What is Artificial Intelligence?

CMPUT 366: Intelligent Systems

P&M Chapter 1

Don't Come to Campus

- **All** of Computing Science's courses are online-only this semester
- CSC and Athabasca Hall are **closed**
  - You can only come if you are explicitly required to by an instructor
  - Even in that case, the Chair and/or Dean need to sign off

Intelligent Systems

- This course is about constructing **intelligent agents**.
- But what does that **mean**?
  - Smarter than the smartest genius?
    - (wait, what does "smart" mean?)
  - Able to do things that computers are pretty bad at?
  - Able to trick a human into thinking it's another human?
- We'll try to define both **intelligent** and **agent** more formally

Lecture Outline

1. Course Logistics
2. What is Artificial Intelligence?
3. AI Seminar!

Course Essentials

Course information: https://eclass.srv.ualberta.ca/course/view.php?id=68187

- This is the main source for information about the class
- Slides, readings, assignments, deadlines

Lectures: Mondays, Wednesdays, and Fridays, 11:00-11:50am on Zoom

- Lectures will be recorded and posted on eClass

eClass Discussion forum for **public** questions about assignments, lecture material, etc.

Email: james.wright@ualberta.ca for **private** questions

- (health problems, inquiries about grades)

Office hours: After lectures on Mondays & Fridays, or by appointment

- TA office hours will be announced on Friday

CMPUT 366 in One Slide

- Focus on **intelligent agents**
  - Intelligence
  - Agents
- Survey methods to construct such agents
  - classic
  - contemporary
- **This is *not* a reinforcement learning class**
  - Reinforcement learning class is CMPUT 397

Readings

We will draw from a lot of texts for this class. *BUT*, they are all available online for free:

[P&M] David Poole and Alan Mackworth,
Artificial Intelligence: Foundations of Computational Agents, 2nd edition.

[Bar] David Barber, *Bayesian Reasoning and Machine Learning*.

[GBC] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*.

[S&B] Richard S. Sutton and Andrew G. Barto,
Reinforcement Learning: An Introduction, 2nd edition.

[S&LB] Yoav Shoham and Kevin Leyton-Brown,
Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.

Readings for each lecture are listed on the schedule and on eClass.

Evaluation

Grade breakdown

- Assignments: 30%
- Midterm exam: 30%
- Final exam: 40%

Late assignments

- 20% deducted per day

Missed assignments or exams

- **Provide a note** from doctor, academic advisor, etc.
- Assignments score will be **reweighted** to exclude missed assignments
- If the midterm exam is missed, the mark from the **final exam** will be used in its place
  - i.e., grade will be 30% assignments, 70% final exam

Assignments

- There will be **four assignments** (roughly every 3 weeks)

- Types of questions:
  - **Short answer**: definitions, distinctions, etc.
    "What is a Nash equilibrium?"
  - **Model construction**: "*Represent XYZ as a graph search problem*"
  - **Algorithmic considerations**: "*What would be an appropriate algorithm to answer XYZ? Why?*"
  - Small **implementation** task

- Assignments are submitted electronically (via eClass)

Collaboration Policy

Detailed version on the syllabus

You are **encouraged to discuss assignments** with other students:

1. You must **list** everyone you talked with about the assignment.
2. You **may not** share or look at each other's **written work or code**.
3. You must **write up** your solutions individually

Individual work only on **exams**: No collaboration allowed

Academic Conduct

- Submitting someone else's work as your own is **plagiarism**.
- So is helping someone else to submit your work as their own.
- We report **all cases** of academic misconduct to the university.
- The university takes academic misconduct **very seriously**.
  Possible consequences:
  - Zero on the assignment or exam (virtually guaranteed)
  - Zero for the course
  - Permanent notation on transcript
  - Suspension or expulsion from the university

Spot checks

- I won't be using a proctoring service for exams
- Instead, we will use **spot checks**
  - After every exam, some students will be selected to **verbally explain** their answers to a TA
  - If you can't explain how you got your answer, you may not get credit for the question

Getting chosen for a spot check is not an accusation of cheating

Prerequisites

- Comfort with or interest in formal, **mathematical/algorithmic reasoning**
- Basic **probability**: random variables, expectations, conditional probability. (There will be a refresher lecture)
- Basic **calculus**: gradients, partial derivatives, vector norms
- Basic **graph theory**: Nodes, edges
- Ability to program in **Python**
    - Most assignments will have a programming component
    - TAs will run a refresher session

What is Artificial Intelligence?

