

CSCI 5582 Artificial Intelligence

Lecture 2
Jim Martin

CSCI 5582 Fall 2006

Today 8/31

- Review
- Intelligent agents
- Administrative stuff
- Turing
- Social agents

CSCI 5582 Fall 2006

Review: Our Framework

- AI is concerned with the creation of artifacts that...
 - Do the right thing
 - Given what their circumstances and what they know

CSCI 5582 Fall 2006

Intelligent Agents

- What is an agent?
- What makes an agent rational?

CSCI 5582 Fall 2006

Ideal Rational Agents

"... should take whatever action is expected to maximize its performance measure on the basis of its percept sequence and whatever built-in knowledge it has"

Key points:

- Performance measure
- Actions
- Percept sequence
- Built-in knowledge

CSCI 5582 Fall 2006

Agents as Functions

A mapping

from some relevant set of conditions (past actions, current sensors, etc)
to an action

CSCI 5582 Fall 2006

Implementation

- Table-based
- Reflex-based
- Model-based
- Goal-based
- Utility-based

CSCI 5582 Fall 2006

Table-based Agents

- What are they?
- What's wrong with them?
- What's right about them?


CSCI 5582 Fall 2006

Reflex-based Agents

- What are they?
- What's good about them?
- What's wrong with them?
- Are they fundamentally different from table-based agents?

CSCI 5582 Fall 2006

Reflex Agents




CSCI 5582 Fall 2006

Model-based Agents

- What's wrong with pure reflex?

CSCI 5582 Fall 2006

Model-based Agents



CSCI 5582 Fall 2006

Goal-based Agents

Agents that take actions in the pursuit of a goal or goals.

CSCI 5582 Fall 2006


Goals

You can think of goals in a number of different ways:

- As a specific state of the world
- As a set of states that satisfy some criteria
- As an operational test that applies to states and says whether or not they satisfy a goal criteria

CSCI 5582 Fall 2006

Goal-based Agents



CSCI 5582 Fall 2006

Goals and the Future

Goals introduce the need to reason about the future or other hypothetical states. It may be the case that none of the actions an agent can currently perform will lead to a goal state.

What should it do?

CSCI 5582 Fall 2006

Utility-based Agents


Agents that take actions that make them the most happy in the long run.

More formally agents that prefer actions that lead to states with higher utility.

Utility-based agents can reason about multiple goals, conflicting goals, and uncertain situations.

CSCI 5582 Fall 2006

Utility-based Agents



CSCI 5582 Fall 2006

Administration

- See me after class if you need to sign up for this class.
- Questions?

CSCI 5582 Fall 2006

Homework

- Two parts
 1. Answer a simple question
 2. Write a simple Python program
- Due on 9/7

CSCI 5582 Fall 2006

Simple Question

- What's the population of Boulder?
 - Answer the question using the Web
 - Describe how you answered the question
 - Give a brief overview of the design of a system that could do what you did to find the answer

CSCI 5582 Fall 2006

Program

- Write a simple program that determines whether or not a mobile is balanced.
 - Mobiles have two rods with weights on them.
 - Weights are either simple weights or other mobiles.
 - A mobile is balanced if
 - the torque on its arms are the same and
 - every component mobile is balanced

CSCI 5582 Fall 2006

Errors

- What kind of errors can your program make?
 - Thinking that an unbalanced mobile is balanced
 - Type I, false positive
 - Thinking that a balanced mobile is unbalanced
 - Type II, false negative

CSCI 5582 Fall 2006

HW Format

- For programming assignments, you should submit a hardcopy listing in a format similar to that created by a2ps.
- When I request your code electronically I'll usually ask for an email *.py attachment with a name like lastname-something.py.
 - Don't tar, zip, uuencode or anything like it.
- For writing assignments, you should submit output formatted using something like LaTeX, MS Word.

CSCI 5582 Fall 2006

Turing Test

Turing (1950) was interested in the following question:

- Can machines think?

But he immediately decides that answering this question directly is hopeless.

CSCI 5582 Fall 2006

Turing Test

Instead Turing proposes a game with three participants:

- A computer
- A human questioner/player
- A second human participant

CSCI 5582 Fall 2006

The Game

- Typed input/output only
- Any kind of question is fair.
- The player poses questions to the computer/other human.
- Can the player reliably distinguish the computer from the human?

CSCI 5582 Fall 2006

Passing the Test

- What would it take for a machine to pass the test?
- What would it mean if a machine passed the test?

CSCI 5582 Fall 2006

Social Interaction

- Turns out that in some sense it's easy to pass the Turing test
- People are prone towards attributing human qualities to all manner of sufficiently complex technology.

CSCI 5582 Fall 2006

Nass and Reeves

People

- Are polite to computers
- Respond emotionally to computers
 - Criticism and praise
- Attribute human qualities to computers based on surface attributes
 - Name of systems
 - Male vs. female voices in TTS systems

CSCI 5582 Fall 2006

Implications

- Whether or not we set out to build intelligent interactive agents people expect computers to act like we do.
- So we may as well build them so they meet those expectations.

CSCI 5582 Fall 2006

Next Time

- We'll start on state space search
- Finish reading Chapters 1, 2 and 3
- Finish the first assignment by Thursday

CSCI 5582 Fall 2006
