

# **Introduction to Artificial Intelligence (AI)**


**CPSC 322 Lecture 1**

# People


## Instructor

**Jordon Johnson** (jordon@cs.ubc.ca; office ICCS 255)

**TAs** (*contact info on course website*)



Raymond  
Li



Egor  
Peshkov



Lena  
Podina



Cindy  
Tu



Matthew  
Wilson

# About Me

M.Sc. 2018 (UBC CS)

Sci-fi/fantasy fan

Favorite book: **The Hitchhiker's Guide to the Galaxy**

I **love** questions

**Horrible at remembering names...**

So I don't mind if you forget mine!



# Course Essentials

- **Course website:**

[www.cs.ubc.ca/~jordon/teaching/cpsc322/2019w2](http://www.cs.ubc.ca/~jordon/teaching/cpsc322/2019w2)

- Course information and policies
- Schedule and **lecture slides**

- **Lectures:** (*grouped by week*)

- Slides will be posted at the beginning of each week
- Most lectures will include iClicker questions (register on Canvas)



- **Textbook:** *Artificial Intelligence*, 2nd Edition, by Poole & Mackworth.

- Free online (link on course website)

# Course Essentials

- **Piazza** : [Sign-up link on course website and/or Canvas](#)
  - Course announcements
  - Discussion board
  - **FAQ** (*a.k.a.* “Jordon’s wall of NO”)
  - [Use email for private questions](#) (e.g., grade inquiries, concession requests)



- **Canvas**: assignments, solutions, other files

- **Aispace** : online tools for learning Artificial Intelligence <http://aispace.org/>

- Under development here at UBC!
- **We will use this often**



# Course Elements

- **Assignments:** 20%
- **Midterm:** 30%
- **Final:** 50%
- **iClickers:** up to 2% bonus (1% participation, 1% correct answers)

**You must pass the final (or the weighted average of the midterm and final) to pass the course**

**If your final grade is higher than your midterm grade, the exams may be re-weighted in favor of the final (specific mechanics TBD)**

# Assignments

- There will be **four** assignments in total (+ Assignment 0)
  - **Not** weighted equally
  - Submit on Canvas as a single PDF file
  - Only the most recent submission will be graded
  - **Handwritten/hand-drawn work will not be graded**
- **Group work**
  - You are **encouraged** to work with a partner
    - “Search for Teammates” post on Piazza
    - **5% bonus** on assignments (caps at 100% per assignment)
    - You should **not** simply divide the assignment up between you

# Assignments: Late Days

- Hand in by 11:59pm on due date (on Canvas)
- You get **four late days** 😊
  - See course website for details
- If you don't claim late days, **you lose 20% per day**
- **Assignments more than 4 days late will not be accepted**
- Exceptions specified in assignment descriptions on Canvas



# Missing Assignments / Midterm / Final

Hopefully late days will cover most issues with submitting assignments

Unfortunately, **more serious issues** can arise

- Contact the instructor as soon as possible
- Details about standard concessions are on the course website

# How to Get Help?

- Go to **office hours**
  - **Poll for office hour availability is on Piazza**
  - Office hours will be posted on Piazza
  - You may schedule an **appointment** if you cannot attend office hours or in special circumstances
- Use the course **discussion board** on **Piazza** for questions on course material (so keep reading from it !)
  - Please **search before you post**

# Academic Conduct

Policies on the course website (with link to department policy)

**If you are in doubt whether the line is crossed:**

- Talk to instructor or the TAs
- Ignorance of the rules will **not** be a sufficient excuse for breaking them

**My advice:** better to skip an assignment (or miss out on some iClicker points) than to have “**academic misconduct**” recorded on your transcript and additional penalties as serious as expulsion from the university!

# My approach to teaching

- I value **conceptual understanding** over **memorization**
- If you prioritize **understanding** over **grades**, then you increase the likelihood of getting **BOTH**.
- If you prioritize **grades** over **understanding**, then you increase the likelihood of getting **NEITHER**.

# My approach to teaching

- I will sometimes leave things **underspecified**
  - Eg. I may ask you to solve a problem without telling you the precise steps to follow or what formats to use
- This is usually **intentional**
  - the point is for you to figure things out using the course concepts and your own **reasoning/problem solving** skills

# My approach to teaching

- I like answers that are:
  - **Correct**
  - **Easy to read**
  - As **short** as possible while including **important** information
- I may **penalize** answers that include information that is not directly related to the question

# To Summarize

- All the course logistics/policies are described in the course Webpage

[www.cs.ubc.ca/~jordon/teaching/cpsc322/2019w2](http://www.cs.ubc.ca/~jordon/teaching/cpsc322/2019w2)

(And summarized in these slides)

- Make sure you carefully read and understand them!

**Any questions before we  
start the material?**

# What is Intelligence?

*"For instance, on the planet Earth, man had always assumed that he was more intelligent than dolphins because he had achieved so much—the wheel, New York, wars and so on—whilst all the dolphins had ever done was muck about in the water having a good time. But conversely, the dolphins had always believed that they were far more intelligent than man—for precisely the same reasons."*

Douglas Adams, **The Hitchhiker's Guide to the Galaxy**




# What is Intelligence?

# What is Artificial Intelligence?

**Two definitions that have been proposed:**

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



# Thinking and Acting Humanly

## Model the cognitive functions of human beings

- Humans are our only example of intelligence: we should use that example!