1. Think like humans	2. Act like humans
3. Think rationally	4. Act rationally

Two dimensions:

- Reasoning vs. acting
- Mimicking humans vs. rationality

1. Thinking Humanly

Model the **cognitive processes** of humans

Benefits:

- We know humans are intelligent!
  Why not learn from that example?

- Understanding human cognition is scientifically valuable in itself.

Drawbacks:

- Cognitive science is really hard!

- Humans often think in ways that we wouldn't call "intelligent"

2. Acting Humanly

The Turing Test:

- Don't try to define exactly what makes a system intelligent
- If you can act intelligently enough that people **can't tell you apart** from other people, then you are effectively intelligent

Drawbacks:

- Is acting exactly like a person really what we want?
  (We already know how to make more people...)
- Don't people often behave pretty unintelligently?

3. Thinking Rationally

Rationality: An ideal of what intelligent cognition **should** do

Benefits:

- Leads to more effective agents
- Not just "whatever people do, even when that's terrible"
- Philosophically important! What *is* rational thinking?

Drawbacks:

- Difficult to define formally! What *is* rational thinking?

4. Acting Rationally

Rational action: Doing what is most likely to best achieve our goals

Benefits:

- More clearly defined than human behaviour
- When human behaviour is irrational, we'd usually prefer the rational behaviour
    - Or would we? Counter-examples?
- Rational **behaviour** is also easier to define than rational **thought**

What is Artificial Intelligence?

	Reasoning	Acting
Like Humans	1. Think like humans	2. Act like humans
Rationally	3. Think rationally	4. Act rationally

Questions:

1. Which of these definitions do you find most **convincing**?
2. What is **missing** from these definitions?

Rational Agents

An **agent** is a system that **acts** in an **environment**
to achieve **goals** or optimize **preferences**.

