What is Artificial Intelligence?

CPSC 322 Lecture 1

January 4, 2006

Lecture Overview

Logistics

What is Artificial Intelligence?

What is an Agent?

Essentials

- Course web page: http://cs.ubc.ca/~kevinlb/teaching/cs322
  - This is where most information about the course will be posted, most handouts (e.g., slides) will be distributed, etc.
  - Check it often!
- Textbook: *Computational Intelligence, 2nd Edition*, by Poole, Mackworth and Goebel. Still under development.
  - it's free!
  - it's only available electronically
- WebCT: used for textbook, discussion board
  - Use the discussion board for questions about assignments, material covered in lecture, etc, rather than email
  - If you do use email, you'll just be asked to repost to the message board, and your answer will take longer!
- My office hours: Thursdays 2:15 — 3:45; TA hours TBA

Course Elements

Grading:

- Assignments: 20%
  - 3–4; not necessarily weighted equally
  - you get 3 late days
- Midterm: 30%
- Final: 50%

Academic Conduct

Submitting the work of another person as your own (plagiarism) constitutes academic misconduct, as does disallowed communication with others (either as donor or recipient):

- Assignments are to be done alone. You may not submit any solution not written by yourself, look at another student's solution (including solutions from assignments completed in the past), or previous sample solutions, and you may not share your own work with others.
- You may discuss your solutions and design decisions with your fellow students. You can talk about the assignments, but you cannot look at or copy other people's answers.

Academic misconduct is very serious, and is subject to penalties ranging from a grade of zero to indefinite suspension from the University. More information is on the course web page.

What is Artificial Intelligence?

Some definitions that have been proposed:

- Systems that think like humans
- Systems that think rationally
- Systems that act like humans
- Systems that act rationally

Thinking Humanly

Model the cognitive functions of human beings

- Humans are our only example of intelligence: we should use that example!
- But... humans often act in ways that we don't consider intelligent
- And... we have to spend most of our effort on studying how people's minds operate, rather than thinking about what intelligence ought to mean in various domains.

Acting Humanly

The Turing Test

- Don't try to come up with a list of characteristics that computers must satisfy to be considered intelligent
- Instead, use an operational definition: consider them intelligent when people can't tell them apart from other people

The original test involved typing back and forth; the Total Turing Test includes a video signal to test perception too

- But... is acting just like a person what we really want?
- For example, don't people often do things that we *don't* consider intelligent?

Thinking Rationally

Rationality: an abstract "ideal" of intelligence, rather than "whatever humans do"

- Example: ancient Greeks invented syllogisms: argument structures that always yield correct conclusions given correct premises
  - This led to logic, which we'll discuss in this course
- Example: a rational player will always win or tie when she plays tic-tac-toe, while some humans lose
  - I hope all of you are at least this rational, however...

But... can we characterize what rational thought ought to look like in a clear (formal) way? People have tried, and haven't really succeeded...

Acting Rationally

AI should aim to build agents: artifacts that are able to act rationally in their environments

- they act appropriately given goals and circumstances
- they are flexible to changing environments and goals
- they learn from experience
- they make appropriate choices given perceptual and computational limitations

We'll emphasize this definition of AI

- rationality is more cleanly defined than human behavior, so it's a better metric
- when human behavior is not rational, often we'd prefer rationality
  - Example: you wouldn't want a shopping agent to make impulsive purchases!
- It's often easier to define rational action than rational thought

What is an agent?

It has the following characteristics:

- It is situated in some environment, which it is able to observe (perhaps imperfectly) and in which it is able to act (perhaps within constraints)
- It has goals or preferences
- It has prior knowledge or beliefs, and has some way of updating these beliefs based on new experiences

