CS440/ECE 448 Lecture 1: Introduction to AI

Mark Hasegawa-Johnson, 1/2022

Outline

- Syllabus: CoViD, Text, Grades, Quizzes, MPs, Exams, and Project
- A brief history of Artificial Intelligence
- Artificially Intelligent Agents
  - Words to know: Performance, Environment, Actions, Sensors (PEAS)
  - Words to know: Rational, Human-like, Autonomous
- Environments in which an AI can operate
  - Words to know: Observable, Deterministic vs. Stochastic, Episodic vs. Sequential, Static vs. Dynamic, Discrete vs. Continuous, Single-Agent vs. Multi-Agent, Known vs. Unknown
- Outline of the rest of this course

Course Intro: Syllabus

- Web page: https://courses.engr.illinois.edu/ece448/sp2022/
- CampusWire
- CoViD policy
- Textbook
- Grading
- Quizzes, MPs, and Exams
- Graduate Section: Project

CampusWire

Add yourself to CampusWire if you're not already added: https://campuswire.com/p/G7F96162F, Code 9270.

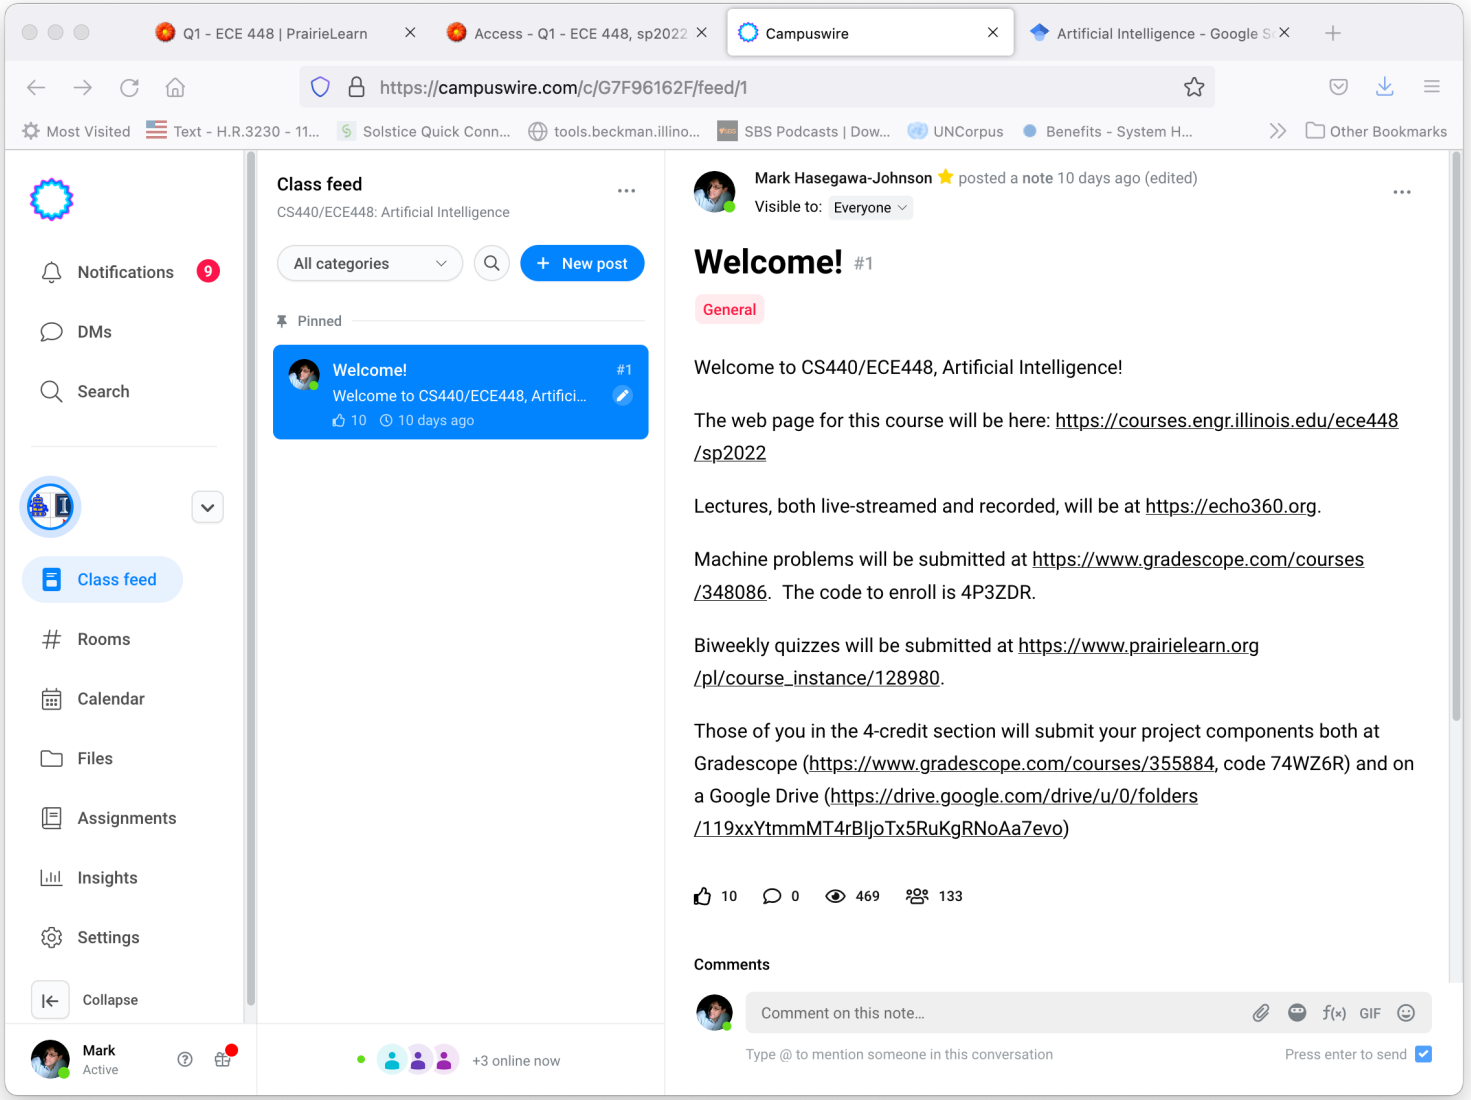

CoViD Policy

- You can watch lectures in any of these three ways, any time:
  - In person, beginning 1/31
  - In real-time, streaming live at echo360.org, beginning 1/31
  - In recordings, available on echo360.org within a few hours after class
- CoViD is NOT usually a valid reason for a quiz or MP deadline extension. All assignments are available for you a week early. If you miss the deadline, our standard late penalty is extremely lenient (5% per day).
- CoViD IS a valid reason for an exam or project extension. If you can take it on the regular exam day (by zoom), please do; if not, we'll arrange a make-up.

Textbook

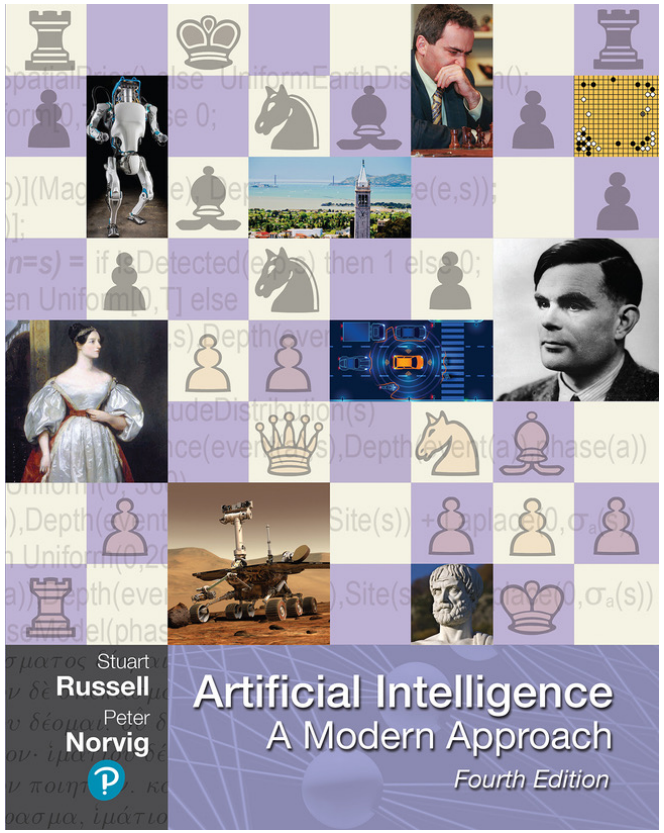

Artificial Intelligence, A Modern Approach: Fourth Edition by Russell & Norvig

- 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.

Grading: 3-credit

- Exams: 40%
- Machine Problems (MPs): 45%
- Quizzes: 15%

Grading: 4-credit

- Exams: 30%
- MPs: 33%
- Quizzes: 12%
- Project: 25%

Quizzes

- Released ten days before due.
- Try each problem as often as you like ("Try a New Variant" button).
- Late penalty 5% per day.
- You may submit any time during the semester for 50% credit.
- Late penalty is not waived for travel, trauma, interviews, religious holidays, or for any illness or injury that leaves you with the use of your hands.

