CS440/ECE 448 Lecture 1: Introduction to AI

Mark Hasegawa-Johnson, 1/2023

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

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?
- Autonomous: Is it capable of revising its sensor→action mapping in response to changes in the environment?
- Conscious: Does it sense, feel, and know the history and present of all of its relationships to other people and to the universe as a whole?

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?

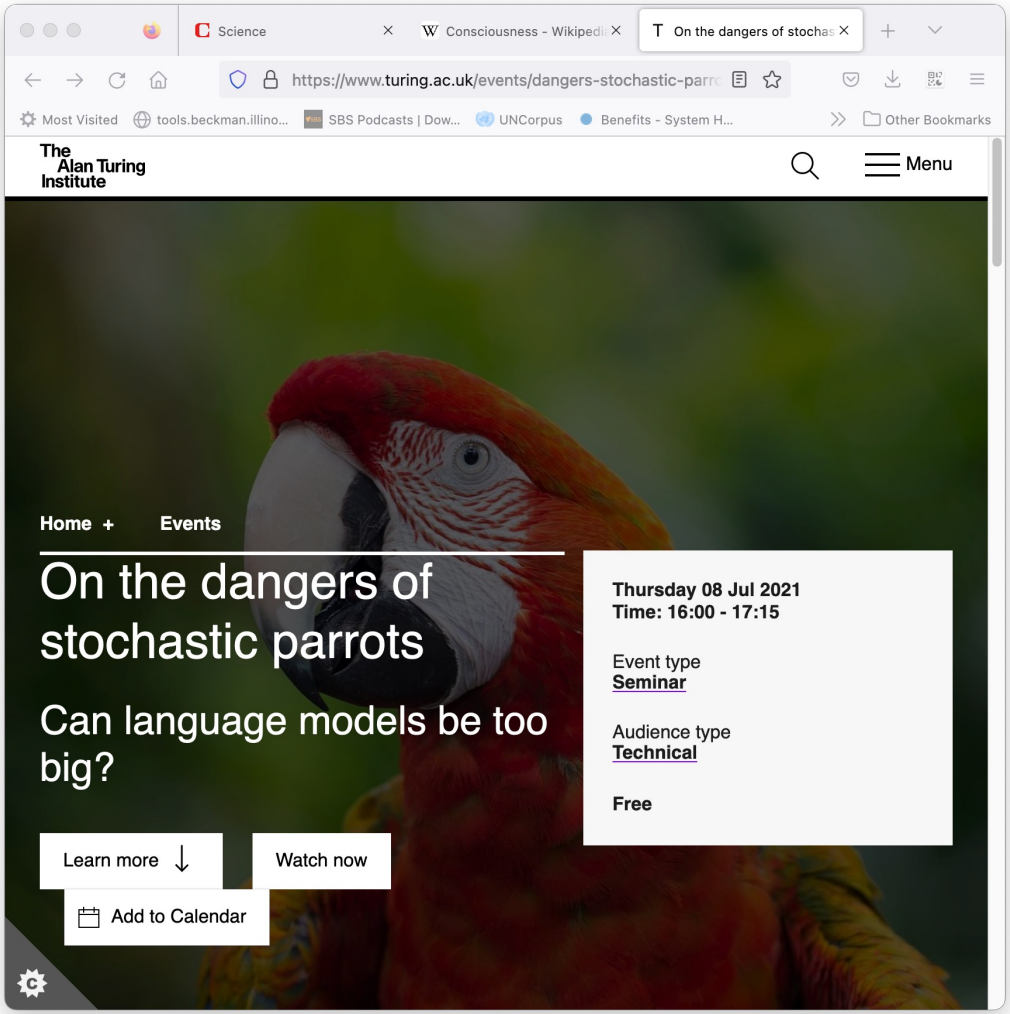

Ad from the Turing Institute for a lecture by Emily Menon Bender about the article she co-wrote with Gebru, McMillan-Major, and Gebru.

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.