## 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**
- **Total Turing Test** includes a video signal

# Thinking and Acting Humanly

## Problems

- Humans often think/act in ways that we don't consider intelligent
  - is acting just like a person what we really want?
- A detailed model of how people's minds operate is not yet available

# Thinking Rationally

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

- Ancient Greeks invented *syllogisms*: argument structures that always yield correct conclusions given correct premises
  - This led to **logic** and **probabilistic reasoning**, which we'll discuss in this course
- But is correct sound reasoning always enough?

# Acting (&thinking) Rationally

This course will emphasize a view of AI as building **agents**: artifacts that are able to think and act **rationally** in their environments


Rationality is **more cleanly defined** than human behavior, so it's a better design objective

*(Eg: “intelligent” vacuum cleaner: maximize area cleaned, minimize noise and electricity consumption)*

Agents that can **answer queries**, **plan actions** and **solve complex problems**

And when you have a rational agent you can always tweak it to make it irrational!

# Agents acting in an environment



# What is an agent?

It has the following characteristics:

More of these &  
more sophisticated  
=> more intelligent

- 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** (*provide an answer, buy a ticket*)
- It has **goals or preferences** (*possibly of its user*)
- It may have **prior knowledge or beliefs**, and some way of **updating beliefs** based on new experiences (to reason, to make inferences)



# Examples

Which of these things is an **agent**,  
and why or why not?

- A soccer-playing robot?
- A rock?
- Machine Translator?
- A thermostat?
- A dog?
- A car?

Which of these things are **intelligent agents**,  
and why or why not?

# Acting (and thinking) Rationally

This course will emphasize a view of AI as building **agents**: artifacts that are able to think and 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
- They **gather information** (if cost less than expected gain)

# About Machine Learning

- Machine learning (ML) is becoming more and more commonplace as a collection of techniques to allow agents to perform tasks
- **However**, ML and AI are not the same thing
- This course will not cover ML; for that, you want **CPSC 340**

# A rough CPSC 322 overview

		Environment	
Problem		Deterministic	Stochastic
Static	Constraint Satisfaction	<i>Variables + Constraints</i> Search Arc Consistency Local Search	
	Query	<i>Logics</i> Search	<i>Bayesian (Belief) Networks</i> Variable Elimination
Sequential	Planning	<i>STRIPS</i> Search	<i>Decision Networks</i> Variable Elimination

*Representation*  
Reasoning Technique



# AI magazine



YAHOO! RESEARCH

Autonomous Vehicle



AI Paintings



AI Music Composition & Performance



Robot Tour Guide

Characters for Virtual Worlds



Humanoid Robot



Social Simulation Game



Smart Environmental Controls



Intelligent Tutoring System



Vehicle Navigation System



Smart Desk with Gesture Recognition



Robots for Education



Disease Diagnosis



Drug Design

Smart Wheelchair



Handwriting & Sketch Recognition



How Can AI Systems Solve Problems Creatively?



Robotic Surgery



Recommender System



Fraud Detection



Web Search



Spam Filtering



Machine Translation



Leibniz



Aristotle



Descartes



Lovelace



Russell



Turing



Whitehead



Leibniz



Whitehead



Leibniz



Whitehead



See the AI timeline and more at  
[www.aaai.org/AIlandscape](http://www.aaai.org/AIlandscape)

# The AI Landscape

David Foray, Indiana University, Poster Development Committee Chair  
Poster: [aaai.org/AIlandscape](http://aaai.org/AIlandscape) | [www.GiacomoMarchesi.com](http://www.GiacomoMarchesi.com)



# In class activity

- Work in pairs searching the web to find **an interesting example of fielded (or experimental) intelligent agents**
- Try to find something different from the usual suspects (Alexa, Siri, Watson, etc.)
- **AAAI is the main AI association**, so feel free to start there, but you encounter plenty on a daily basis!

# Answer the following questions (take notes)

1. What does the application do? (e.g.,. control a spacecraft, perform medical diagnoses, provide intelligent help for computer users, shop on eBay, perform translation)
2. List some of the application's : goals/preferences; observations that it needs about the environment; types of actions that it performs
3. What **AI technologies** does the application use (e.g.,. belief networks, Markov models, semantic networks, heuristic search, constraint satisfaction, planning)
4. Why is it intelligent? Which aspects make it an intelligent system?
5. Is it an experimental system or a fielded system (i.e., used in a real world setting)?
6. Is there evidence on how well the application performs?

# Assignment 0

- Available on Canvas
- Also posted on Piazza (for those on the waitlist)
- Due January 18 (after the add/drop deadline)
- You may work in pairs, but each partner must submit separately (**THIS TIME ONLY**)
- You may **not** use late days, and late submissions will **not** be accepted
- The good news: **you already started it!**