

CS440/ ECE 448 Lecture 1: Introduction to AI

Mark Hasegawa-Johnson, 1/2024

Lecture slides: CC0



Public domain image, https://commons.wikimedia.org/wiki/File:Robonaut_and_astronaut_hand_shake.jpg

Outline

- What is Artificial Intelligence?
 - Human-like? Rational? Predictive? Conscious?
 - Seven things an AI should be able to do
 - Environments in which an AI can operate
- Syllabus
 - Text
 - Web Page, Office Hours, and CampusWire
 - Grades: Quizzes, MPs, Exams, and Project
 - Lectures

What Is Artificial Intelligence?

- Human-like: Is it able to communicate with human beings, and explain or demonstrate its reasoning to them?
- Rational: Does it always act in a manner that maximizes its expected performance metric?
- Predictive: Is it capable of predicting the future?
- Conscious: Does it have an internal model of itself as an actor in all of its memories of the past and predictions of the future?

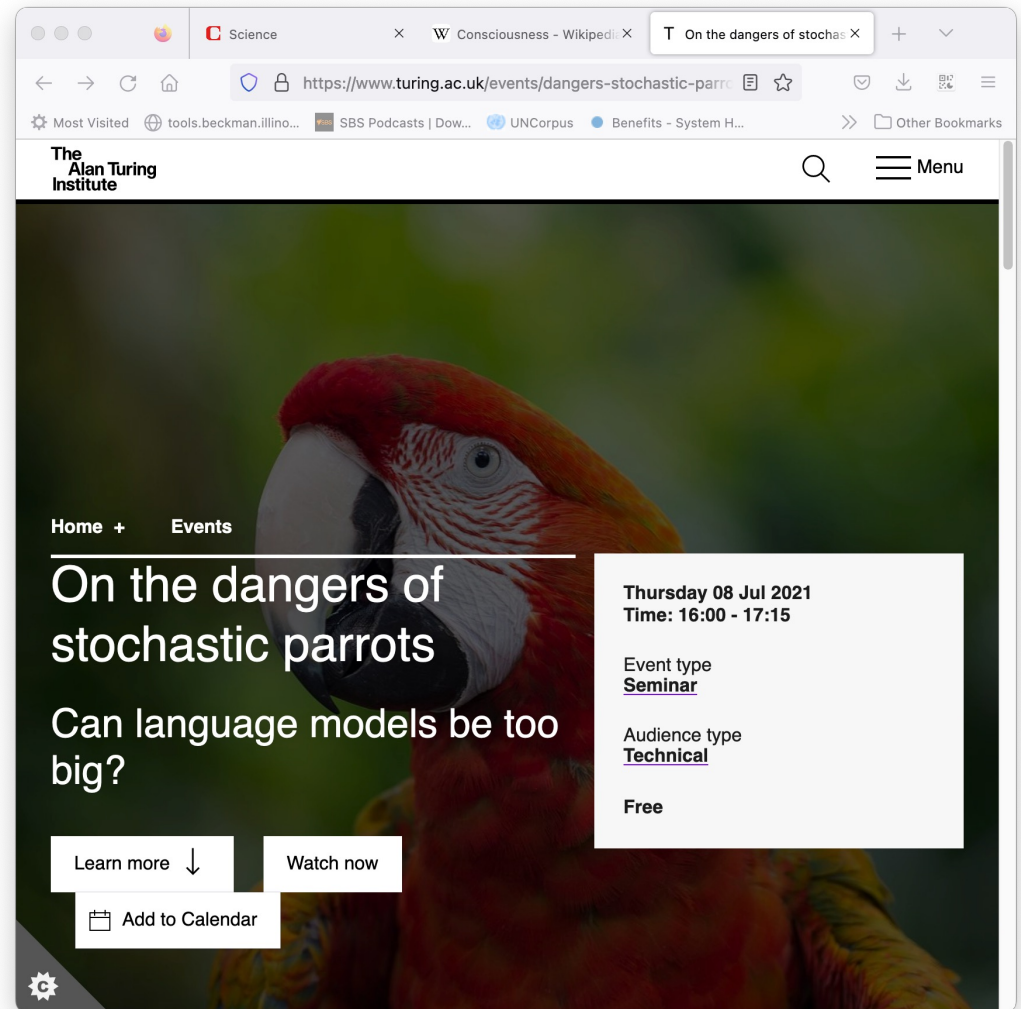
Human-likeness

The Turing test proposes that an AI is intelligent if a human interviewer can't tell whether it is human.

Modern AI routinely fools humans.

It does so by performing as a "Stochastic Parrot:" given a prompt, the AI repeats what a well-read human would have said in response to that prompt.

Is that intelligence?



The screenshot shows a web browser window with the URL <https://www.turing.ac.uk/events/dangers-stochastic-parrots>. The page is for an event titled "On the dangers of stochastic parrots" at The Alan Turing Institute. The event is scheduled for Thursday 08 Jul 2021, from 16:00 to 17:15. It is a free seminar with a technical audience type. The main text on the page asks "Can language models be too big?". There are buttons for "Learn more", "Watch now", and "Add to Calendar". The background of the page features a close-up image of a red and yellow macaw parrot.

Ad from the Turing Institute for a lecture by Emily Menon Bender about the article she co-wrote with Gebru and McMillan-Major.

Rationality

- It has been argued that the Turing test is too human-centric. How would we know if an Elephant is intelligent? An Extraterrestrial? A Robot?
- The most commonly proposed alternative is **rationality**: the quality of being guided by reasons.