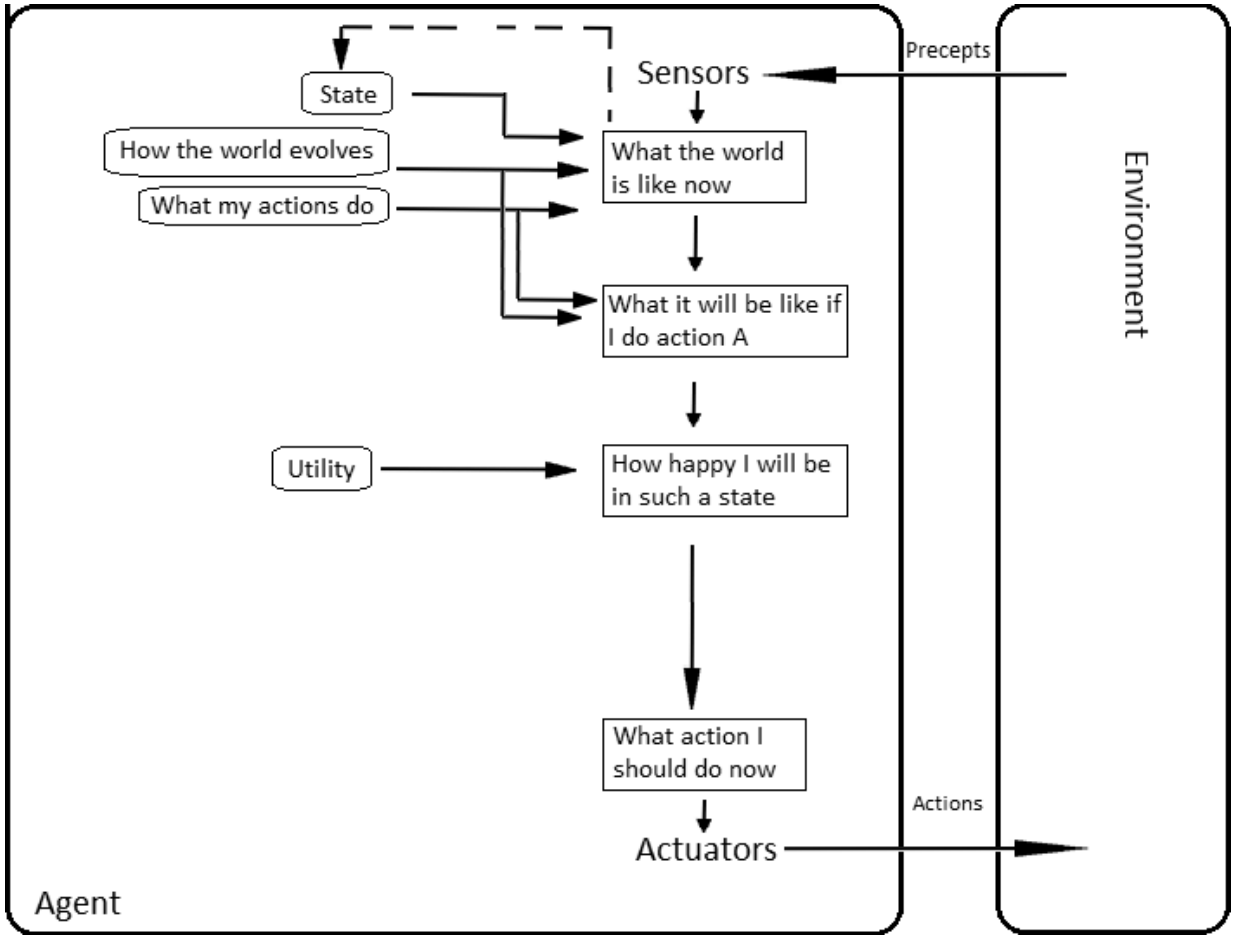

Autonomy

An agent is “autonomous” if it is capable of revising its own behavior in response to changes in the environment.

The quadrupedal military robot Cheetah, an evolution of BigDog (pictured), was clocked as the world's fastest legged robot in 2012. Open-source image, DARPA strategic plan 2007.

But do those things make it “intelligent”?

If an agent is:

- Human-like,
- Rational, and
- Autonomous…

Does that make it intelligent?

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 we currently have no way to test whether a given agent is conscious.

