

# Agents, Environments, and Goals

Traditional AI begins with some simple premises:

- An intelligent **agent** lives in a particular **environment**.
- An intelligent agent has **goals** that it want to achieve.

Where do an agent's goals come from?

- automatic, just a part of being an agent
  - hard-wired into agent's body by its programmer
  - "hard-wired" into agent's body by genetics

a goal of self-preservation, which leads the agent to seek sustenance and shelter

- elective, of the agent's own volition

e.g., doing a crossword puzzle, finding a lost ticket

# Satisfying Goals Through Reflex

An agent can satisfy its goal with a very simple algorithm:

1. Sense the world.
2. Act.

This approach assumes that the agent can react to the world based just on what it senses from its environment.

## Quick Exercise:

- Is reflex ever the intelligent way to achieve a goal? If yes, identify a situation where it is. If not, why?
- Identify a situation where reflex is not the most intelligent way to proceed.

Reflex is often the right way to achieve a goal.

Sometimes it fails because the world is complex.  
Sometimes it fails because the cost of failure is too high.

# Satisfying Goals Through Deliberation

An agent can satisfy its goal with a slightly more complex algorithm:

1. Sense the world.
2. **Choose an action.**
3. Act.

Choosing an action can be almost as simple as reflex:

Look up the right action in a table.

Choosing an action can also be arbitrarily complex: An agent can plan ahead before choosing its action.

# The Costs of Deliberation

But planning ahead has its own costs.

1. The agent must be able to imagine the effect of each action that it can take. How will the world change if I do X?
2. The agent must be able to keep in its mind all the possibilities that it's considering while choosing its action.

Another of AI's basic premises:

**Agents are resource-limited.**

The size of an agent's memory is bounded.

The amount of time available is usually bounded.

# A Random Action Selector

INPUT: a goal to achieve

OUTPUT: an action to take

- STEPS
1. Choose one of the available actions.
  2. If the action achieves the goal, then stop. Otherwise, go to Step 1.

This algorithm is called **Generate and Test**.

## Quick Exercise:

- Is Generate and Test ever the intelligent way to achieve a goal? If yes, identify a situation where it is. If not, why?
- Identify a situation where Generate and Test is not the most intelligent way to proceed.

Generate and Test may be the only way to proceed. If the agent doesn't know what the effect of an action will be, then it can't plan ahead based on that action.

But this isn't a *systematic* way to proceed.

# An Algorithm for Systematic Search

INPUT:       the starting situation       (the **start state**)  
              a goal to achieve  
              a search strategy

OUTPUT:     a sequence of actions (called **operators**)  
              that transforms the problem's initial state  
              into a goal state OR an announcement that  
              no such sequence can be found

## STEPS

1. Initialize the set of states to be considered to include the start state.
2. Repeat the following:
  - a. If the set of unconsidered states is empty, then announce failure.
  - b. *Choose a state* to consider, based on the search strategy.
  - c. If the state chosen is the goal state, then return the sequence of actions that leads to this state.
  - d. Add to the set of unconsidered states all of the states that can be reached from the current state by doing one action.