CC-BY 2.5, <https://commons.wikimedia.org/wiki/File:Loxodontacyclotis.jpg>

Human-likeness + Rationality = Rational Prediction?


- There is a growing consensus among AI experts for the following rough definition of intelligence, which combines the most important attributes of human-likeness and rationality:

Intelligence is the ability to consistently predict the future, with low error, based on rational inference from past events.

- By this definition, a large language model is intelligent within its own domain (the domain of predicting the next word of text) and could be made into a general intelligence by simply generalizing its domain.
- Although philosophers increasingly agree with this definition, it is unsatisfying to many people. Why?

By “intelligent,” do we mean “conscious”?

- Science fiction leads us to think that an “intelligent” agent should be one that is self-aware in the way we are: conscious.
- ... but philosophers increasingly argue that an agent can be intelligent without being conscious, and vice versa, and
- ... the most scientifically defensible definitions of consciousness leave open the possibility that even if an AI were conscious, it might be very different than we are,
- ... e.g., in the most defensible theories, consciousness does not necessarily imply desires of any kind, not even wishes concerning the length of one’s own future life.



Frankenstein, 1831 edition. Public domain image.

Outline

- What is Artificial Intelligence?
 - Human-like? Rational? Autonomous? Conscious?
 - Seven things an AI should be able to do
 - Environments in which an AI can operate
- Syllabus
 - Text
 - Web Page, Office Hours, and CampusWire
 - Grades: Quizzes, MPs, Exams, and Project
 - Lectures

Seven things an AI should be able to do

Without any need for consciousness...

- Make rational (reason-guided) decisions
- Learn
- Plan (solve problems)
- Understand what it has learned (make inferences)
- Communicate using natural language
- Perceive its environment
- Act on its environment

Outline of this Course

Topics for this semester will roughly follow the “seven things an AI should be able to do.”

