

Lecture 1: Introduction

CS 486/686: Introduction to Artificial Intelligence

Plan for Today

- Introductions
- Logistics for the course
- What is AI?
 - What will **this** course cover?

Who Am I?

- My name is Kate Larson. I am a Professor at UWaterloo and a Research Scientist with Google DeepMind
- Research interests include multiagent systems, reinforcement learning, cooperative AI, game theory
- See more at <https://www.cs.uwaterloo.ca/~klarson>



Introductions: Meet Your Peers

- In the next two minutes introduce yourself to someone you don't know.
- Talk about courses, co-op, summer activities, the headaches involved in finding somewhere to live, extra-curricular activities, graduation, jobs, etc,...

Course Logistics

- Instructor: Kate Larson (DC 2518)
- Lectures
 - Monday/Wednesday
 - 8:30-9:50am in MC 2038
 - 10:00-11:20am in MC 2038
- Office Hours:
 - Mondays 1:00-2:00pm in DC2518 or by appointment
 - You can always ask questions on Piazza as well

Course Logistics

- Teaching Assistants
 - Liliana Hotsko (lhotsko)
 - Bihui Jin (b27jin)
 - Yuxuan Li (y624li)
 - Alina Lytovchenko (alytovch)
 - Hala Sheta (hsheta)
 - Zhiyuan Sun (z93sun)

Course Logistics

- Website

- <https://www.cs.uwaterloo.ca/~klarson/teaching/F26-486>
- Schedule + logistics

- Communications

- LEARN (Lecture Notes, Course Information, Assignments, Assignment Submission)
- Piazza: Link for signing up is on the course website

- Texts

- Artificial Intelligence: A Modern Approach by S. Russell and P. Norvig
- Artificial Intelligence: Foundations of Computational Agents by D. Poole and A. Mackworth
- Reinforcement Learning: An Introduction by R. Sutton and A. Barto

Evaluation

CS 486

- 4 Assignments 16%
- 10 Chat Assignments 10%
- Midterm (Oct 29) 30%
- Final Exam 44%

CS 686

- 4 Assignments 12%
- 10 Chat Assignments 5%
- Midterm (Oct 29) 30%
- Final Exam 44%
- Project Proposal 0%
- Project 9%

Chat Assignments

- You will be asked to teach an AI system (Chrysalis) what you have learned by answering a set of questions.
- Marks are awarded by the number of questions you answer (i.e. this is through participation).
- Chrysalis will participate in a quiz on your behalf and you will get feedback to help identify concepts that you might want to further study. Your grade does not depend on Chrysalis' performance on the quiz.
- See the course webpage and Learn for information on how to set up your account, etc

GenAI and Assignments

- The recent advances in generative AI are exciting and have opened many avenues for AI applications and research.
- While we encourage you to explore and experiment with generative AI, these models are not to be used in assignments unless the assignment clearly states that they are allowed. If the use of generative AI is allowed on an assignment (or part of an assignment) then proper documentation, citation, and acknowledgement is required
- You are accountable for the content and accuracy of all work you submit in this class, including any supported by generative AI

What is AI?

What is AI? (different definition from R&N)

The exciting new effort to make computers that think... machines with minds in the full and literal sense [Haugeland 85] [The automation of] activities that we associate with human thinking, such as decision making, problem solving, learning [Bellman 78]	The study of mental faculties through the use of computational models [Charniak & McDermott 85] The study of computations that make it possible to perceive, reason and act [Winston 92]
The art of creating machines that perform functions that require intelligence when performed by a human [Kurzweil 90] The study of how to make computers do things at which, at the moment, people are better [Rich&Knight 91]	A field of study that seeks to explain and emulate intelligent behavior in terms of computational processes [Schalkoff 90] The branch of computer science that is concerned with the automation of intelligent behavior [Luger&Stubblefield93]

My Favorite Definition

“Artificial Intelligence is that activity devoted to making machines intelligent, AND intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”


Nils J. Nilsson

What is AI?

- Systems that think **like humans**
 - Cognitive science
 - Fascinating area but we will not be covering it in this course
- Systems that **think rationally**
 - Aristotle: What is the correct thought process
 - Systems that reason in a logical manner
 - Systems that do inference correctly

What is AI?