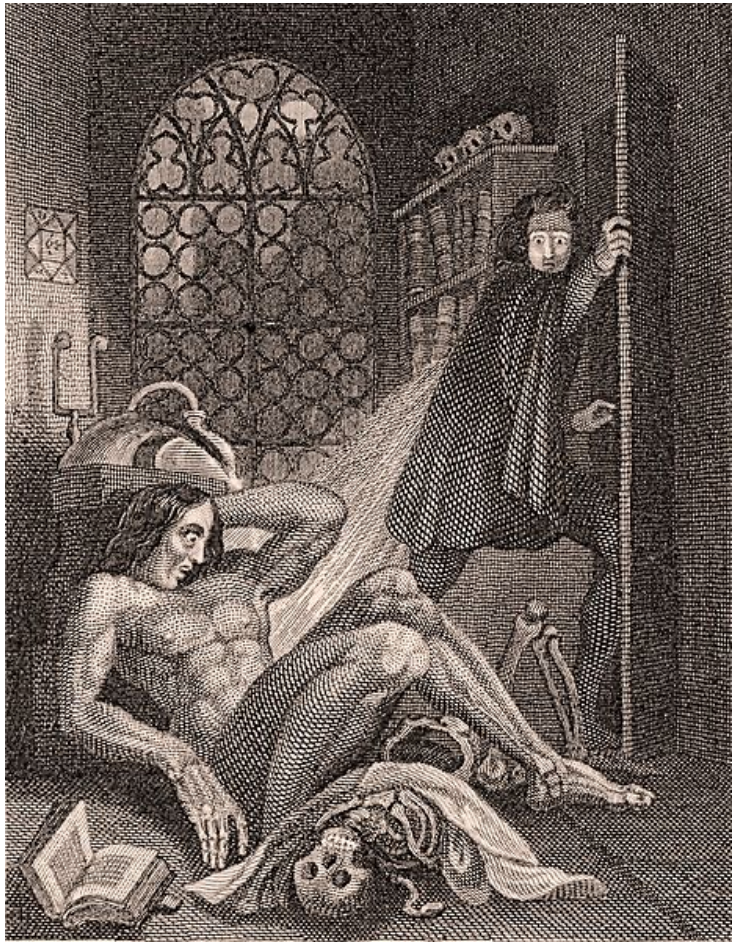

Frankenstein, 1831 edition. Public domain image.

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	18-Jan	Welcome & Intro		2.1-5
Deciding		F	20-Jan	Random Variables		12.1-7
	2	M	23-Jan	Decision Theory		16.1-4
		W	25-Jan	Naïve Bayes	Probability	20.2.1-2;23.1.1
		F	27-Jan	Game Theory		18.2
	3	M	30-Jan	Fairness		27.3.3
Learning		W	1-Feb	Learning	Naïve Bayes	19.1-5;7
		F	3-Feb	Linear Regression		19.6
	4	M	6-Feb	Linear Classifiers		19.6
		W	8-Feb	Multilayer Networks	KNN	21.1-2
		F	10-Feb	PyTorch		Pytorch tutorial
	5	M	13-Feb	Optimization		4.1-5
		W	15-Feb	Privacy	Neural Nets	27.3.2
		F	17-Feb	Exam 1 Review		
	6	M	20-Feb	Exam 1		
Planning		W	22-Feb	Search		3.1-4
		F	24-Feb	A* Search		3.5-6
	7	M	27-Feb	Minimax		5.1-3
		W	1-Mar	AI Safety	Search	27.3.7
Understanding		F	3-Mar	Logic		7.1-7
	8	M	6-Mar	Ontology		10.1-6
		W	8-Mar	Bayesian Networks	Logic	13.1-3
		F	10-Mar	Transparency		27.3.4
Communicating	9	M	20-Mar	HMM		14.1-3
		W	22-Mar	Parsing	Bayes Nets	23.1-4
		F	24-Mar	DL for NLP		24.1-6
	10	M	27-Mar	Convolutional Networks		21.3
		W	29-Mar	Consciousness	HMM	27.2;27.3.6
		F	31-Mar	Exam 2 Review		
	11	M	3-Apr	Exam 2		
Perceiving		W	5-Apr	Computer Vision		25.1-5
		F	7-Apr	Robot Perception		26.1-4
	12	M	10-Apr	Kalman Filter		14.4
		W	12-Apr	Autonomous Weapons	Perception	27.3.1
Acting		F	14-Apr	MDP		17.1-3
	13	M	17-Apr	Model-Based RL		22.1-2
		W	19-Apr	Model-Free RL	MDP	22.3