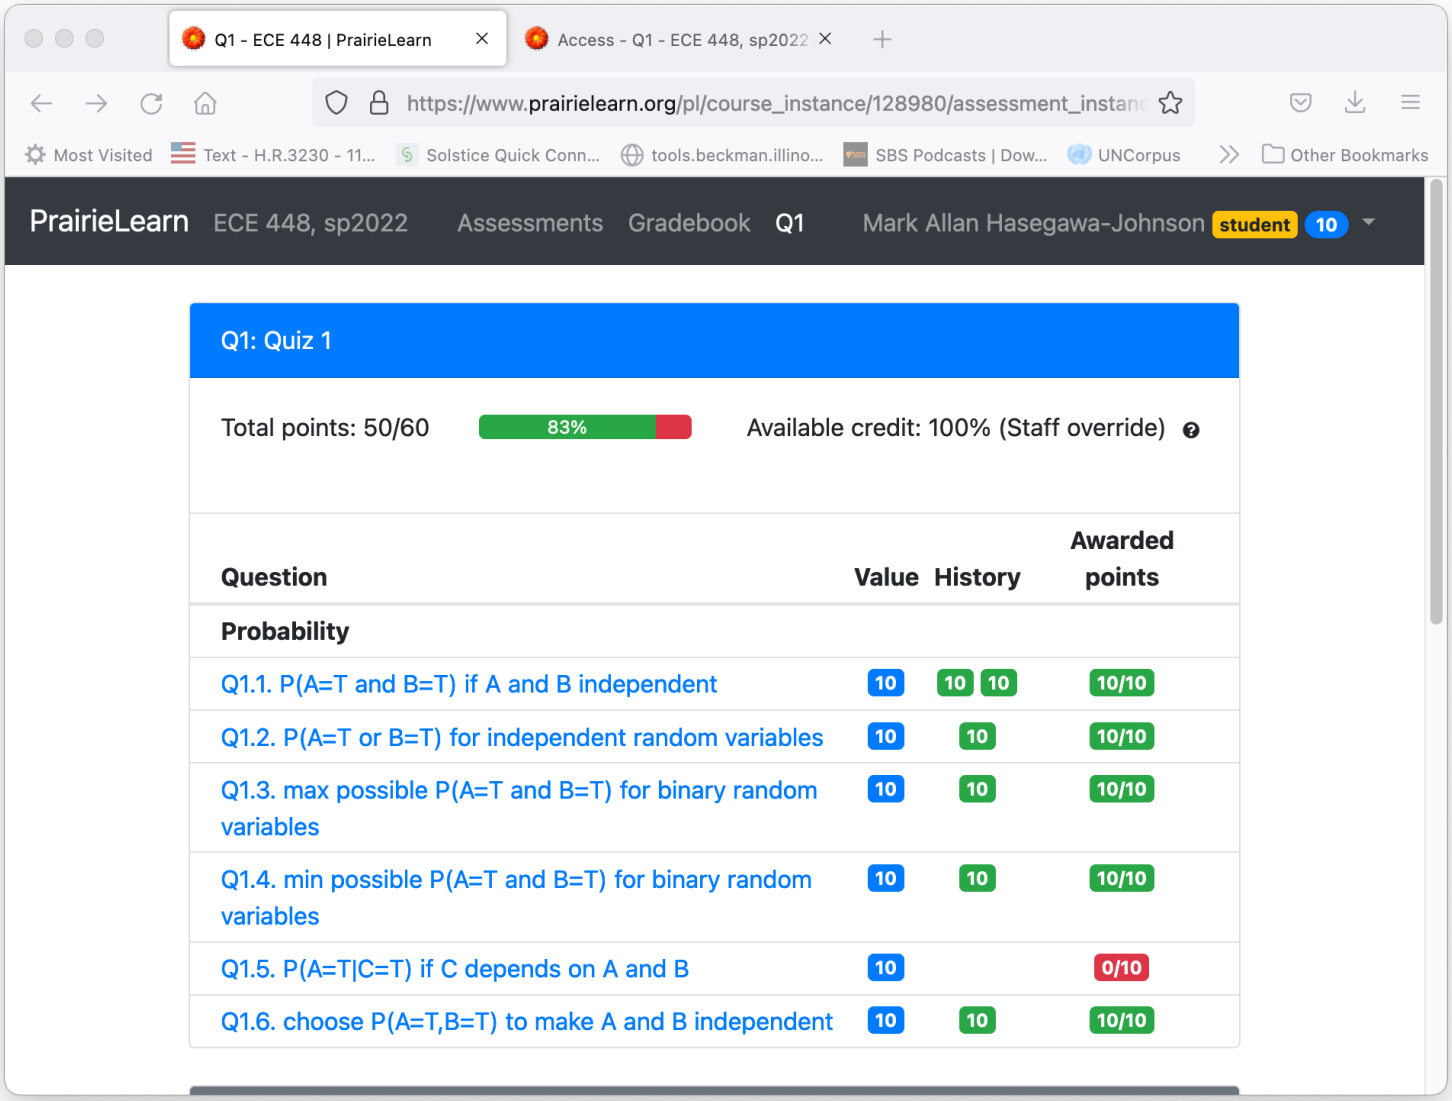

Machine Problems (MPs)

- Released ten days before due.
- Submit results to Gradescope, as often as you like.
- Late penalty up to 5% per day.
- You may submit any time during the semester for 50% credit.
- Late penalty is not waived for travel, trauma, interviews, religious holidays, or for any illness or injury that leaves you with the use of your hands.

