

What is Artificial Intelligence?

CPSC 322 – Intro 1

Textbook §1

Lecture Overview

- 1 Logistics
- 2 What is Artificial Intelligence?
- 3 What is an Agent?

Essentials

- Course web page: <http://www.ugrad.cs.ubc.ca/~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: *Artificial Intelligence: Foundations of Computational Agents*, by Poole and Mackworth. 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. That way others can learn from your questions and comments!
 - Use email for private questions (e.g., grade inquiries or health problems).

Course Elements

Grading:

- Assignments: 20%
- Midterm: 30%
- Final: 50%

Course Elements

Grading:

- Assignments: 20%
- Midterm: 30%
- Final: 50%

If your final grade is $\geq 20\%$ higher than your midterm grade:

- Assignments: 20%
- Midterm: 15%
- Final: 65%

Assignments

- There will be **five** assignments in total
 - counting “assignment zero”, which you’ll get today
 - They will not be weighted equally
- Group work
 - code questions:
 - you can work with a partner
 - written questions:
 - you may **discuss** questions with other students
 - you may not **look at or copy each other’s written work**
 - you’ll be asked to sign an **honour code** saying you’ve followed these rules

Late Days

- You get **four late days**
 - to allow you the flexibility to manage unexpected issues
 - moderate illness, conflicts with other courses, travel, extracurricular obligations, job interviews, etc.
 - additional late days will **not** be granted except under truly exceptional circumstances
- a day is defined as all or part of a 24-hour block of time beginning at 4 PM on the day an assignment is due
- applicable to assignments 0–4
 - subject to a few restrictions
 - **not** applicable to midterm, final!
- if you've used up all your late days, you lose 20% per day

Missing Assignments or Final

Hopefully late days will cover almost all the reasons you'll be late in submitting assignments. However, it's possible that your coursework will be disrupted by something more serious like an extended illness.

- For all such cases, you'll need to **provide a note** from your doctor, psychiatrist, academic advisor, etc.
- If you **miss an assignment**, your score will be reweighted to exclude that assignment
- If you **miss the midterm**, those grades will be shifted to the final
 - thus, your total grade will be 80% final, 20% assignments
- If you **miss the final**, you'll have to write a make-up final as soon as possible.

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.

Assignment 0

- This assignment asks you to describe an AI agent from fiction, and to explain some high-level details about how it works.
- The assignment is available from WebCT
- It must be handed in using `handin`
- It's due in one week.

Lecture Overview

- 1 Logistics
- 2 What is Artificial Intelligence?
- 3 What is an Agent?

What is Artificial Intelligence?

Some definitions that have been proposed:

- 1 Systems that think like humans
- 2 Systems that think rationally
- 3 Systems that act like humans
- 4 Systems that act rationally

1. 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.

2. 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 it intelligent when **people can't tell a computer 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, again, don't people often do things that we *don't* consider intelligent?

3. Thinking Rationally

Rationality: an **abstract “ideal” of intelligence**, rather than “whatever humans do”

- Example: ancient Greeks invented **sylogisms**: 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...

4. 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 design objective
- in cases where 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

Lecture Overview

- 1 Logistics
- 2 What is Artificial Intelligence?
- 3 What is an Agent?**

What is an agent?

It has the following characteristics:

- It is situated in some **environment**
 - does not have to be the real world—can be an abstracted electronic environment
- It can make **observations**
 - perhaps imperfectly
- It is able to **act**
 - perhaps within constraints
- It has **goals or preferences**
- It may have **prior knowledge or beliefs**, and some way of **updating beliefs** based on new experiences

Agents acting in an environment