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

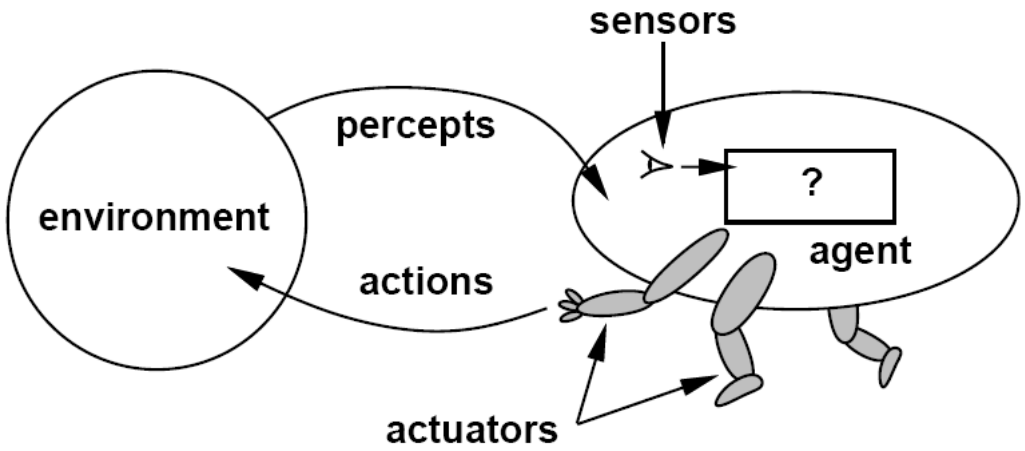

Fully observable vs. Partially observable

- 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?

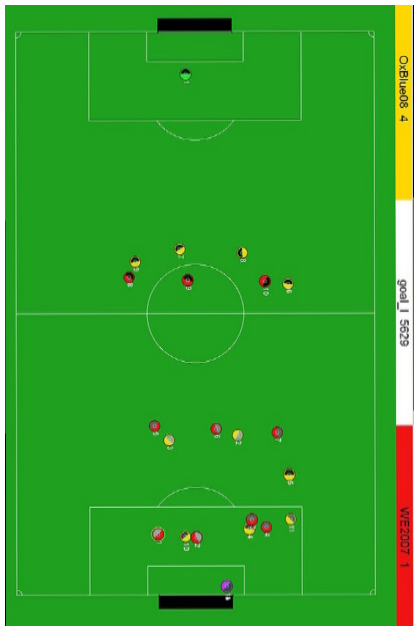

vs.

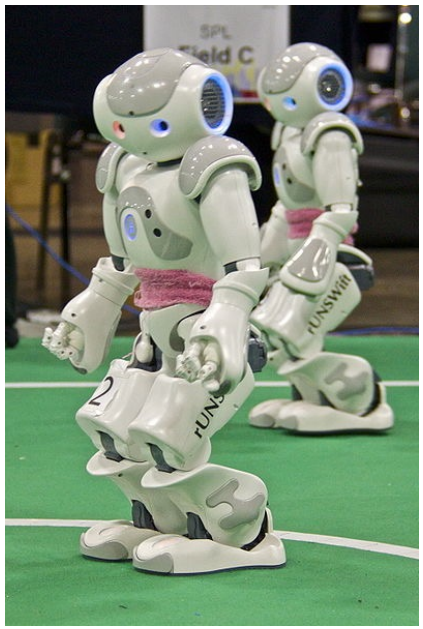

Source: L. Zettlemoyer

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.

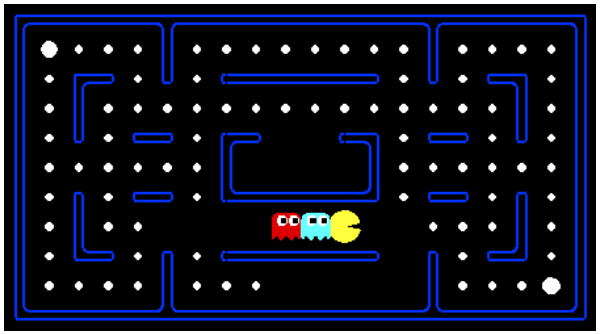

Static vs. Dynamic

- 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.

