What is Artificial Intelligence?

CMPUT 366: Intelligent Systems

P&M Chapter 1

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 webpage: jrwright.info/aicourse/

- This is the main source for information about the class

- Slides, readings, assignments, deadlines

Contacting us:

- Discussion board: piazza.com/ualberta.ca/winter2019/cmput366
  for **public** questions about assignments, lecture material, etc.

- Email: james.wright@ualberta.ca
  for **private** questions (health problems, inquiries about grades)

- TA office hours:

  - Maryam: Thursdays 3:30pm to 5:00pm

  - Ji: TBD

- Labs: Mondays 5:00pm to 7:50pm (CAB 235)
  No labs or TA office hours this week

- Office hours: After every lecture, or by appointment

Readings

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

1. David Poole and Alan Mackworth,
   Artificial Intelligence: Foundations of Computational Agents, 2nd edition.

2. David Barber, *Bayesian Reasoning and Machine Learning*.

3. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*.

4. Richard S. Sutton and Andrew G. Barto,
   Reinforcement Learning: An Introduction, 2nd edition.

5. Yoav Shoham and Kevin Leyton-Brown,
   Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.

Readings for each lecture are listed on the schedule.

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

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

Assignments

There will be **four** assignments (not weighted equally)

You are **encouraged to discuss** assignment questions with other students:

1. You **may not** share or look at each other's **written work or code**.

2. You must **write up** your solutions individually

3. You must **list** everyone you talked with about the assignment.

Prerequisites

- Comfort with or interest in formal, **mathematical reasoning**

- Basic **probability**: random variables, expectations, conditional probability. (There will be a refresher lecture)