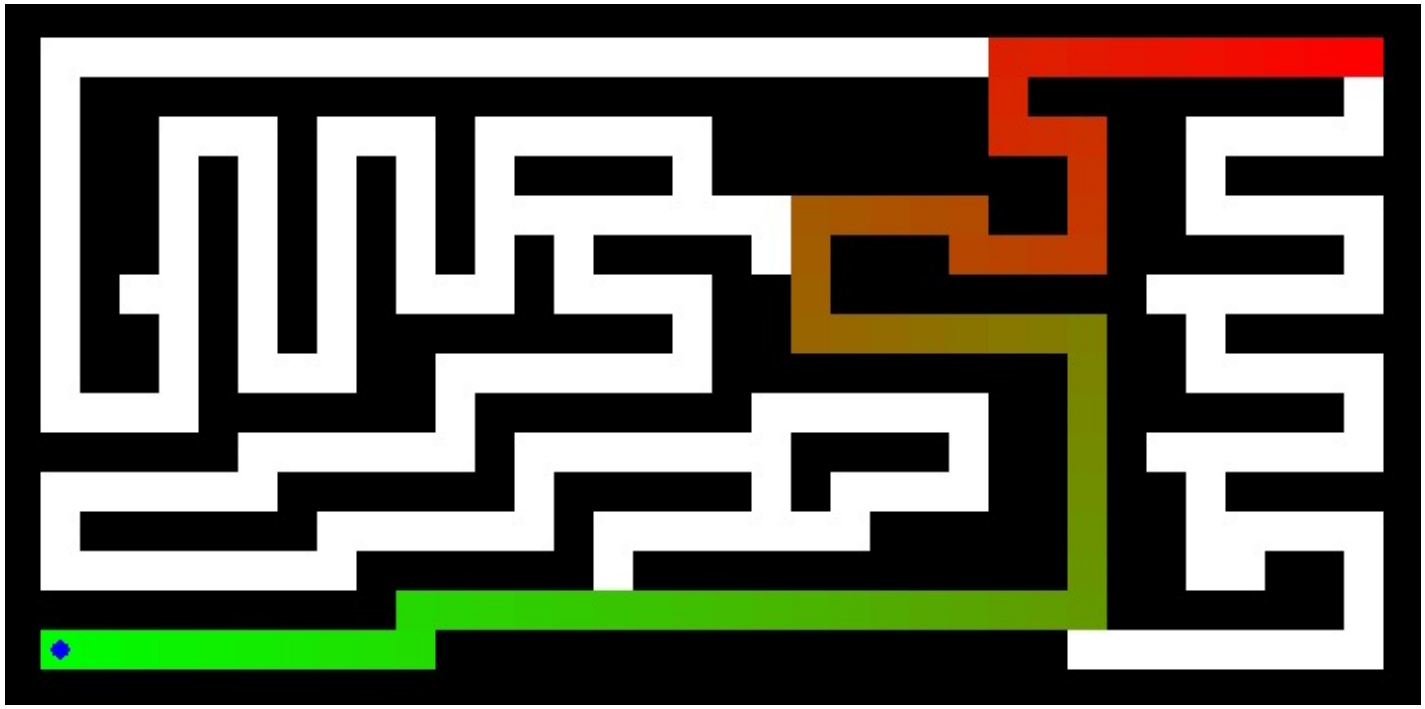

Exams

- Exams will be conducted synchronously, in person and via zoom.
- If you're in the in-person section and want to take the exam on zoom, tell us in advance.
- Those on zoom will need to turn your webcams on.

Zoom participants of the conference "Bubrikh readings: a documented vernacular", 2020.
Public domain image, AKA MBG, 2020.

Project

- 4-credit students do a final project.
- There are six stages in the final project.
- Each stage is graded (by me) and peer-reviewed (which does not affect the reviewee's grade).