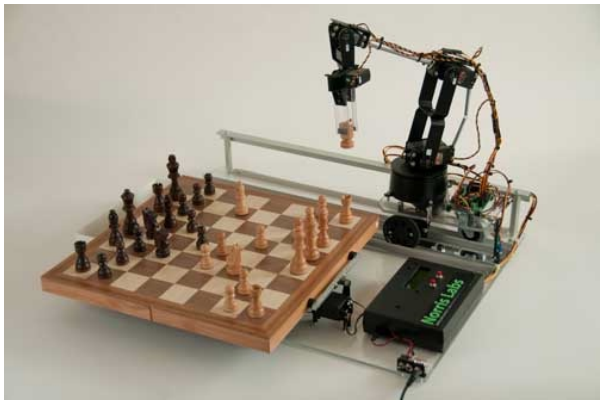

Single-agent vs. Multi-agent

- Is an agent operating by itself in the environment?

vs.

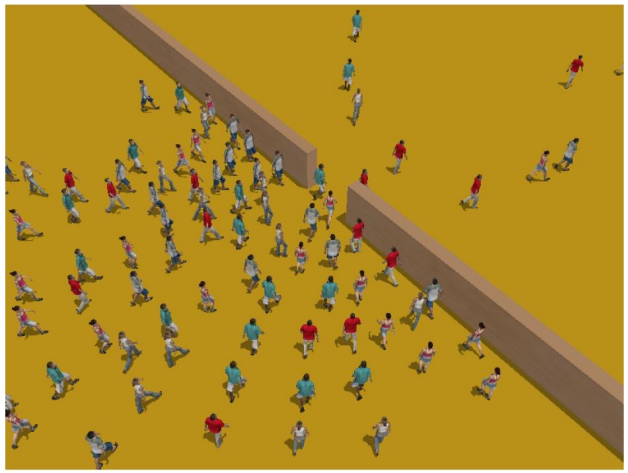

Known vs. Unknown

- 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/129874/

Join the course

Take the quiz called “18-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

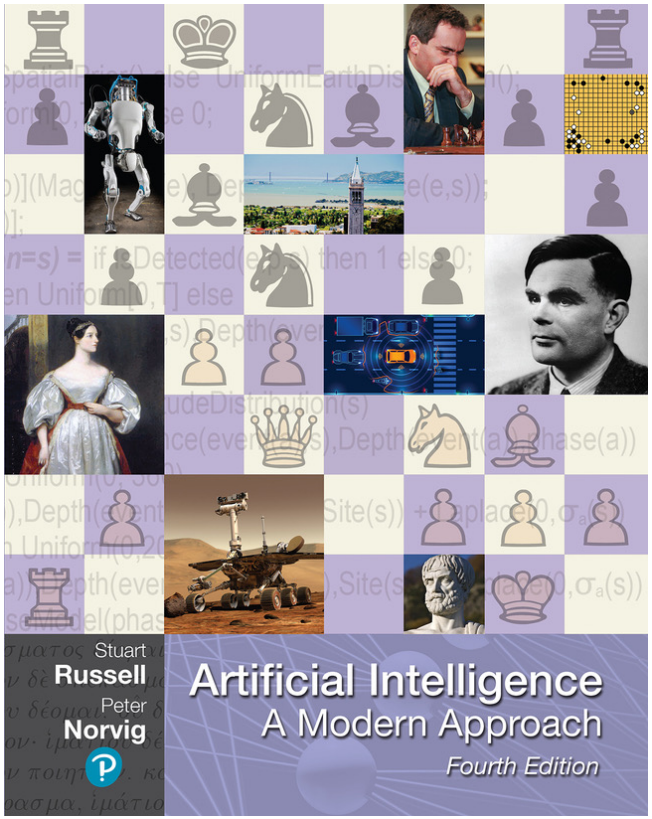

Artificial Intelligence, A Modern Approach: Fourth Edition by Russell & Norvig

- Readings will be specified for each lecture.
- Material will only show up on exams if it has first appeared in the lecture slides, and usually, on a quiz or MP. Textbook is a backup, if you want deeper understanding.

