

Artificial Intelligence

Topic 2


Intelligent Agents

- ◇ What are agents?
- ◇ Rational agents
- ◇ Agent functions and programs
- ◇ Types of agents

Reading: Russell and Norvig, Chapter 2

1. What are agents?

agent — *perceives* environment through *sensors*
— *acts* on the environment through *effectors*



Examples...

Percepts	light, sound, solidity, ...
Sensors human robot	eyes, ears, skin, ... infra-red detectors, cameras, microphone, accelerometers, ...
Effectors human robot	hands, legs, voice, ... grippers, wheels, speakers, ...
Actions	pickup, throw, speak, ...

2. Rational agents

Recall: rational agent tries to...

- “do the right thing”
- act to achieve goals

The “right thing” can be specified by a *performance measure* defining a numerical value for any environment history

Rational action: whichever action maximizes the expected value of the performance measure given the percept sequence to date

Rational \neq omniscient

Rational \neq clairvoyant

Rational \neq successful

3. Agent functions and programs

An agent can be completely specified by an *agent function* mapping percept sequences to actions

(In principle, one can supply each possible sequence to see what it does. Obviously, a lookup table would usually be immense.)

One agent function (or a small equivalence class) is rational

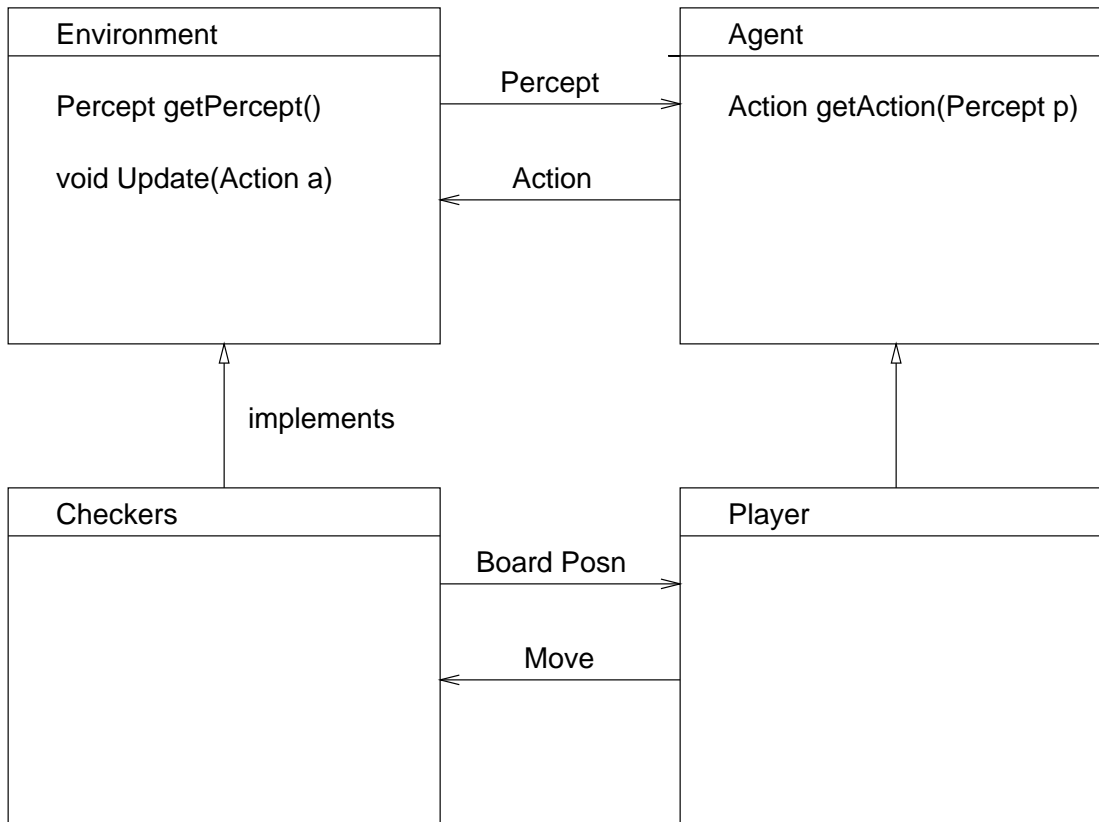
Aim: find a way to *implement* the rational agent function *concisely* and *efficiently*