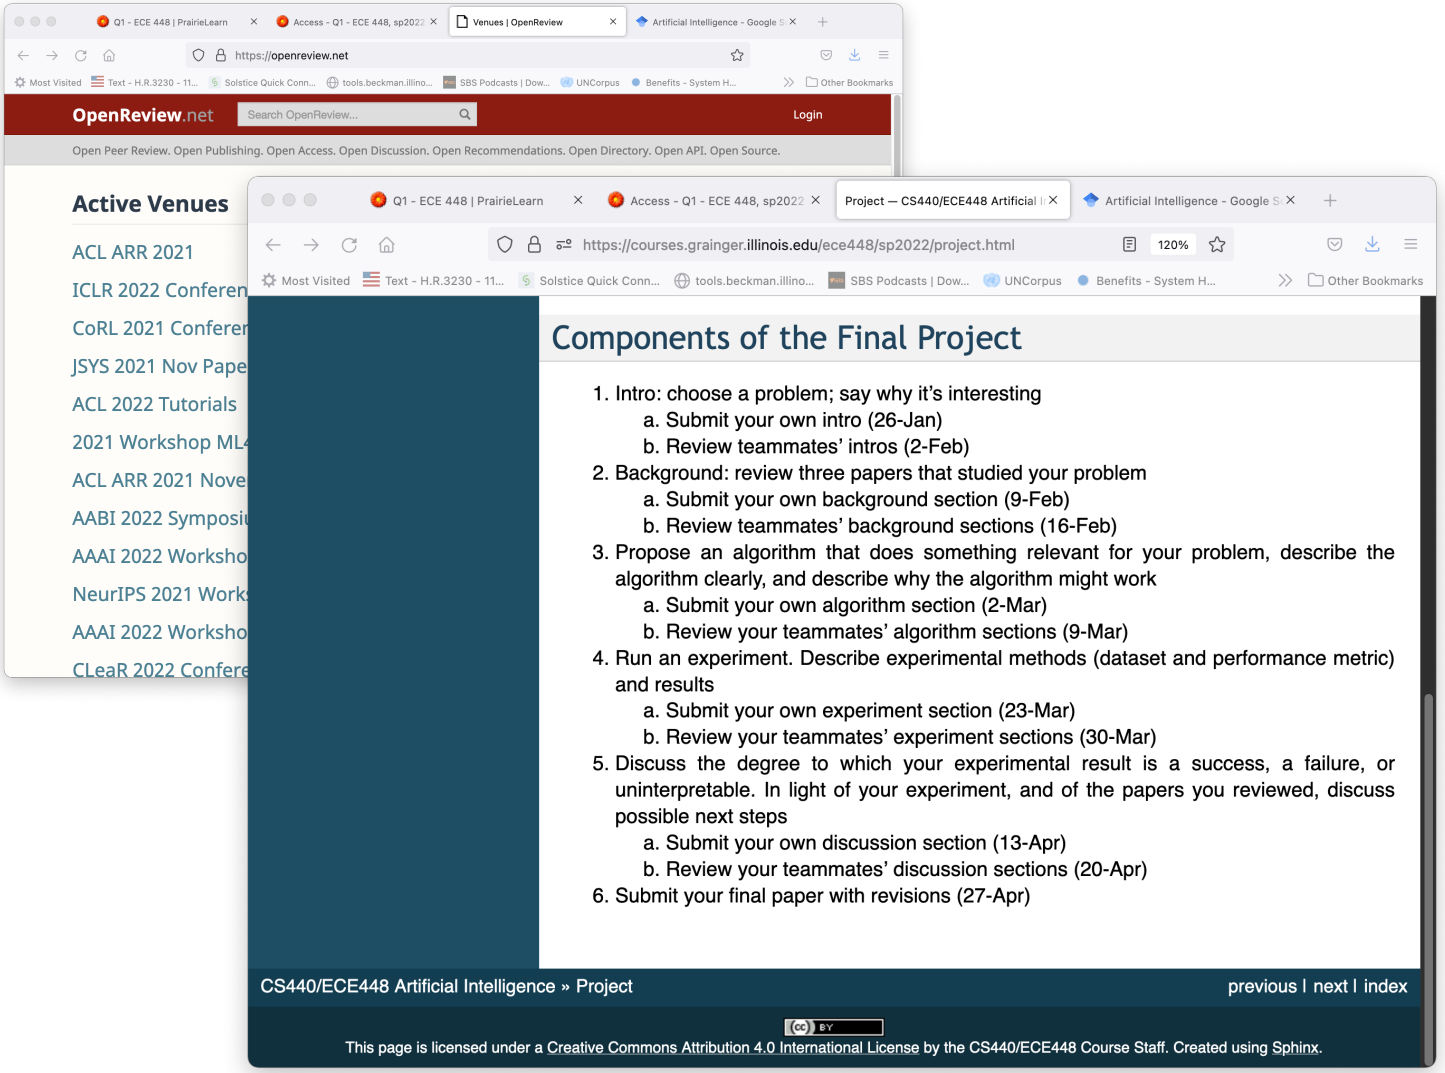

Outline

- Syllabus: CoViD, Text, Grades, Quizzes, MPs, Exams, and Project
- A brief history of Artificial Intelligence
- Artificially Intelligent Agents
  - Words to know: Performance, Environment, Actions, Sensors (PEAS)
  - Words to know: Rational, Human-like, Autonomous
- Environments in which an AI can operate
  - Words to know: Observable, Deterministic vs. Stochastic, Episodic vs. Sequential, Static vs. Dynamic, Discrete vs. Continuous, Single-Agent vs. Multi-Agent, Known vs. Unknown
- Outline of the rest of this course

A brief history of Artificial Intelligence

- 1954-1981: Symbolic reasoning
  - ”Artificial Intelligence” coined by John McCarthy in 1954
- 1982-2011: Probabilistic reasoning
  - Louis Liporace published continuous-observation hidden Markov models in 1982
- 2012-present: Deep learning
  - AlexNet won ImageNet in 2012

John McCarthy at a conference in 2006.

Outline

- Syllabus: CoViD, Text, Grades, Quizzes, MPs, Exams, and Project