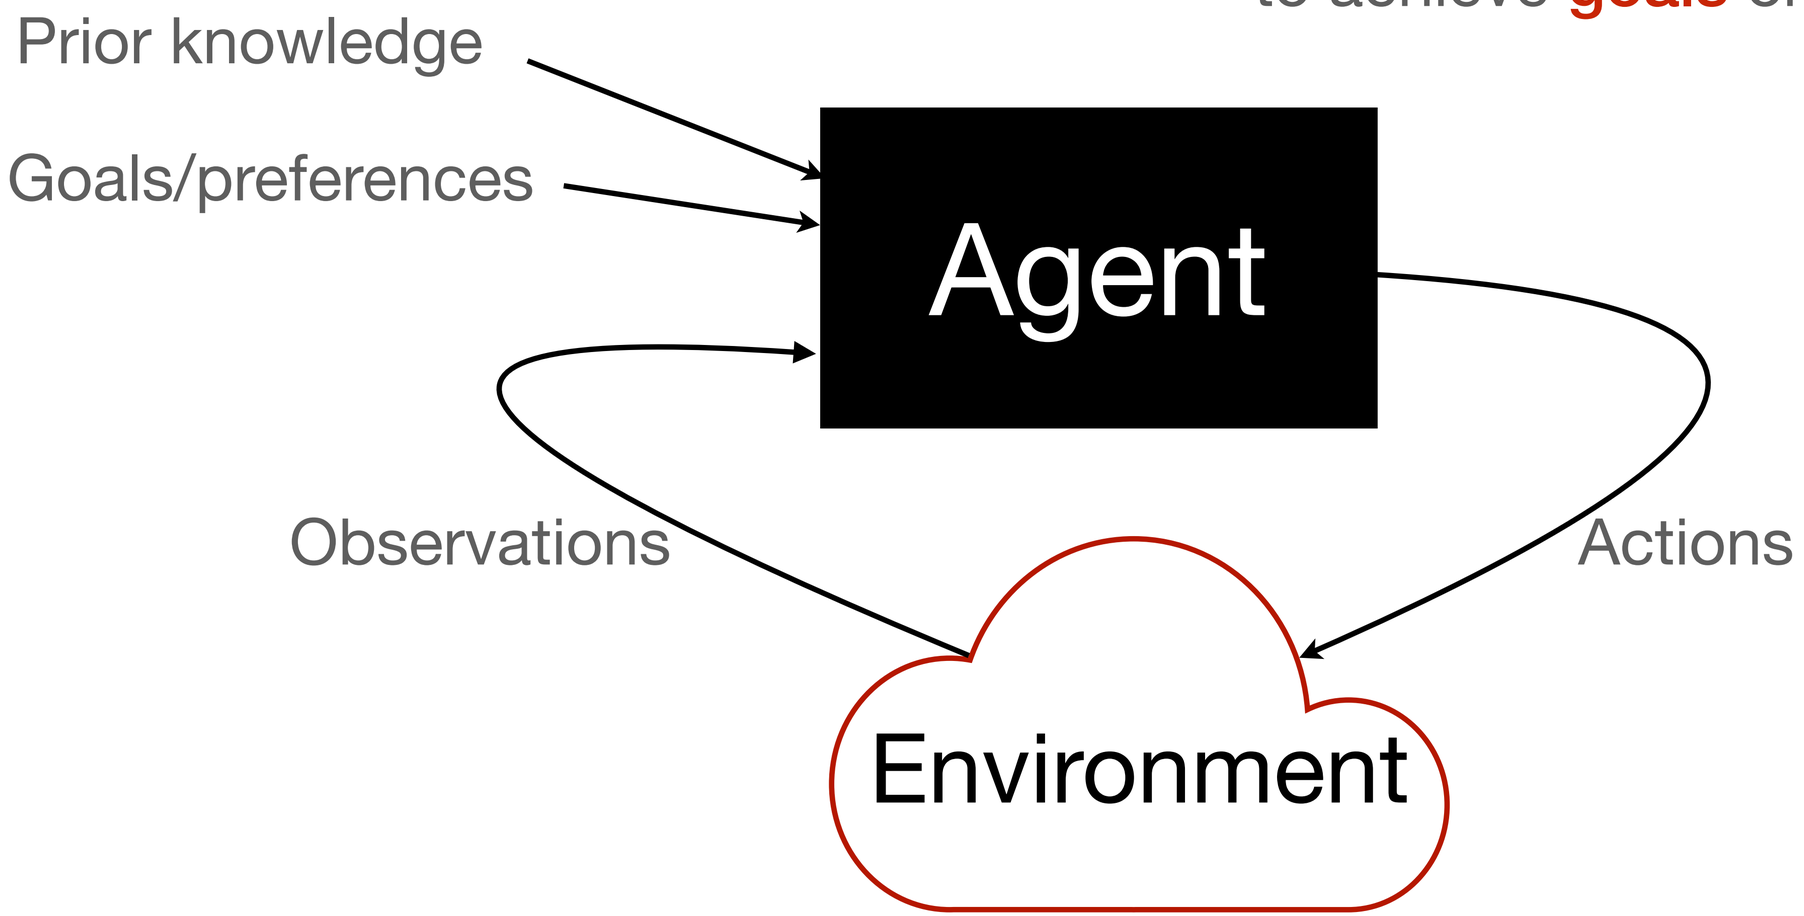

Course Topics

- Search
- Reasoning Under Uncertainty
- Causality
- Supervised Learning
- Deep Learning
- Reinforcement Learning
- Multiagent Systems

Summary

- Course details on **eClass**: https://eclass.srv.ualberta.ca/course/view.php?id=68187
- This course will focus on the construction of **rational agents**
  - **Agent:** System that **acts** in an **environment** to achieve **goals**
  - **Rational action:** Do what **best** achieves explicit goals

AI Seminar

What: Great talks on cutting-edge AI research
External (e.g., DeepMind, IBM) and internal speakers

When: Fridays at noon

~~But come at 11:45 for free pizza / good seats~~

Where: ~~CSC 3-33~~ Online Zoom meeting
Website: sites.google.com/ualberta.ca/ai-seminar/

Announcements: Sign up for mailing list (bottom of webpage)