Webpage: https://courses.grainger.Illinois.edu/ece448

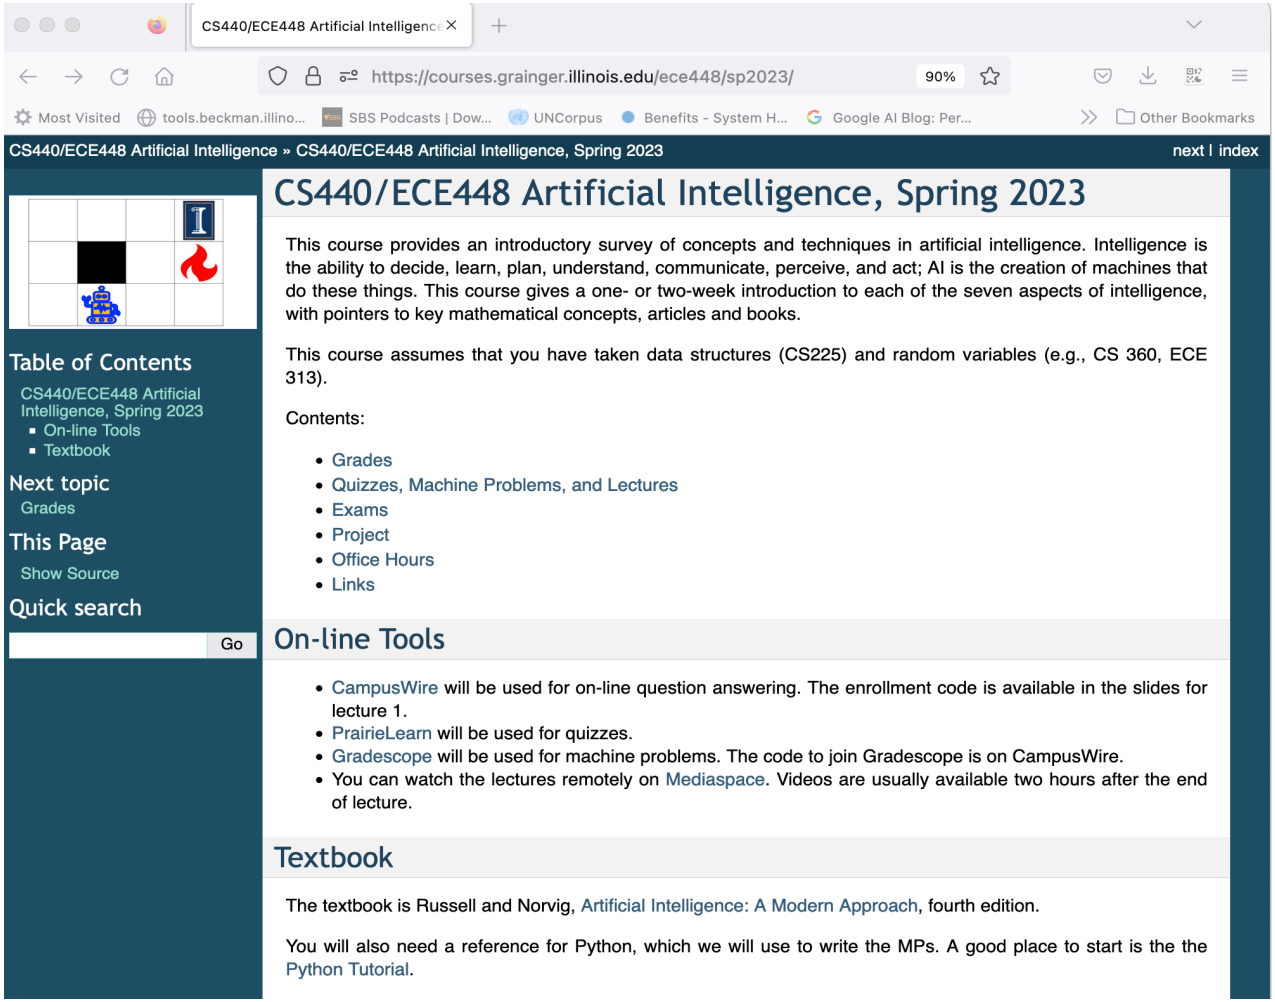

Office Hours

Office hours will be:

Every M,T,W,R

5:00-7:00 pm

Starting this Thursday, 1/19

In-person only

ECEB 2036

CampusWire

Add yourself to CampusWire if you're not already added: https://campuswire.com/c/G2DB4EC0D,

Code 9176.

