

Lecture 1: Course Introduction

CS486/686 Intro to Artificial Intelligence

2023-5-9

Pascal Poupart
David R. Cheriton School of Computer Science



Outline

- Introduction to Artificial Intelligence
- Course website and logistics

Instructors

Pascal Poupart
(Professor and CIFAR AI Chair)



Sriram Ganapathi Subramanian
(Postdoc Vector Institute)



Artificial Intelligence (AI)

- What is **AI**?

Artificial Intelligence (AI)

- What is **AI**?
- What is **intelligence**?

Artificial Intelligence (AI)

- What is **AI**?
- What is **intelligence**?

Webster says: a. the capacity to acquire and apply knowledge. b. the faculty of thought and reason.

Artificial Intelligence (AI)

- What is **AI**?
- What is **intelligence**?
- What features/abilities do humans (animals? animate objects?) have that are indicative or characteristic of intelligence?

Webster says: a. the capacity to acquire and apply knowledge. b. the faculty of thought and reason.

Artificial Intelligence (AI)

- What is **AI**?
- What is **intelligence**?
- What features/abilities do humans (animals? animate objects?) have that are indicative or characteristic of intelligence?
- *abstract concepts, mathematics, language, problem solving, memory, logical reasoning, emotions, morality, ability to learn/adapt, etc...*

Webster says: a. the capacity to acquire and apply knowledge. b. the faculty of thought and reason.

Some Definitions (Russell & Norvig)

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

Some Definitions (Russell & Norvig)

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

What is AI?

- **Systems that think like humans**
 - Cognitive science
 - Fascinating area, but we will not be covering it in this course


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 are the correct thought processes
 - Systems that reason in a logical manner
 - Systems doing inference correctly

What is AI?

- **Systems that behave like humans**

- Turing (1950) “Computing machinery and intelligence”
- 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 in the following 50 years
- Suggested major components of AI: knowledge, reasoning, language understanding, learning



What is AI?

- **Systems that act rationally**
 - Rational behavior: “doing the right thing”
 - Rational agent approach
 - Agent: entity that perceives and acts
 - Rational agent: acts so to achieve best outcome

What is AI?


- **Systems that act rationally**
 - Rational behavior: “doing the right thing”
 - Rational agent approach
 - Agent: entity that perceives and acts
 - Rational agent: acts so to achieve best outcome
- This is the approach we will take in this course
 - General principles of rational agents
 - Components for constructing rational agents

Topics we will cover

- **Search**
 - Uninformed and heuristic search
 - Constraint satisfaction problems
- **Reasoning under uncertainty**
 - Probability theory, utility theory and decision theory
 - Probabilistic inference, causal inference
 - Bayesian networks, decision networks, Markov decision processes
- **Learning**
 - Decision trees, statistical learning, neural networks,
 - Reinforcement learning, bandits, causal learning
- **Multiagent systems**
 - Game theory, multi-agent reinforcement learning


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




wavelab.uwaterloo.ca

Reasoning Under Uncertainty



Machine Learning

- Traditional computer science
 - Program computer for every task
- New paradigm
 - Provide examples to machine
 - Machine learns to accomplish tasks based on examples



Economic growth has slowed down in recent years .



Das Wirtschaftswachstum hat sich in den letzten Jahren verlangsamt .
Economic growth has slowed down in recent years .




La croissance économique s' est ralentie ces dernières années .

Three Categories



Supervised learning




Reinforcement learning



Unsupervised learning

Reinforcement Learning Problem



Goal: Learn to choose actions that maximize rewards

Animal Psychology


- Negative reinforcements:
 - Pain and hunger
- Positive reinforcements:
 - Pleasure and food
- Reinforcements used to train animals

- **Let's do the same with computers!**



Game Playing

- Example: Go (one of the oldest and hardest board games)
- **Agent:** player
- **Environment:** opponent
- **State:** board configuration
- **Action:** next stone location
- **Reward:** +1 win / -1 loose



2016: AlphaGo defeats top player Lee Sedol (4-1)
Game 2 move 37: AlphaGo plays unexpected move (odds 1/10,000)

A brief history of AI

- 1943-1955: Initial work in AI
 - McCulloch and Pitts produce Boolean model of the brain
 - Turing's "Computing machinery and intelligence"
- Early 1950's: Early AI programs
 - Samuel's checker program, Newell and Simon's Logic Theorist, Gerlenter's Geometry Engine
- 1956: Happy birthday AI!
 - Dartmouth workshop attended by McCarthy, Minsky, Shannon, Rochester, Samuel, Solomonoff, Selfridge, Simon and Newell


A brief history of AI

- 1950's-1969: Enthusiasm and expectations
 - Many successes (in a limited way)
 - LISP, time sharing, resolution method, neural networks, vision, planning, learning theory, Shakey, machine translation,...
- 1966-1973: Reality hits
 - Early programs had little knowledge of their subject matter
 - Machine translation
 - Computational complexity
 - Negative result about perceptrons - a simple form of neural network

A brief history of AI

- 1969-1979: Knowledge-based systems
- 1980-1988: Expert system industry booms
- 1988-1993: Expert system busts, AI Winter
- 1986-2000: The return of neural networks
- 2000-present: Increase in technical depth
 - Probabilities, statistics, optimization, utility theory, game theory, learning theory
- 2010-present: Big data, deep neural networks

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: a rational agent “does the right thing”
- Performance measure – success criteria
 - Evaluates a sequence of environment states
- A **rational agent** chooses whichever action that maximizes the **expected** value of its performance measure **given the percept sequence to date**
 - Need to know performance measure, environment, actions, percept sequence
- Rationality \neq omniscience, perfection, success
- Rationality \rightarrow exploration, learning, autonomy

PEAS

- Specify the **task environment**:
 - **P**erformance measure, **E**nvironment, **A**ctuators, **S**ensors

Example: Autonomous Taxi

Performance Measure: Safety, destination, legality...

Environment: Streets, traffic, pedestrians, weather...

Actuators: Steering, brakes, accelerator, horn...

Sensors: GPS, engine sensors, video...

Properties of task environments

- Fully observable vs. partially observable
- Deterministic vs. stochastic
- Episodic vs. sequential
- Static vs. dynamic
- Discrete vs. continuous
- Single agent vs. multiagent

Hardest case: Partially observable, stochastic, sequential, dynamic, continuous and multiagent. (Real world)

Examples

Solitaire	Computer Go	Recommender system	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	Multiagent	Multiagent	Multiagent

Many Applications

- credit card fraud detection
- medical assistive technologies
- information retrieval, question answering, conversational agents
- speech recognition, natural language processing
- scheduling, logistics, etc.
- aircraft, pipeline inspection
- Mars rovers, driverless cars
- and, of course, cool robots