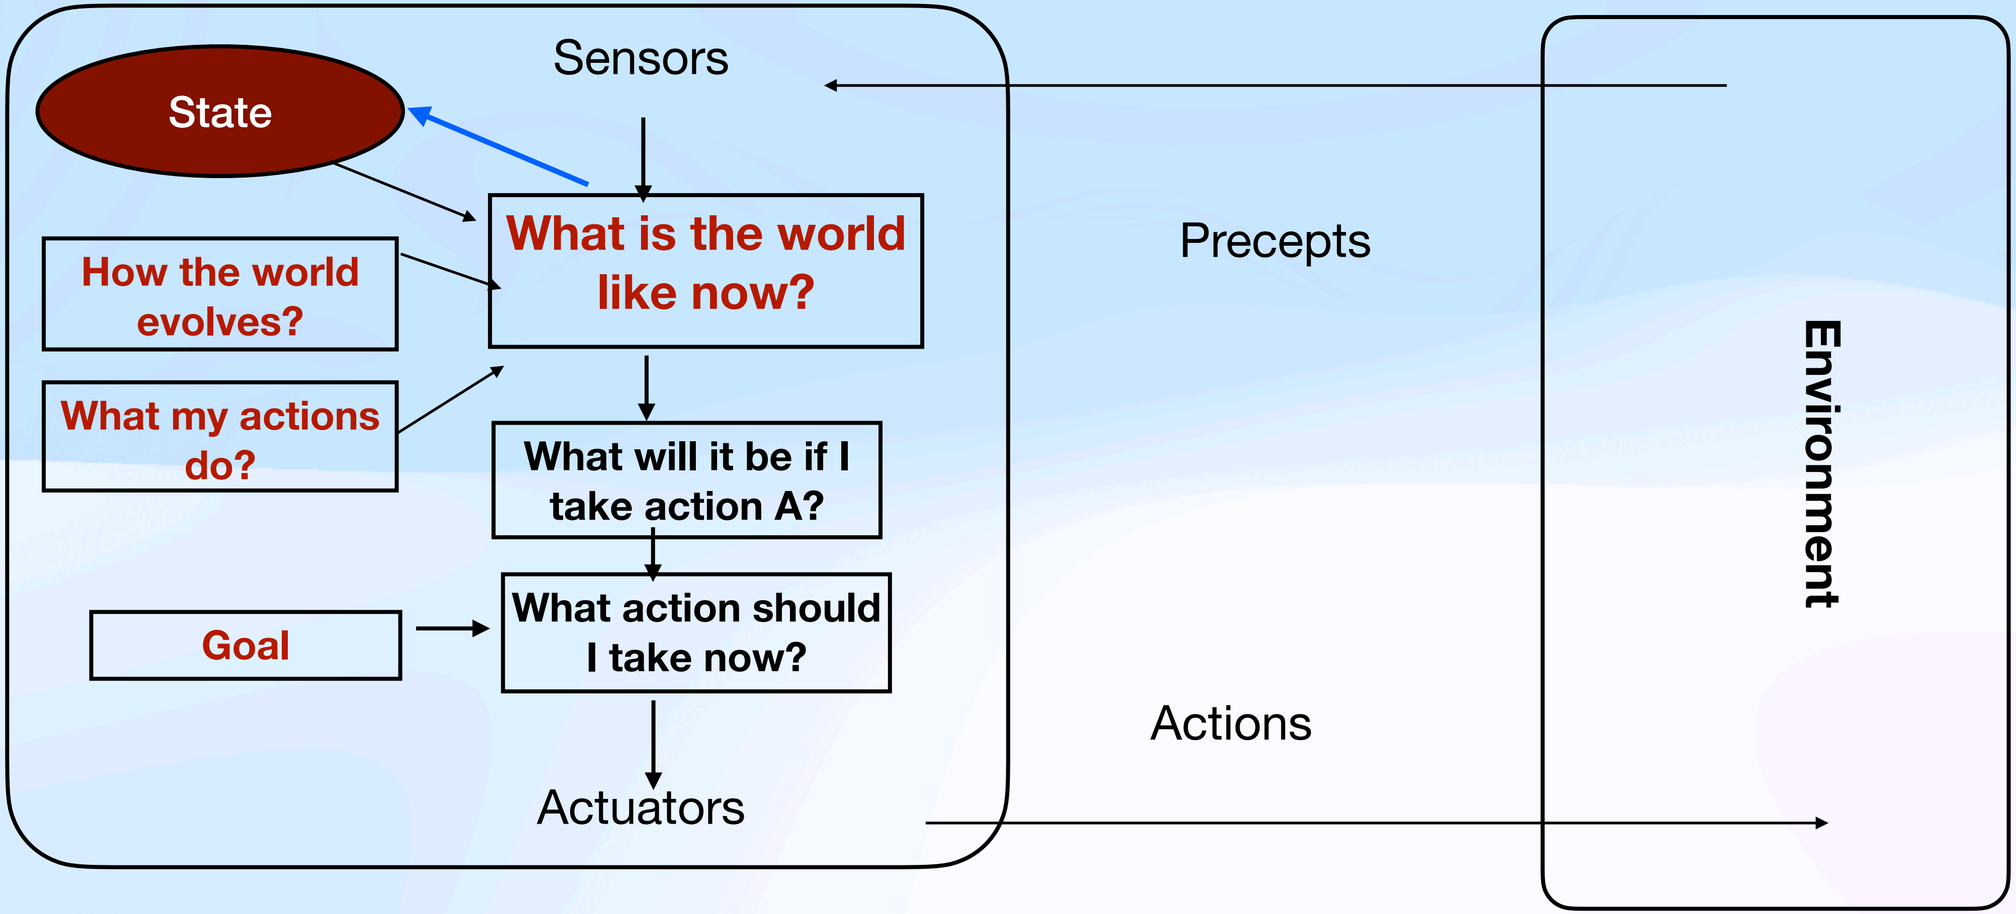

Goal based agent

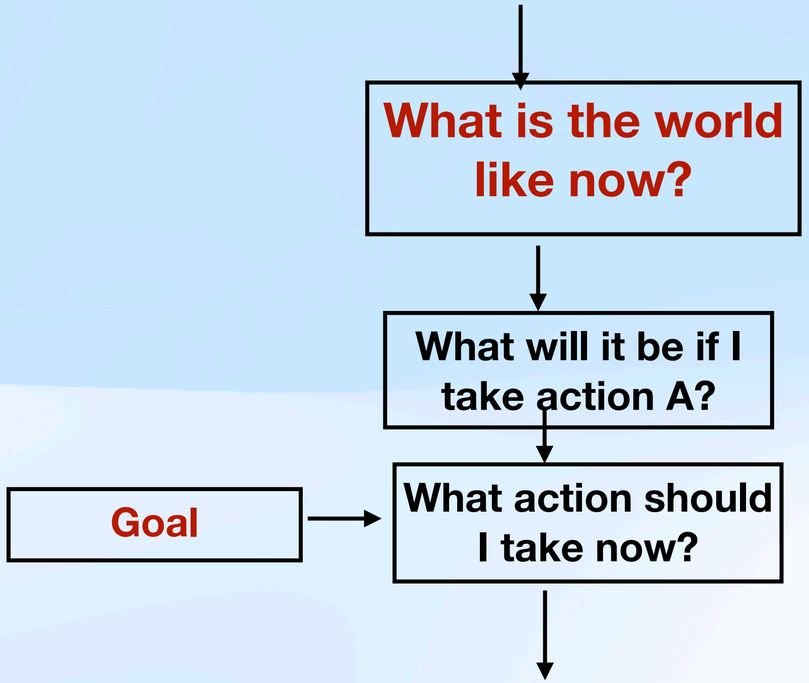

What if achieving the goal is not via a single action but requires a **sequence of actions?**