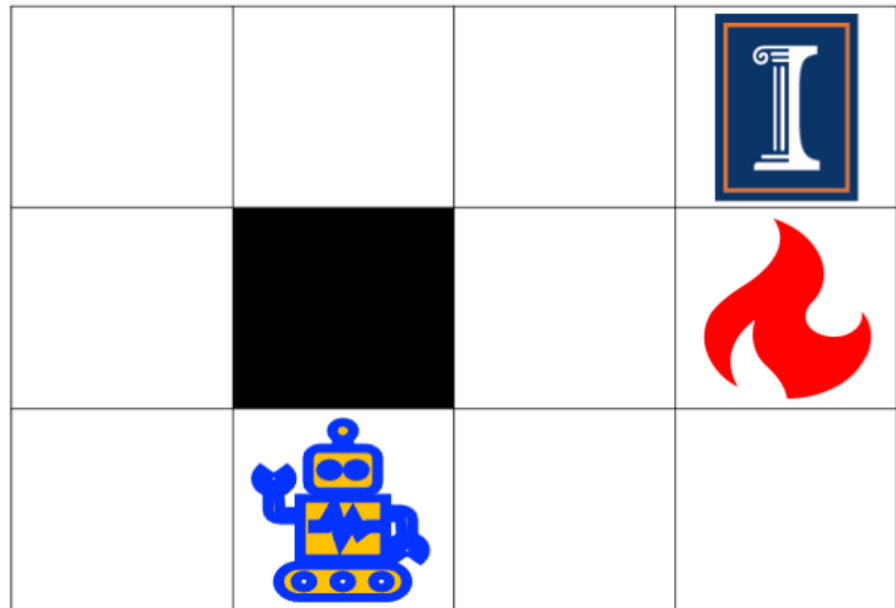
Topic	Week	Day	Quiz	Slides	MP	Reading
Intro	1	W	17-Jan	Welcome & Intro		1.5
Deciding		F	19-Jan	Random Variables		9.1-2
	2	M	22-Jan	Decision Theory		7.2.3;7.4.3
		W	24-Jan	Naive Bayes	Probability	10.2.2;9.6.6
Communicating		F	26-Jan	Bayes Networks		
	3	M	29-Jan	HMM		
		W	31-Jan	Fairness	Naive Bayes	
Learning		F	2-Feb	Learning		
	4	M	5-Feb	Linear Regression		
		W	7-Feb	Linear Classifiers	HMM	
		F	9-Feb	Multilayer Networks		
	5	M	12-Feb	Optimization		
		W	14-Feb	Pytorch Tutorial	Perceptron	
		F	16-Feb	Exam 1 Review		
		M	19-Feb	Exam 1		
Perceiving		W	21-Feb	Computer Vision		
		F	23-Feb	ConvNets & RNN		
Planning	7	M	26-Feb	Explainable AI		
		W	28-Feb	Search	Neural Nets	
		F	1-Mar	A* Search		
	8	M	4-Mar	Markov Decisions		
		W	6-Mar	Minimax	Search	
		F	8-Mar	Expectiminimax		
		M	11-Mar	Spring Break		
		W	13-Mar	Spring Break		
		F	15-Mar	Spring Break		
		M	18-Mar	No Class		
		W	20-Mar	No Class		
		F	22-Mar	Game Theory		
9	M	25-Mar	Iterated Games			

Outline

- What is Artificial Intelligence?
 - Human-like? Rational? Autonomous? Conscious?
 - Seven things an AI should be able to do
 - Environments in which an AI can operate
- Syllabus
 - Text
 - Web Page, Office Hours, and CampusWire
 - Grades: Quizzes, MPs, Exams, and Project
 - Lectures

The abilities of an AI need to be matched to the properties of its environment

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



Fully observable vs. Partially observable

- Do the agent's sensors give it access to the complete state of the environment?
 - For any given world state, are the values of all the variables known to the agent?



VS.



Source: L. Zettlemoyer

Deterministic vs. Stochastic

- Is the next state of the environment completely determined by the **current state** and the **agent's action**?
 - Is the transition model **deterministic** (unique successor state given current state and action) or **stochastic** (distribution over successor states given current state and action)?
 - **strategic**: the environment is deterministic except for the actions of other agents



VS.




