6.825 Techniques in Artificial Intelligence

What is Artificial Intelligence (AI)?

- Computational models of human behavior?
  - Programs that behave (externally) like humans
- Computational models of human "thought" processes?
  - Programs that operate (internally) the way humans do
- Computational systems that behave intelligently?
  - What does it mean to behave intelligently?
- Computational systems that behave rationally!
  - More on this later
- AI applications
  - Monitor trades, detect fraud, schedule shuttle loading, etc.

Agents

Software that gathers information about an environment and takes actions based on that information.

- a robot
- a web shopping program
- a factory
- a traffic control system…

The Agent and the Environment

How do we begin to formalize the problem of building an agent?

- Make a dichotomy between the agent and its environment
- Not everyone believes that making this dichotomy is a good idea, but we need the leverage it gives us.

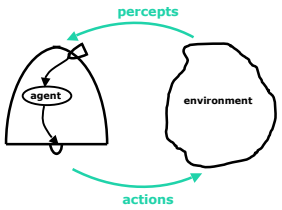

World Model

- A – the action space
- P – the percept space
- E – the environment: $A^* \rightarrow P$
- Alternatively, define
  - S – internal state [may not be visible to agent]
  - Perception function: $S \rightarrow P$
  - World dynamics: $S \times A \rightarrow S$

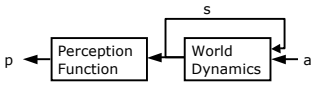

Rationality

- A rational agent takes actions it believes will achieve its goals.
  - Assume I don't like to get wet, so I bring an umbrella. Is that rational?
  - Depends on the weather forecast and whether I've heard it. If I've heard the forecast for rain (and I believe it) then bringing the umbrella is rational.
- Rationality ≠ omniscience
  - Assume the most recent forecast is for rain but I did not listen to it and I did not bring my umbrella. Is that rational?
  - Yes, since I did not know about the recent forecast!
- Rationality ≠ success
  - Suppose the forecast is for no rain but I bring my umbrella and I use it to defend myself against an attack. Is that rational?
  - No, although successful, it was done for the wrong reason.

Limited Rationality

- There is a big problem with our definition of rationality…
- The agent might not be able to compute the best action (subject to its beliefs and goals).
- So, we want to use limited rationality: "acting in the best way you can subject to the computational constraints that you have"
- The (limited rational) agent design problem:

  Find $P^* \rightarrow A$
  - mapping of sequences of percepts to actions
  - maximizes the utility of the resulting sequence of states
  - subject to our computational constraints

Issues

- How could we possibly specify completely the domain the agent is going to work in?
  - If you expect a problem to be solved, you have to say what the problem is!
  - Specification is usually iterative: Build agent, test, modify specification
- Why isn't this "just" software engineering?
  - There is a huge gap between specification and the program
- Isn't this automatic programming?
  - It could be, but AP is so hard most people have given up
  - We're not going to construct programs automatically!
  - We're going to map classes of environments and utilities to structures of programs that solve that class of problem

Thinking

- Is all this off-line work AI? Aren't the agents supposed to think?
- Why is it ever useful to think? If you can be endowed with an optimal table of reactions/reflexes ($P^{*!}$ A) why do you need to think?
- The table is too big! There are too many world states and too many sequences of percepts.
- In some domains, the required reaction table can be specified compactly in a program (written by a human). These are the domains that are the target of the "Embodied AI" approach.
- In other domains, we'll take advantage of the fact that most things that could happen – don't. There's no reason to pre-compute reactions to an elephant flying in the window.

Learning

- What if you don't know much about the environment when you start or if the environment changes?
  - Learn!
  - We're sending a robot to Mars but we don't know the coefficient of friction of the dust on the Martian surface.
  - I know a lot about the world dynamics but I have to leave a free parameter representing this coefficient of friction.
- Part of the agent's job is to use sequences of percepts to estimate the missing details in the world dynamics.
- Learning is not very different from perception, they both find out about the world based on experience.
  - Perception = short time scale (where am I?)
  - Learning = long time scale (what's the coefficient of friction?)

Classes of Environments

- Accessible (vs. Inaccessible)
  - Can you see the state of the world directly?
- Deterministic (vs. Non-Deterministic)
  - Does an action map one state into a single other state?
- Static (vs. Dynamic)
  - Can the world change while you are thinking?
- Discrete (vs. Continuous)
  - Are the percepts and actions discrete (like integers) or continuous (like reals)?

Example: Backgammon (http://www.bkgm.com/rules.html)

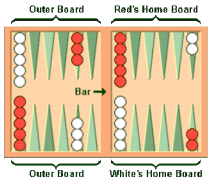

Backgammon is a game for two players, played on a board consisting of twenty-four narrow triangles called points. The triangles alternate in color and are grouped into four quadrants of six triangles each. The quadrants are referred to as a player's home board and outer board, and the opponent's home board and outer board. The home and outer boards are separated from each other by a ridge down the center of the board called the bar.

The points are numbered for either player starting in that player's home board. The outermost point is the twenty-four point, which is also the opponent's one point. Each player has fifteen stones of his own color. The initial arrangement of stones is: two on each player's twenty-four point, five on each player's thirteen point, three on each player's eight point, and five on each player's six point.

Both players have their own pair of dice and a dice cup used for shaking. A doubling cube, with the numerals 2, 4, 8, 16, 32, and 64 on its faces, is used to keep track of the current stake of the game.

Backgammon-Playing Agent

- Action space – A
  - The backgammon moves
    - Motor voltages of the robot arm moving the stones?
    - Change the (x,y) location of stones?
    - Change which point a stone is on? ["Logical" actions]
- Percepts – P
  - The state of the board
    - Images of the board?
    - (x,y) locations of the stones?
    - Listing of stones on each point? ["Logical" percepts]

Backgammon Environment

- Accessible?
  - Yes!
- Deterministic?
  - No! Two sources of non-determinism: the dice and the opponent
- Static?
  - Yes! (unless you have a time limit)
- Discrete?
  - Yes! (if using logical actions and percepts)
  - No! (e.g. if using (x,y) positions for actions and percepts)
  - Images are discrete but so big and finely sampled that they are usefully thought of as continuous.

Example: Driving a Taxi

Recitation Exercise: Think about how you would choose –

- Action space – A?
- Percept space – P?
- Environment – E?

Structures of Agents

- Reflex ("reactive") agent
  - No memory

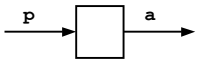

- What can you solve this way?
  - Accessible environments
    - Backgammon
    - Navigating down a hallway

Structures of Agents

- Agent with memory

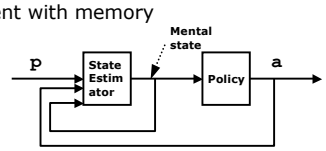

- State estimator/Memory
  - What we've chosen to remember from the history of percepts
  - Maps what you knew before, what you just perceived and what you just did, into what you know now.
- Problem of behavior: Given my mental state, what action should I take?

Planning Agent Policy

Planning is explicitly considering future consequences of actions in order to choose the best one.

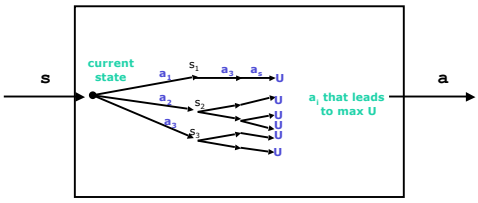

- "Let your hypotheses die in your stead." – Karl Popper