Search techniques

Planning

Vacuum world

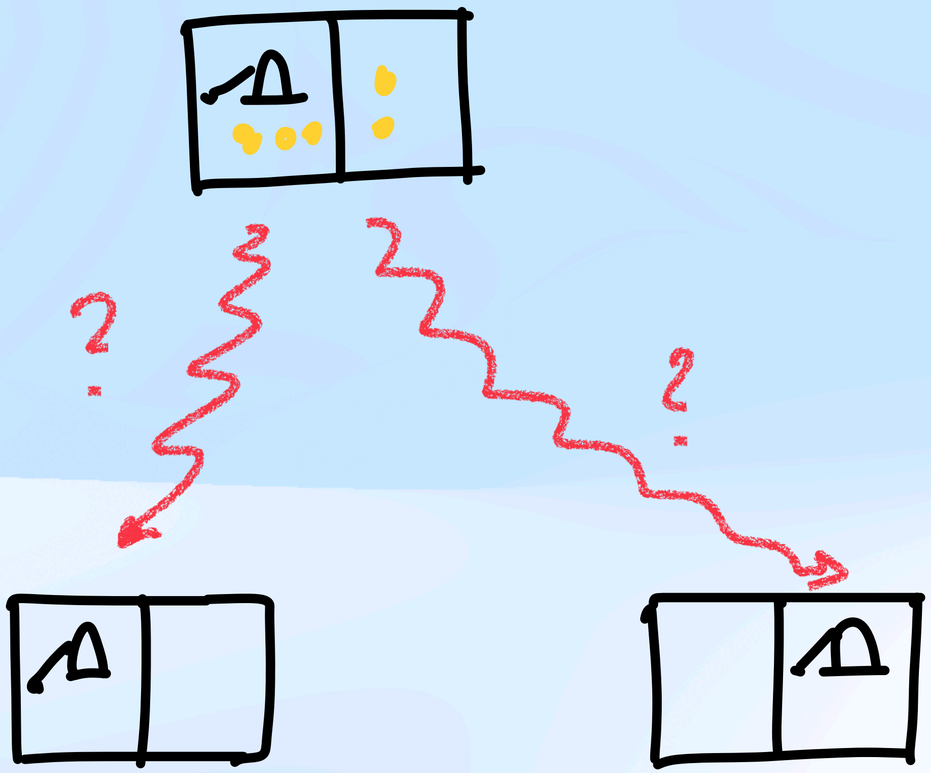

Model it as a state space search problem

- Set of states
- Initial state
- Goal state(s) or isGoal function
- Actions
- Transition model
- Action cost function

Vacuum world

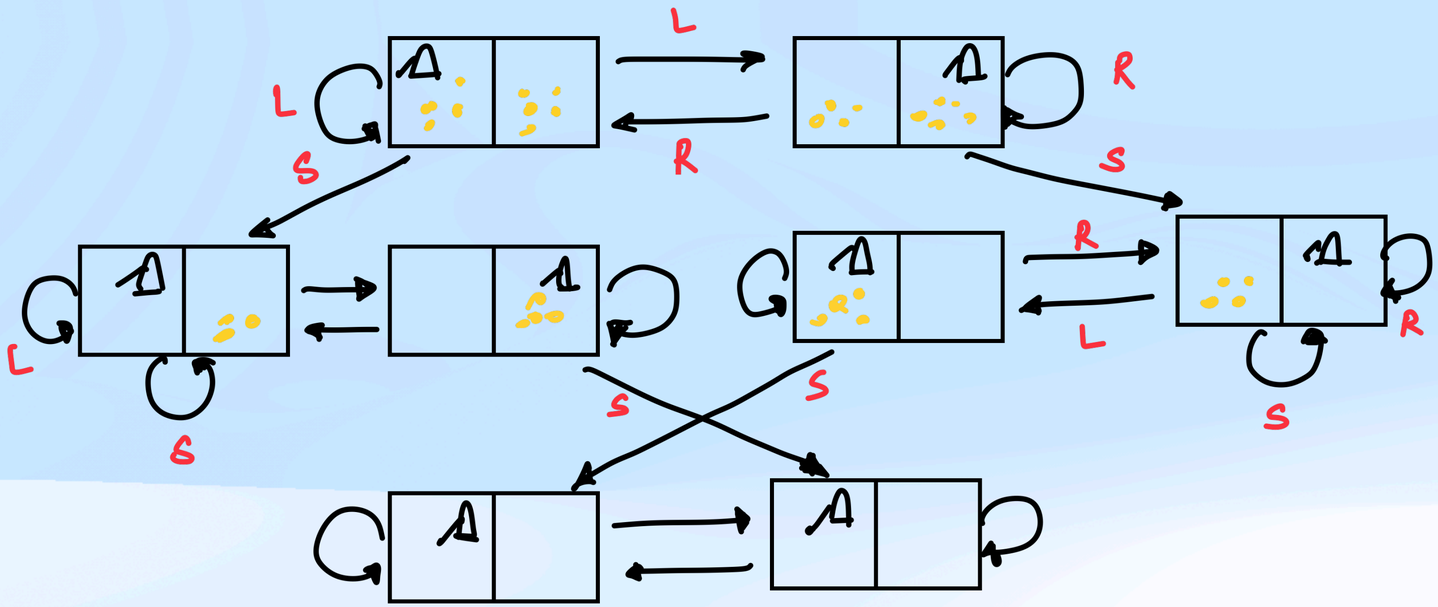

Discussion points from class

- For an example with n cells, how many states are there in the state space?
- In a 2 cell world, if we start with initial state where both cells are clean, does the agent need to go to other cell to know that it is clean?
  - No. We do not insist that the agent to knows that the two cells are clean. Agent knowledge versus the state being a goal state

Notion of a state

State: All information about the environment

Is all information relevant?

State: All information necessary to make a decision to solve the problem

Representation of states

Type	State representation
Atomic	States are indivisible, no internal structure
Factored	States are made up of variables that take values
Structured	States have variables and objects and their relations