Episodic vs. Sequential

- Is the agent's experience divided into unconnected episodes, or is it a coherent sequence of observations and actions?
 - Does each problem instance involve just one action or a series of actions that change the world state according to the transition model?



VS.



Static vs. Dynamic

- Is the world changing while the agent is thinking?




vs.




Discrete vs. Continuous

- Does the environment provide a countable (discrete) or uncountably infinite (continuous) number of distinct percepts, actions, and environment states?
 - Are the values of the state variables discrete or continuous?
 - Time can also evolve in a discrete or continuous fashion
 - “Distinct” = different values of utility



VS.




Single-agent vs. Multi-agent

- Is an agent operating by itself in the environment?



vs.



Known vs. Unknown

- Are the rules of the environment (transition model and rewards associated with states) known to the agent?
 - Strictly speaking, not a property of the environment, but of the agent's state of knowledge



vs.



Quiz question

Go to https://us.prairielearn.com/pl/course_instance/147925/assessment/2390712

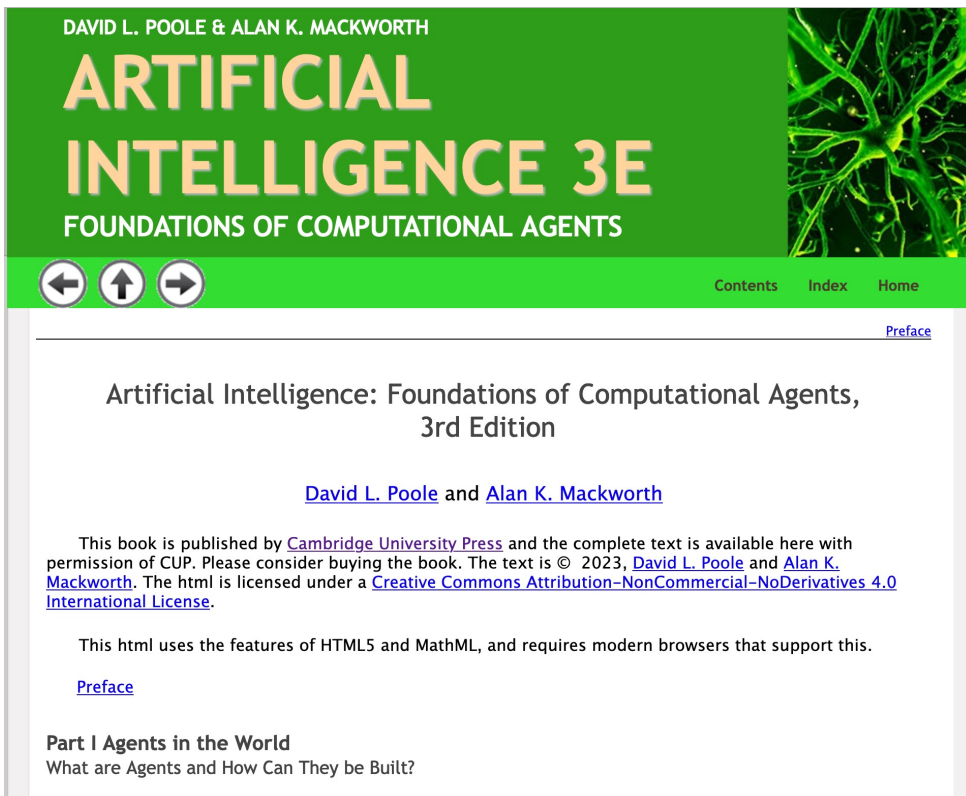
Join the course

Take the quiz called “17-Jan”

Outline

- What is Artificial Intelligence?
 - Human-like? Rational? Autonomous? Conscious?
 - Seven things an AI should be able to do
 - Environments in which an AI can operate
- Syllabus
 - Text
 - Web Page, Office Hours, and CampusWire
 - Grades: Quizzes, MPs, Exams, and Project
 - Lectures