- A brief history of Artificial Intelligence
- Artificially Intelligent Agents
  - Words to know: Performance, Environment, Actions, Sensors (PEAS)
  - Words to know: Rational, Human-like, Autonomous
- Environments in which an AI can operate
  - Words to know: Observable, Deterministic vs. Stochastic, Episodic vs. Sequential, Static vs. Dynamic, Discrete vs. Continuous, Single-Agent vs. Multi-Agent, Known vs. Unknown
- Outline of the rest of this course

Artificially Intelligent Agents

- An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators

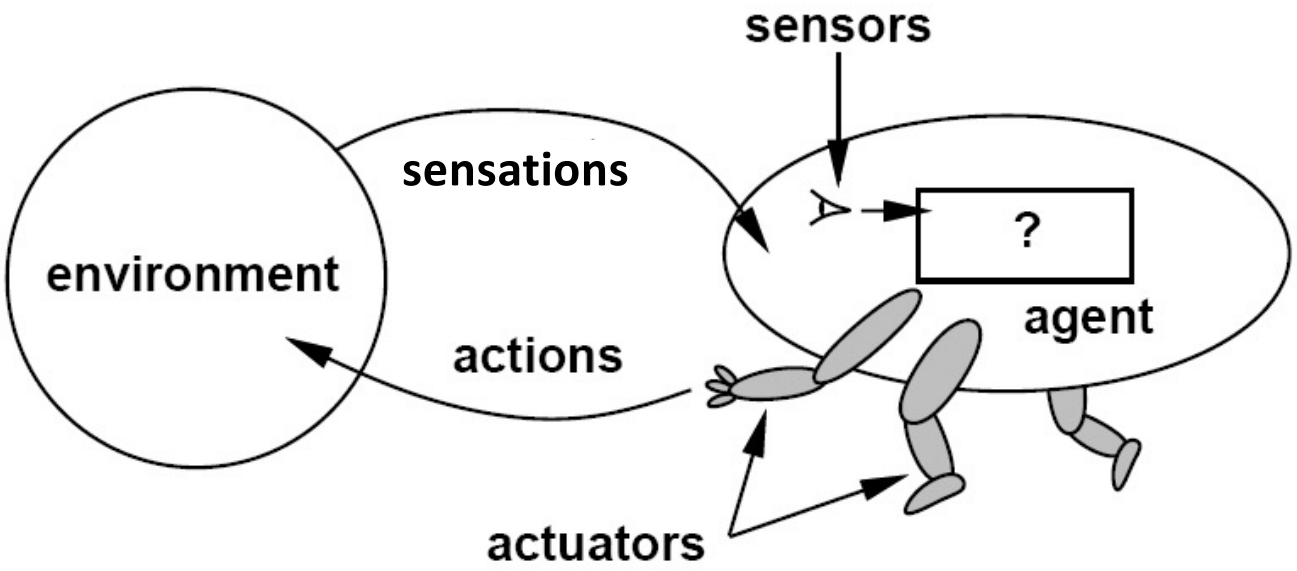

PEAS

- **P**erformance measure: Determined by the system designer, attempts to measure some intuitive description of behavior goodness.
- **E**nvironment: Where will the agent operate?
- **A**ctions: Determined by the system designer, usually trades off cost versus utility
- **S**ensors: Determined by the system designer, usually trades off cost versus utility