Agents acting in an environment

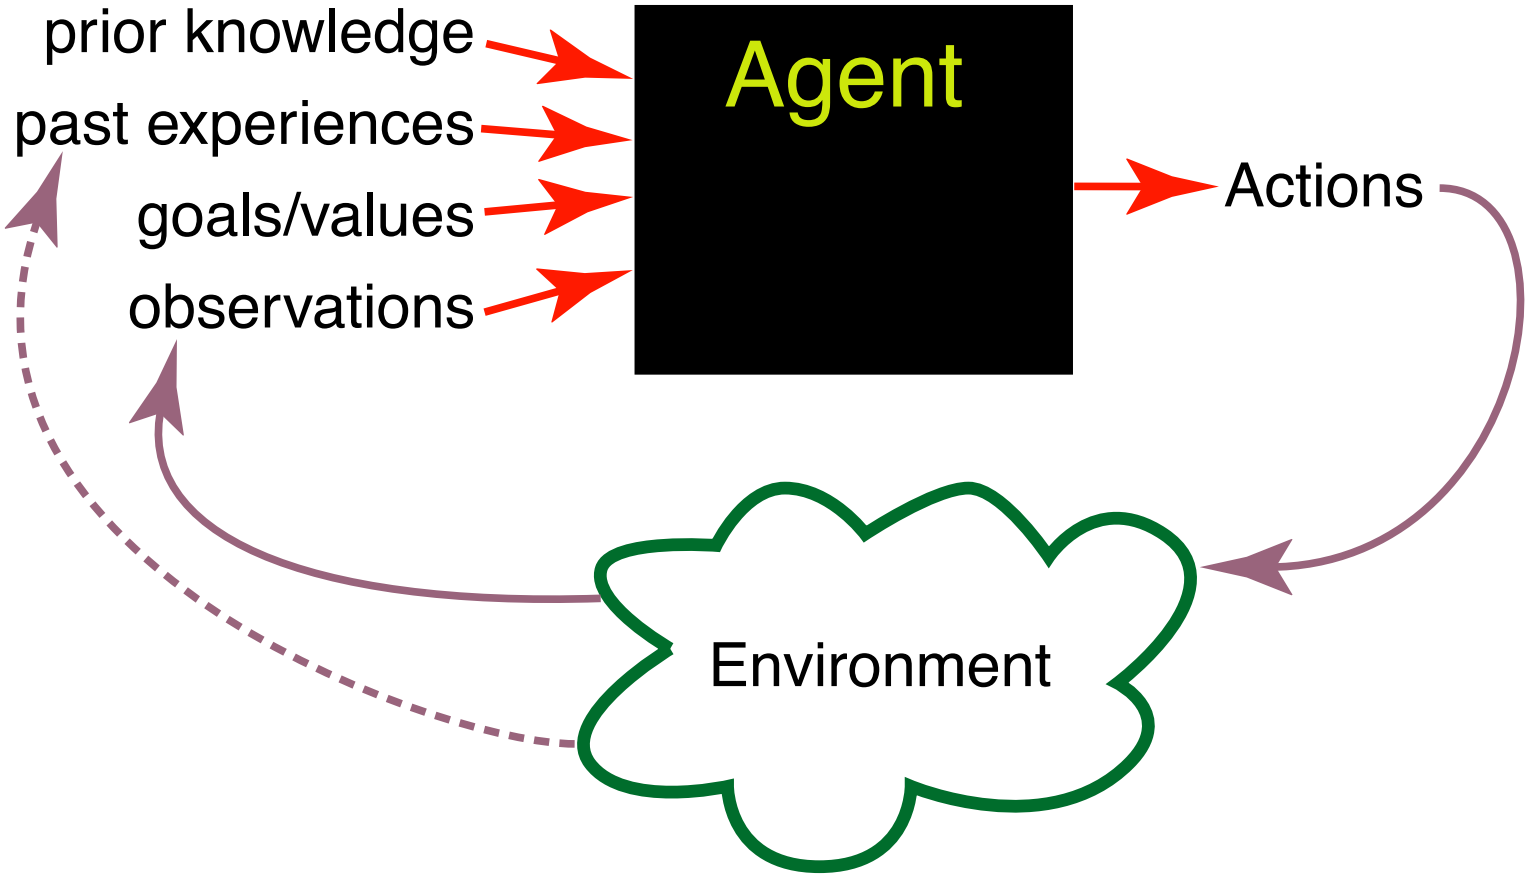

Examples

Which of these things is an agent?

- A soccer-playing robot?
- A rock?
- A Google web crawler?
- A thermostat?
- A dog?
- A car?