Textbook



Artificial Intelligence 3E by Mackworth & Poole

- Readings will be specified for each lecture.
- Readings are optional; required material is on slides, quizzes & MPs
- Textbook is free online, so you may find its coverage of some material helpful

Webpage: <https://courses.grainger.illinois.edu/ece448>

The screenshot shows a web browser window with the URL <https://courses.grainger.illinois.edu/ece448/sp2023/>. The page title is "CS440/ECE448 Artificial Intelligence, Spring 2023". The main content area features a 3x3 grid icon with a blue 'I' and a red flame, a "Table of Contents" section with links to "On-line Tools" and "Textbook", a "Next topic" section for "Grades", a "This Page" section with a "Show Source" link, and a "Quick search" bar. The main text describes the course as an introductory survey of AI concepts and techniques, listing prerequisites (CS225 and ECE 313) and a list of contents including Grades, Quizzes, Machine Problems, Exams, Project, Office Hours, and Links. Below this are sections for "On-line Tools" (listing CampusWire, PrairieLearn, Gradescope, and Mediaspace) and "Textbook" (listing Russell and Norvig's "Artificial Intelligence: A Modern Approach").

CS440/ECE448 Artificial Intelligence, Spring 2023

CS440/ECE448 Artificial Intelligence, Spring 2023

This course provides an introductory survey of concepts and techniques in artificial intelligence. Intelligence is the ability to decide, learn, plan, understand, communicate, perceive, and act; AI is the creation of machines that do these things. This course gives a one- or two-week introduction to each of the seven aspects of intelligence, with pointers to key mathematical concepts, articles and books.

This course assumes that you have taken data structures (CS225) and random variables (e.g., CS 360, ECE 313).

Contents:

- Grades
- Quizzes, Machine Problems, and Lectures
- Exams
- Project
- Office Hours
- Links

On-line Tools

- [CampusWire](#) will be used for on-line question answering. The enrollment code is available in the slides for lecture 1.
- [PrairieLearn](#) will be used for quizzes.
- [Gradescope](#) will be used for machine problems. The code to join Gradescope is on CampusWire.
- You can watch the lectures remotely on [Mediaspace](#). Videos are usually available two hours after the end of lecture.

Textbook

The textbook is Russell and Norvig, [Artificial Intelligence: A Modern Approach](#), fourth edition.

You will also need a reference for Python, which we will use to write the MPs. A good place to start is the [Python Tutorial](#).

CampusWire

Add yourself to CampusWire if you're not already added:

<https://campuswire.com/p/G5603CE97>

Code 4442.

Class feed

CS 440/ECE 448: Artificial Intelligence

All categories



+ New post

Pinned



Welcome to CS 440/ECE 448!

#1

Welcome to CS 440/ECE 448! The c...



0 a day ago

This week



MP extra credit grading question

#2

I have a few questions about MP ext...



1 19 hours ago



Mark Hasegawa-Johnson posted a note a day ago (edited)

Visible to: Everyone

Welcome to CS 440/ECE 448! #1

General

Welcome to CS 440/ECE 448!

The course website is <https://courses.grainger.illinois.edu/ece448/sp2024/>.

The Gradescope site is <https://www.gradescope.com/courses/701752>. The registration code is 6GN3K3.

0 0 619 220

Comments

Grading: Quizzes, MPs, Exams, Project

- Quizzes: 15% of 3-credit grade
 - Every lecture will have a quiz. Do it in class if you can.
 - Due: 23 hours after the end of lecture.
- Machine Problems: 45% of 3-credit grade
 - Every week will have an MP (11 in total)
 - Due: Every Wednesday, at 1:00pm, starting NEXT WEDNESDAY
- Exams: 40% of 3-credit grade
 - Will be held in person
- Literature Review: 100% of the 4th credit
 - Ten quizzes on ten extra articles throughout the semester

Late Policy

- Quizzes, MPs, and Project components may be turned in late for partial credit:

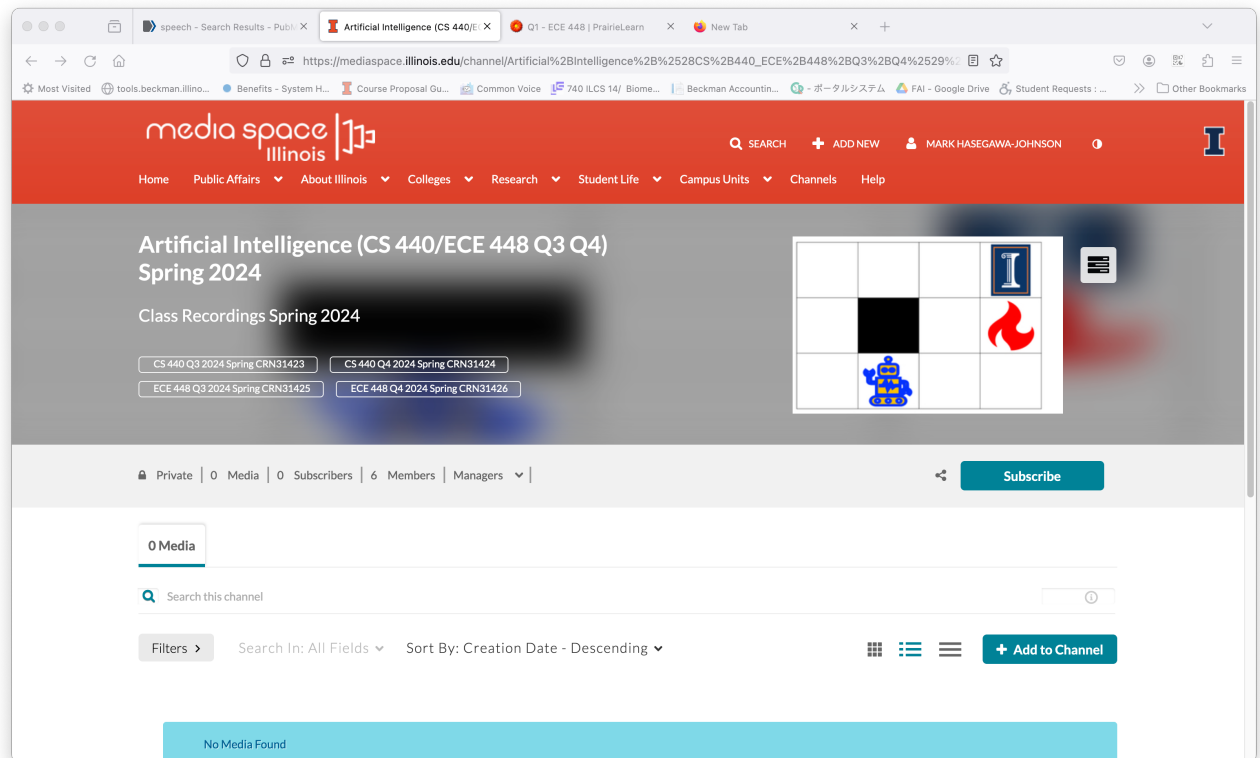
$$\max\left(1 - \frac{t}{20}, 0.5\right)$$

where t is the lateness, in days. This policy is intentionally lenient: if you get sick, you can still turn in your homework late for most of the credit.

- Further exemptions from this late policy are not granted for illness, travel, or any other reason.

Lectures

- Lectures are MWF, 1pm, Lincoln Hall Theater
- Lecture recordings are automatically posted, about 6 hours later, on MediaSpace



Welcome to Artificial Intelligence!

- Get started on MP01
- See you on Friday!



iCub Production Lab,
[https://commons.wikimedia.org/wiki/File:P058324-119830_\(cropped\).jpg](https://commons.wikimedia.org/wiki/File:P058324-119830_(cropped).jpg)