Example: Vacuum Cleaner Agent

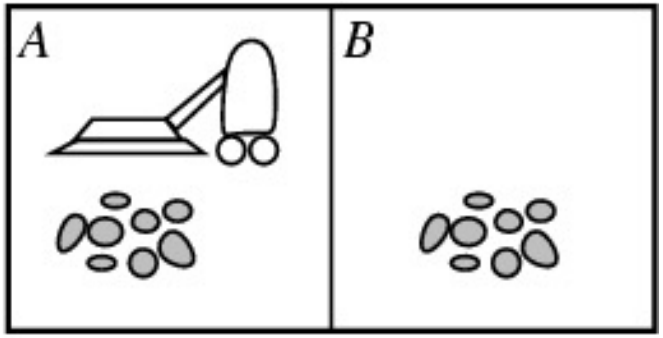

- Performance Measure:
  - Does it get the rooms clean?
  - How quickly?
  - Using how much electricity?
- Environment:
  - A two-room apartment
- Actions:
  - Suck; Move; DoNothing
- Sensors:
  - Is this room dirty or not?
  - (Maybe): which room am I in?

Criteria for judging an agent

- Intelligent: Is it intelligent?
- Rational: Does it always act in a manner that maximizes its expected performance metric?
- Human-like: Is it able to communicate with human beings, and explain or demonstrate its reasoning to them?
- Autonomous: Is it capable of revising its sensor→action mapping in response to changes in the environment?

What is Intelligence?

The word “intelligence” is surprisingly recent. Ancients used it to mean “the universal mind.” Early moderns (e.g., Bacon, Hobbes; 1500s) ridiculed it, and stopped using it.

What is Intelligence?

Charles Spearman popularized the modern definition in his paper "General intelligence objectively determined and measured," American Journal of Psychology 15(2):201-292.

- He showed that test scores are correlated across many subjects and proposed "general intelligence" as the faculty that unifies them.

Spearman's correlation matrix for six measures of school performance. All the correlations are positive, the *positive manifold* phenomenon. The bottom row shows the *g* loadings of each performance measure.[7]

	Classics	French	English	Math	Pitch	Music
Classics	–					
French	.83	–				
English	.78	.67	–			
Math	.70	.67	.64	–		
Pitch discrimination	.66	.65	.54	.45	–	
Music	.63	.57	.51	.51	.40	–
g	.958	.882	.803	.750	.673	.646

https://en.wikipedia.org/wiki/G_factor_(psychometrics)

Intelligence can be learned

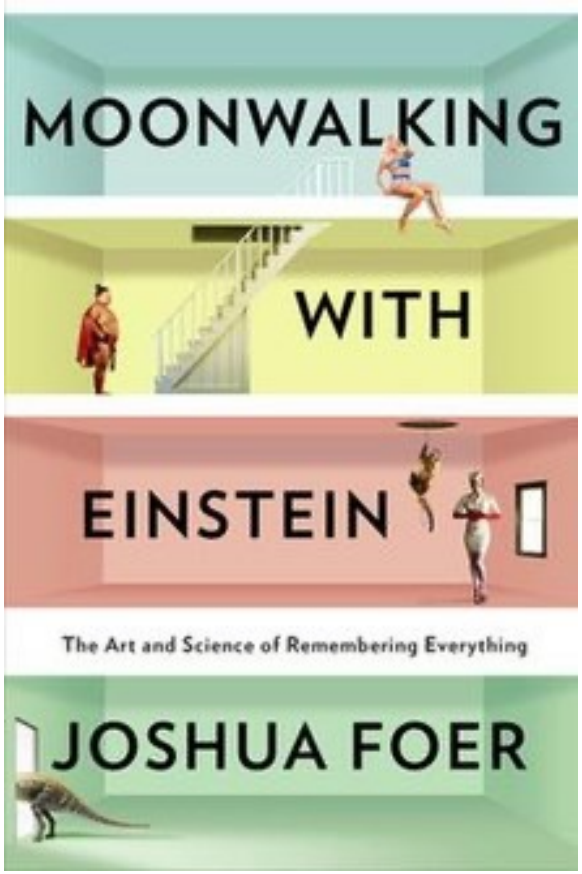

- It is possible to learn new strategies for learning.
- By doing so, it is possible to increase your score on tests of general intelligence.