An *agent program* implements an agent function: takes a single percept as input, keeps internal state, returns an action:

```
function SKELETON-AGENT(percept) returns action
  static: memory, the agent's memory of the world
  memory ← UPDATE-MEMORY(memory, percept)
  action ← CHOOSE-BEST-ACTION(memory)
  memory ← UPDATE-MEMORY(memory, action)
  return action
```

In OO-speak ... ⇒

3.1 Agents in Java



```
class Player implements Agent {
    ...

    Action getAction(Percept percept) {
        Move myNextMove;
        // my selection algorithm
        return myNextMove;
    }
    ...
}
```

4. Types of Agents

An *agent program* accepts percepts, combines them with any stored knowledge, and selects actions.

A *rational agent* will choose actions so as to maximise some *performance measure*. (In practice try to achieve “good” performance.)

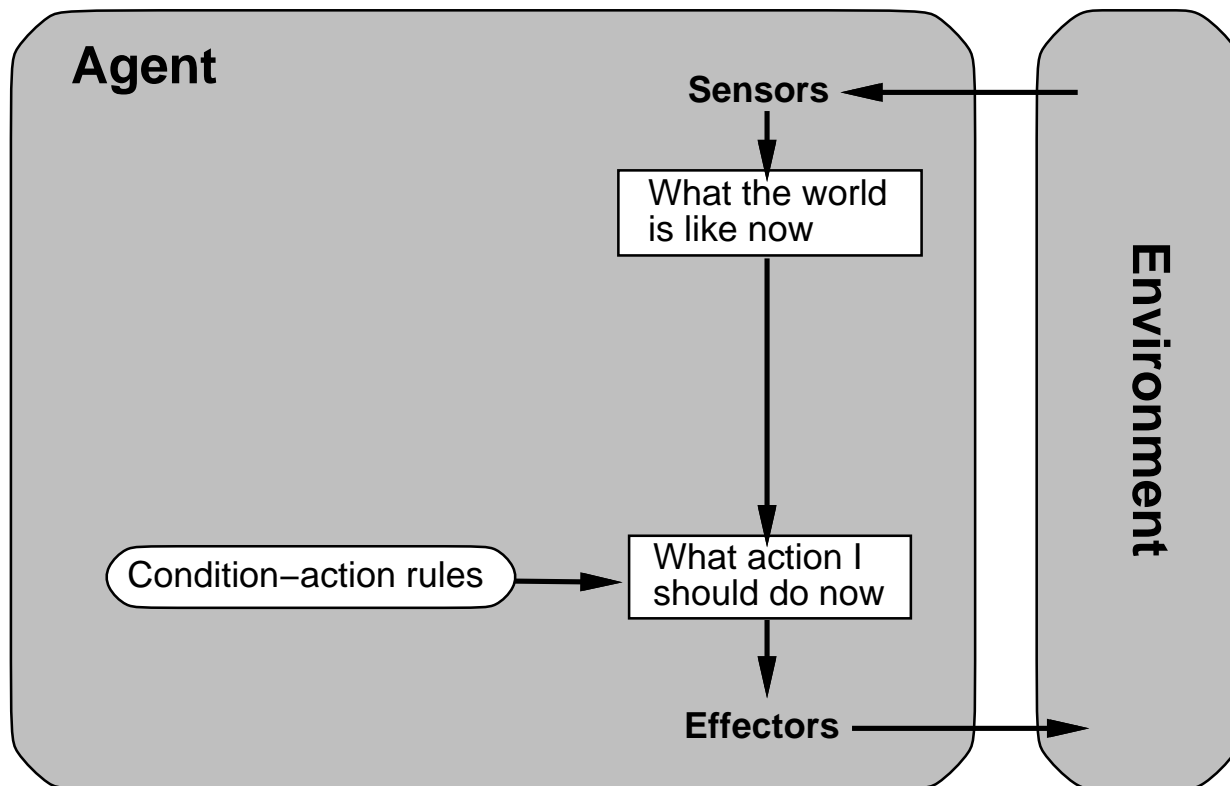
Four basic types in order of increasing generality:

- simple reflex agents
- reflex agents with state
- goal-based agents
- utility-based agents

4.1 Simple Reflex Agents

⇒ choose responses using *condition-action* rules (or *production rules*).

if you see the car in front's brake lights then apply the brakes



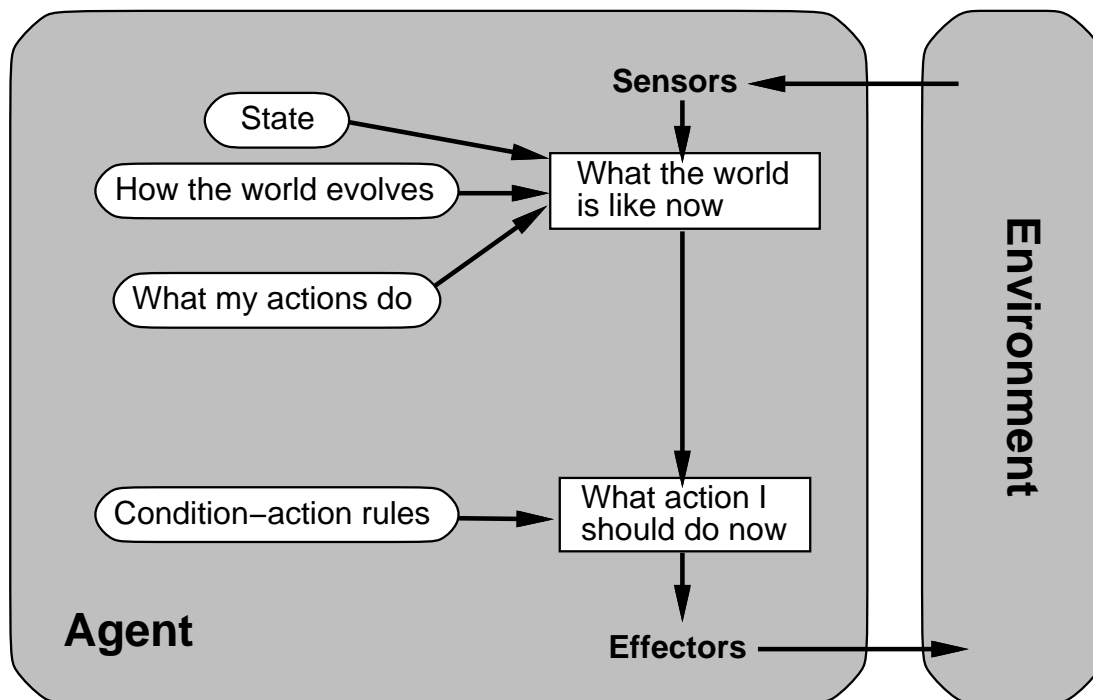
Some researches claim this is how simple life-forms (eg. insects) behave.

4.2 Agents that keep track of the world

While simply reacting to the current state of the world is adequate in some circumstances, most intelligent action requires more knowledge to work from:

- memory of the past
- knowledge about the effects of actions — how the world evolves
- requires internal *state*

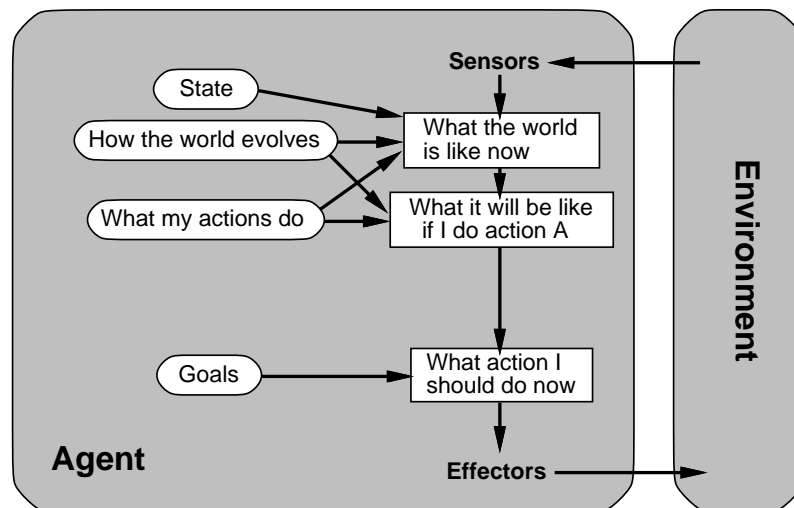
eg. You notice someone ahead signal to the bus. You know that this will cause the bus driver to stop (ideally), and conclude that you should change lanes.