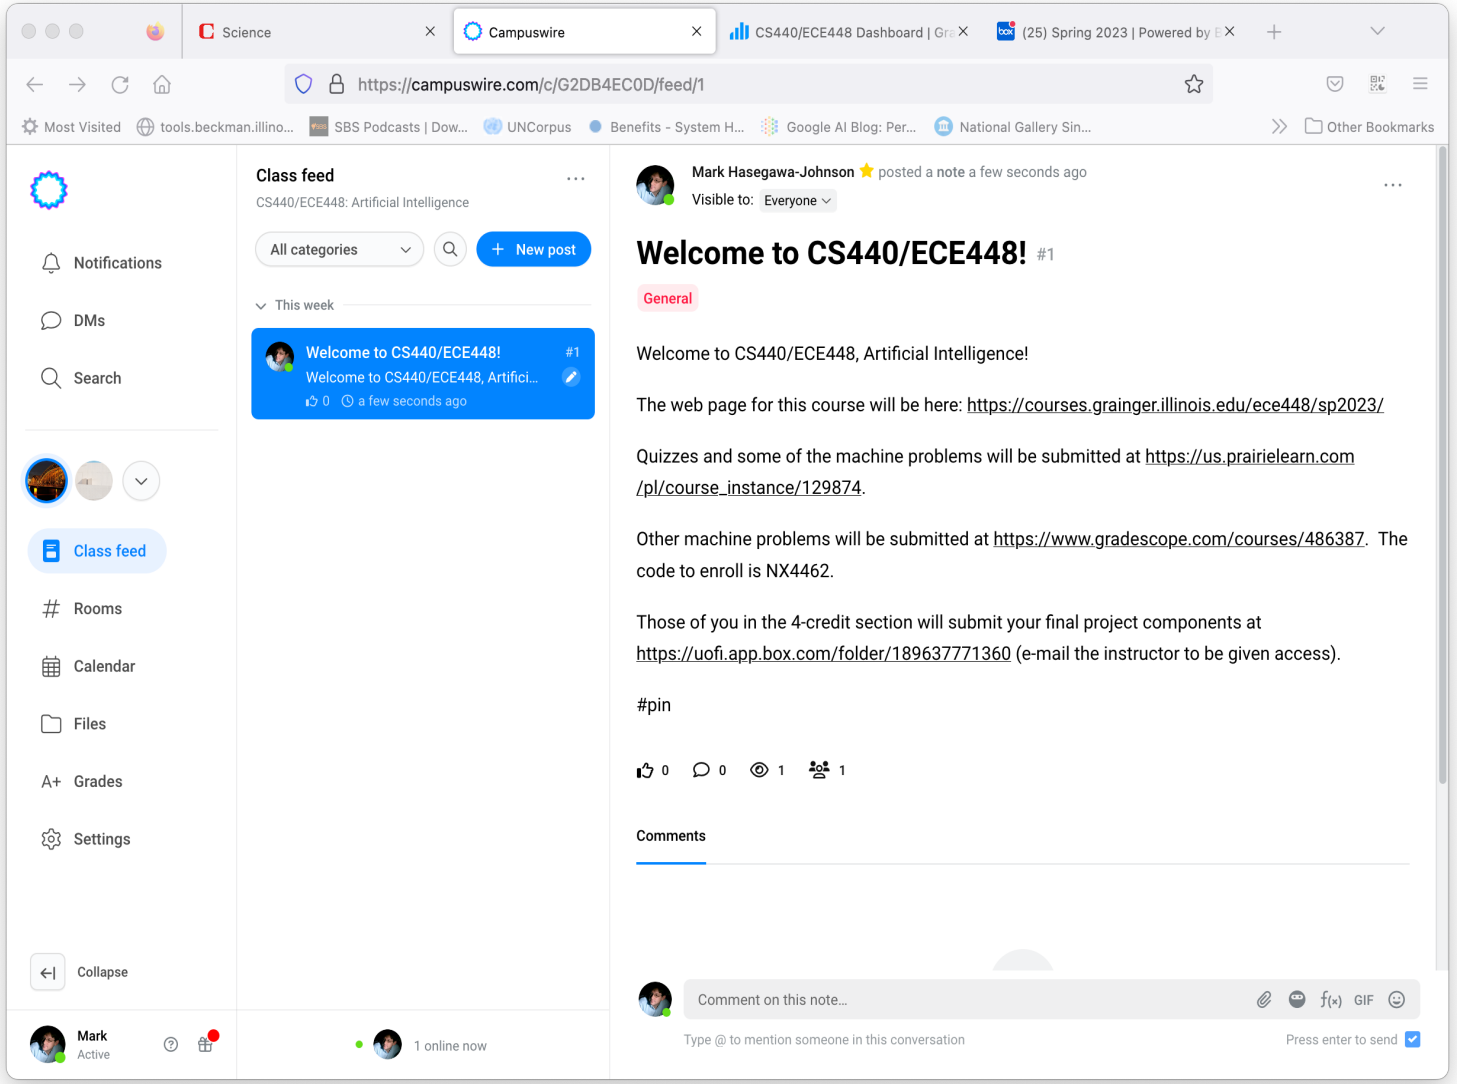

Grading: Quizzes, MPs, Exams, Project

- 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
- Project: 100% of the 4^{th} credit
  - Seven project component deadlines 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