What is “Artificial Intelligence”

We usually mean more than just the ability to learn (ability to learn is covered in “Machine Learning”). We usually mean:

- Ability to learn (a.k.a. intelligence)
- Ability to choose a plan that optimizes my rewards (a.k.a. rationality)
- Ability to communicate my plan to human beings, and behave in a human-like manner around human beings (a.k.a. human-like)
- Ability to modify my own behavior based on things I have learned (a.k.a. autonomous)
- Ability to appropriately trade off my own performance measures against the goals of other agents (a.k.a. morality)

Rationality

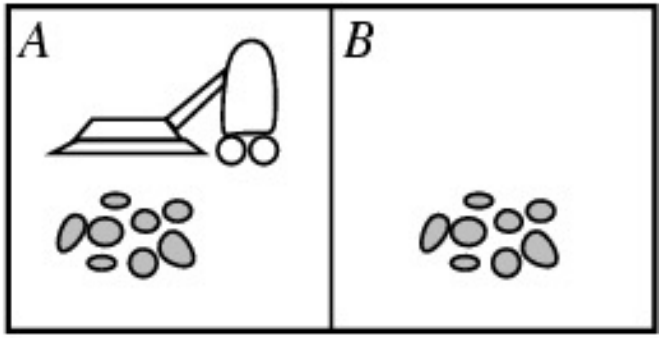

- Performance Measure:
  - Does it get the rooms clean?
  - How quickly?
  - Using how much electricity?
- Then here is a rational agent program:
  1. If the current room is dirty, then Suck.
  2. If I have not been to the other room, Move then Suck.
  3. DoNothing.

Human-likeness

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

…but…

Chatbots have been fooling human interviewers since 1966.

The chatbot Eugene Goostman fooled 30% of judges in 2014, by pretending to be a 13-year-old Ukrainian boy.

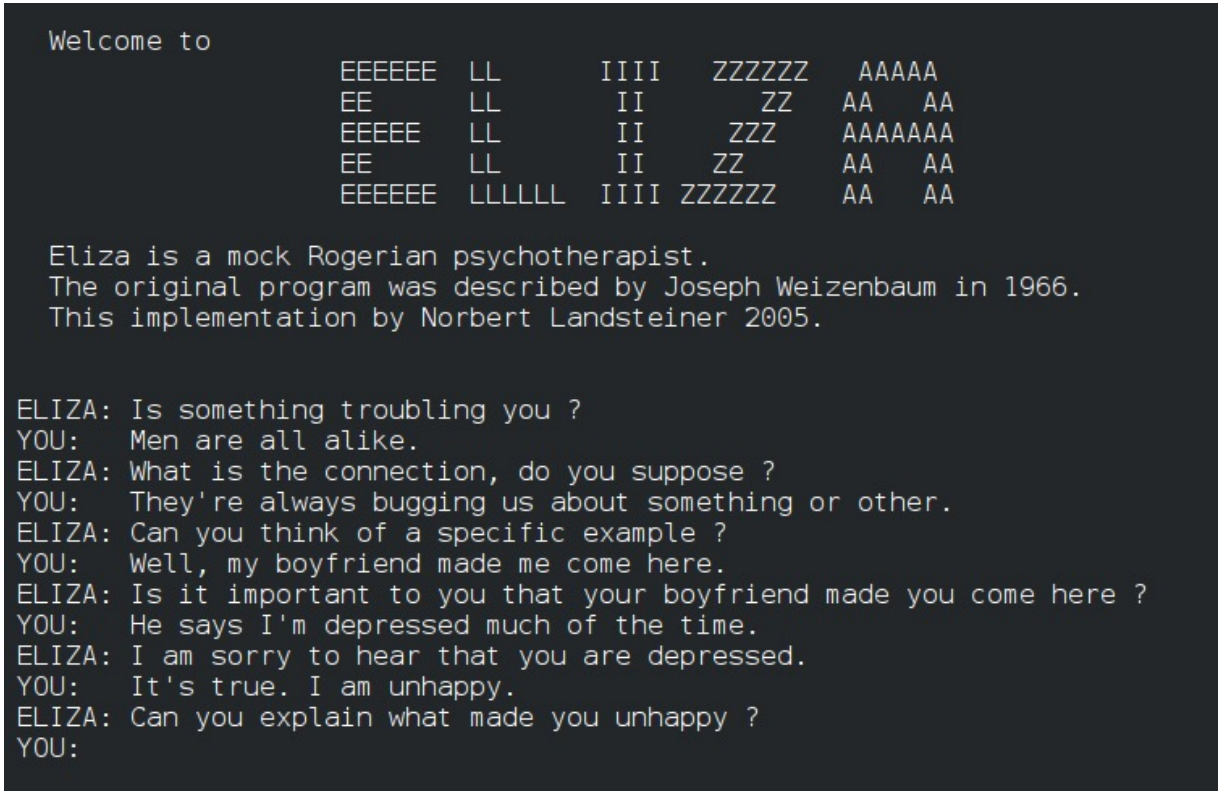

Autonomy

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.

Morality

Many recent authors have argued that an intelligent, rational, human-like, autonomous agent is undesirable if it cannot appropriately weight the needs & goals of other agents, including human beings.

Ethics & morality of AI will be interspersed throughout this course, to give you some idea of current thinking.

Reproduction of the Prague Golem. Public domain image, Thander, 2007

Outline

- Syllabus: CoViD, Text, Grades, Quizzes, MPs, Exams, and Project
- A brief history of Artificial Intelligence
- Artificially Intelligent Agents
  - Words to know: Rational, Human-like, Autonomous
- Environments in which an AI can operate
  - Words to know: Observable, Deterministic vs. Stochastic, Episodic vs. Sequential, Static vs. Dynamic, Discrete vs. Continuous, Single-Agent vs. Multi-Agent, Known vs. Unknown
- Outline of the rest of this course

Properties of Environments

- 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?