- Systems that **behave like humans**
 - Turing (1950), Computing Machinery and Intelligence, Mind
 - <https://academic.oup.com/mind/article/LIX/236/433/986238>
 - Predicted that by 2000 a computer would have a 30% chance of fooling a lay person for 5 minutes
 - Anticipated all major arguments against AI
 - Suggested the major components of AI
 - Knowledge, Reasoning, Language Understanding, Learning



What is AI?

- Systems that **behave rationally**
 - Rational behaviour: “doing the right thing”
 - Rational agent approach:
 - Agent: entity that perceives and acts
 - Rational Agent: acts so as to achieve best outcome
- This is the focus of this course
 - General principles of rational agents
 - Components for constructing rational agents

What We Will Cover

- **Search**

- Uninformed and heuristic search
- Constraint satisfaction
- Multiagent and game search

- **Reasoning and Decision Making Under Uncertainty**

- Probability theory, decision theory
- Probabilistic inference and causal inference
- Bayes Nets, Hidden Markov Models, Markov Decision Processes

- **Learning**


- Supervised learning and neural networks (note that CS480 covers these topics in more depth)
- Bandits, Reinforcement learning (CS 486 focus)

- **Multiagent Systems**

- Game theory
- Multiagent reinforcement learning


Applications

Search




7	1			3		6	8
	6	3	2	5			
8					6	5	3
			8	9		7	
2			1	4			9
	9		5	7			
1	8	4					6
				8	2	1	4
5	3	6			9		2

Probabilistic Reasoning and Machine Learning



Credit: <https://developer.nvidia.com/blog/how-to-get-better-outputs-from-your-large-language-model/>



Multiagent Systems



Reasoning Under Uncertainty



Reinforcement Learning

DeepMind AI Reduces Google Data Centre Cooling Bill by 40%


30 JULY 2016
Richard Evans, Jim Gao

[Share](#)



From smartphone assistants to image recognition and translation, machine learning already helps us in our everyday lives. But it can also help us to tackle some of the world's most challenging physical problems - such as energy consumption. Large-scale commercial and industrial systems like data centres consume a lot of energy, and while much has been done to [slow the growth of energy use](#), there remains a lot more to do given the world's increasing need for computing power.

Agents and Environments



Agents include humans, robots, softbots, thermostats,....

- The agent function maps percepts to actions $f:P \rightarrow A$
- The agent program runs on the physical architecture to produce f

Rational Agents

- Recall that a rational agent “does the right thing”
- We need some way of figuring out what “the right thing is”
 - Performance measures/success criteria
 - This could be goal achievement, resource consumption, loss minimization,...
 - A rational agent chooses which ever action that maximizes the **expected value of its performance measures given the percept sequence to date**
 - You need to know the performance measure, environment, actions, percept sequence
- Rationality is not omniscience, perfection, or success.
- Rationality involves exploration, learning, autonomy

PEAS

- Specify the task environment
 - Performance measure, Environment, Actuators, Sensors

Example: Autonomous Taxi

- Performance Measure: Safety, destination, legality,...
- Environment: Streets, traffic, pedestrians, weather,...
- Actuators: Steering, brakes, accelerator, horn,...
- Sensors: GPS, engine sensors, lidar systems, video,...

Properties of the Task Environment

Properties of the task environment determine what AI approach is most suitable

- Fully Observable vs **Partially Observable**:
 - Does the agent's sensors give it access to the complete state of the environment?
- Deterministic vs **Stochastic**
 - Is the next state completely determined by the current state and action executed?
- Episodic vs **Dynamic**
 - Does the current decision or action influence future decisions or actions?
- Discrete vs **Continuous**
 - How are states, time, actions modelled?
- Static vs **Dynamic**
 - Is the environment changing as the agent is planning what to do next?
- Single Agent vs **Multiagent**

Properties of the Task Environment

Solitaire	Computer Go	Recommender Systems	Autonomous Driving
Fully Observable	Fully Observable	Partially Observable	Partially Observable
Deterministic	Deterministic	Stochastic	Stochastic
Sequential	Sequential	Episodic	Sequential
Static	Static	Dynamic	Dynamic
Discrete	Discrete	Discrete	Continuous
Single agent	Multi agent	Multi agent	Multi agent