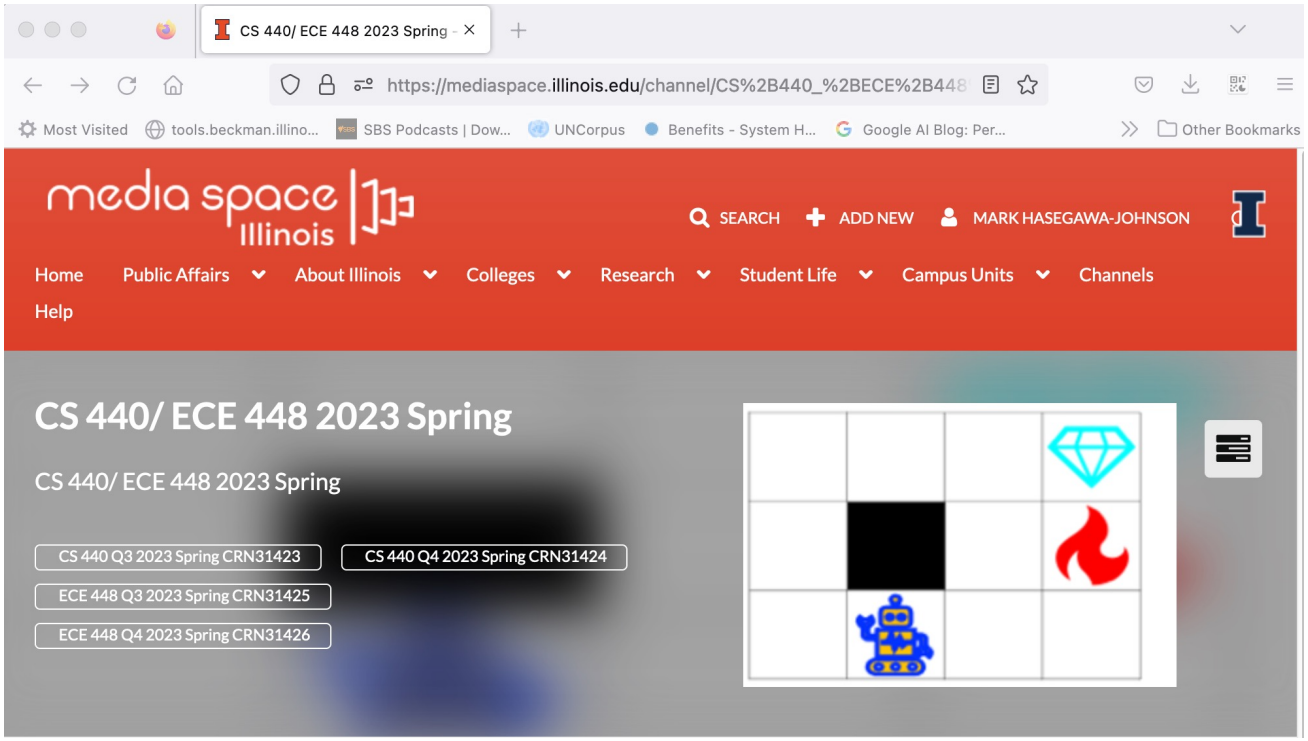

Welcome to Artificial Intelligence!

- Come to office hours tomorrow (1/19), 5pm in ECEB 2036, to meet some of your teaching assistants
- Get started on MP01
- See you on Friday!

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