4.3 Goal-based agents

Reacting to evolving world may keep you from crashing your car, but it doesn't tell you where to go!

Intelligent beings act in such a way as to try to achieve some *goals*.



Some goals are pretty simple. eg. Star Trek — “*to go where no-one has gone before*”

Some are more complex and require *planning* to achieve them. eg. Star Wars — *to defeat the Empire* — find a potential Jedi knight, ship him off to see Yoda, teach him to use the force, etc

Planning is a fundamental problem of AI. Usually requires *search* through the available actions to find an appropriate sequence.

Some people even say *AI is really search!*

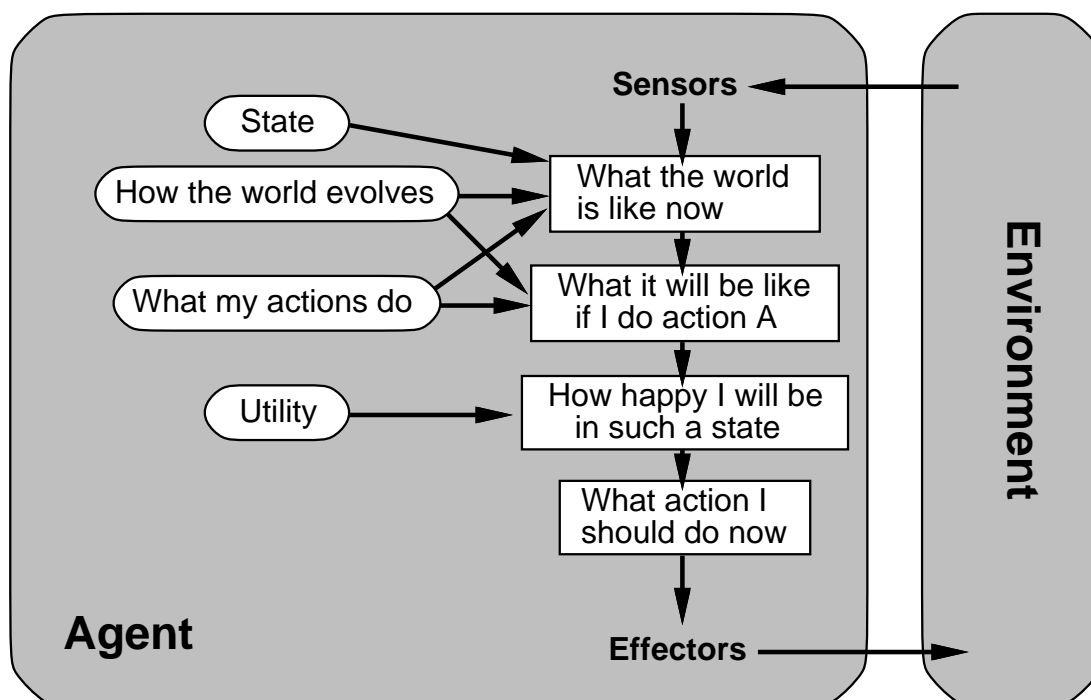
4.4 Utility-based agents

Two views:

1. a successful agent is only required to satisfy objective goals (emotions are a hindrance)
2. subjective measures such as happiness, security (safety), etc are important for success

eg. Luke could maximise his material benefits by turning to the dark side of the force, but he'd be very unhappy :-)

Researchers call this measure of happiness *utility* (since it sounds more scientific)



5. Next...

- Conceptualising the environment
 - states, actions (operators), goals
- The fundamental skills of an intelligent agent
 - problem solving and search

The End