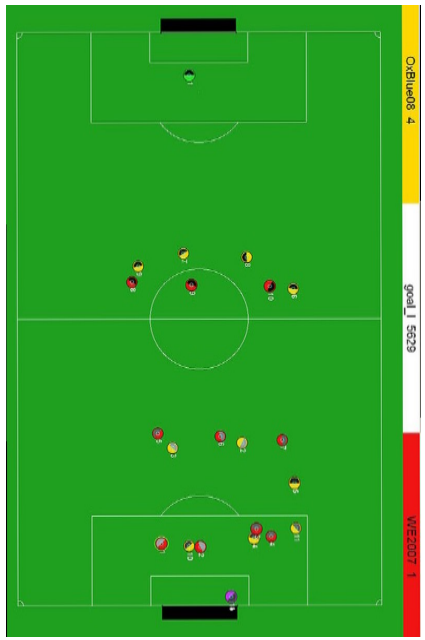

vs.

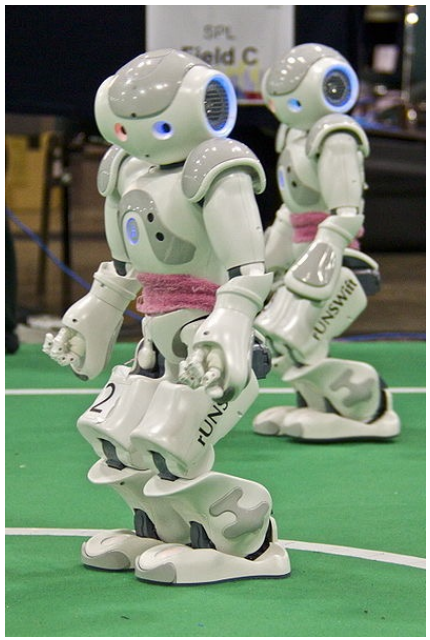

Source: L. Zettlemoyer

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.

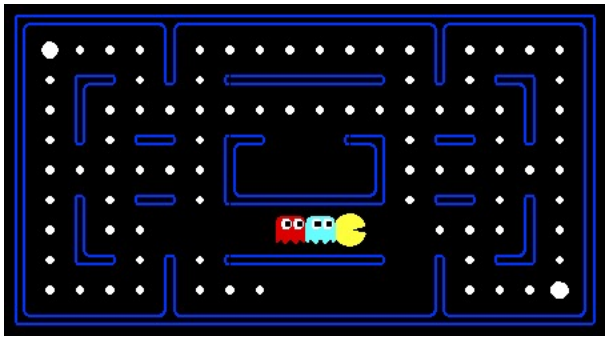

Static vs. dynamic

- 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.

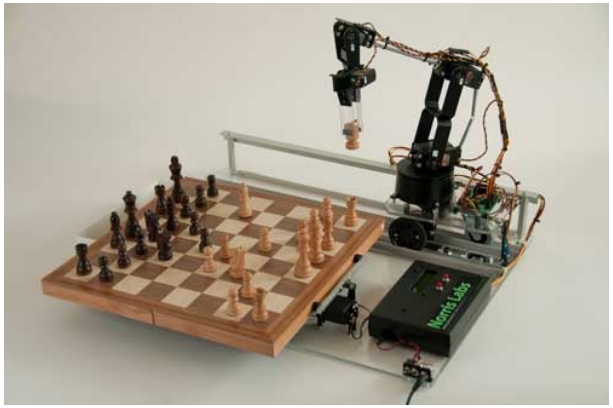

Single-agent vs. multiagent

- Is an agent operating by itself in the environment?

vs.

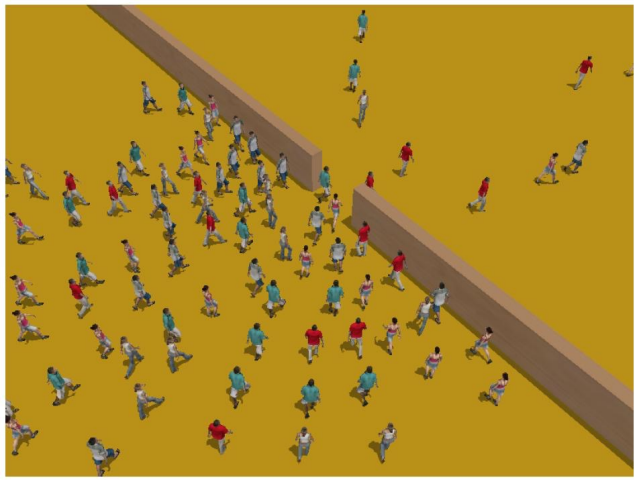

Known vs. unknown

- 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.

Outline

- Syllabus: CoViD, Text, Grades, Quizzes, MPs, Exams, and Project
- A brief history of Artificial Intelligence
- Artificially Intelligent Agents
  - Words to know: Rational, Human-like, Autonomous
- Environments in which an AI can operate
  - Words to know: Observable, Deterministic vs. Stochastic, Episodic vs. Sequential, Static vs. Dynamic, Discrete vs. Continuous, Single-Agent vs. Multi-Agent, Known vs. Unknown
- Outline of the rest of this course

Outline of the Course

1. Episodic environments: Deterministic vs. Stochastic
   - MP1: Naïve Bayes spam filter
2. Episodic environments: Discrete vs. Continuous
   - MP2: Neural Network image classifier
3. Sequential environments
   - MP3: A*-Search maze solver
4. Sequential environments: Fully vs. Partially Observable
   - MP4: Hidden Markov Model text parser
5. Sequential environments: Single-agent vs. Multi-agent
   - MP5: Alpha-beta chess player
6. Sequential environments: Known vs. Unknown
   - MP6: Reinforcement learning video game player