- Basic **calculus**: gradients, vector norms

- Ability to program in **Python**

  - Most assignments will have a programming component

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**

Rational Agents

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

Prior knowledge

Goals/preferences

Agent

Observations

Actions

Environment
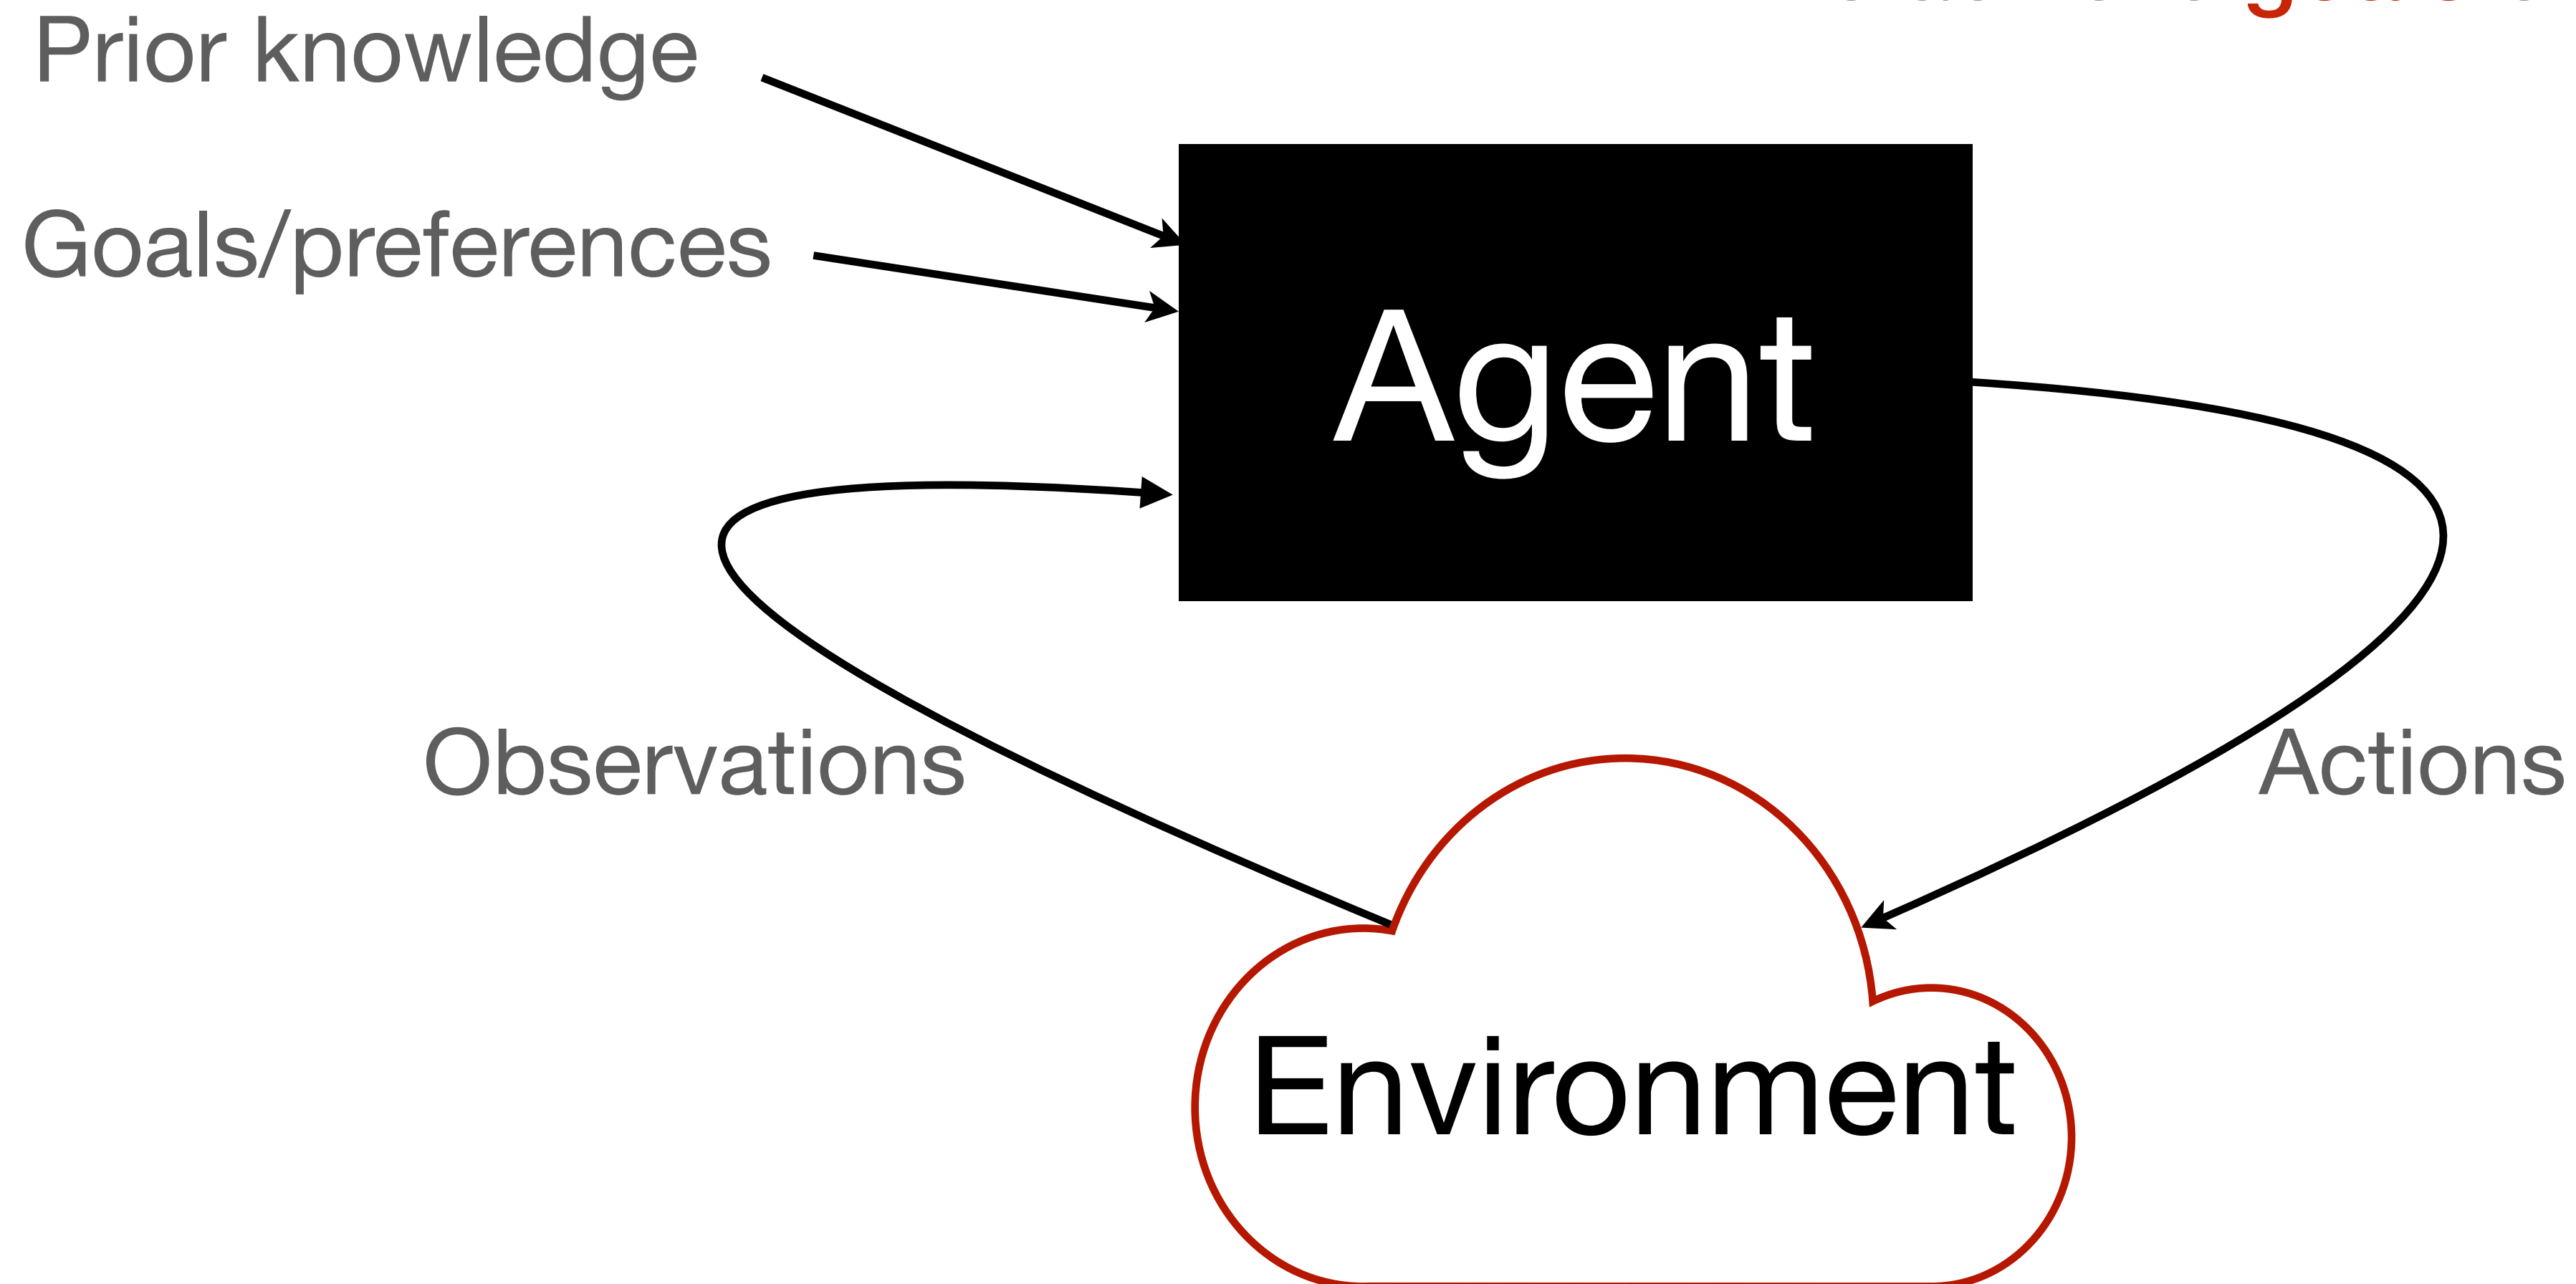
Course Topics

- Search

- Reasoning Under Uncertainty

- Causality (I hope!)

- Supervised Learning

- Deep Learning

- Reinforcement Learning

- Multiagent Systems

Summary

- Course details on the **website**: <u>jrwright.info/aicourse</u>

- 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
(Also free pizza!)

When: Fridays at noon
Where: CSC 3-33
Calendar: www.cs.ualberta.ca/~ai/cal/

Announcements: Sign up for **ai-seminar**
www.mailman.srv.ualberta.ca/