CS440/ECE448: Artificial Intelligence
Lecture 1: Course Intro


Mary Shelley

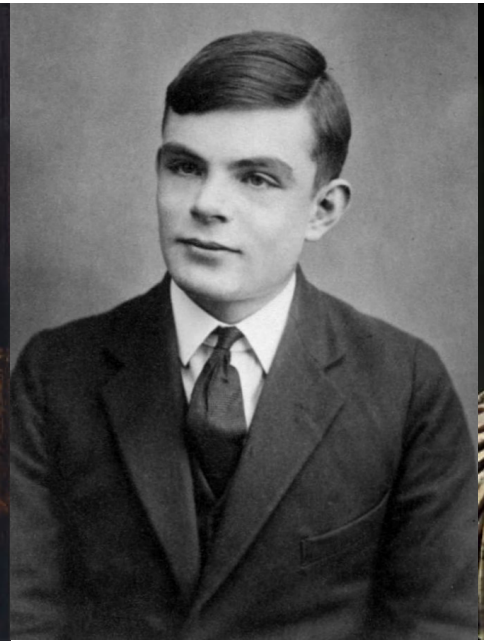
Alfred Turing

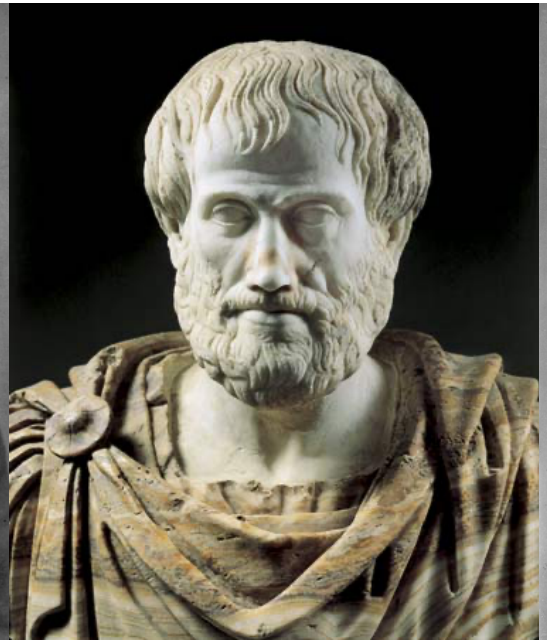
Aristotle

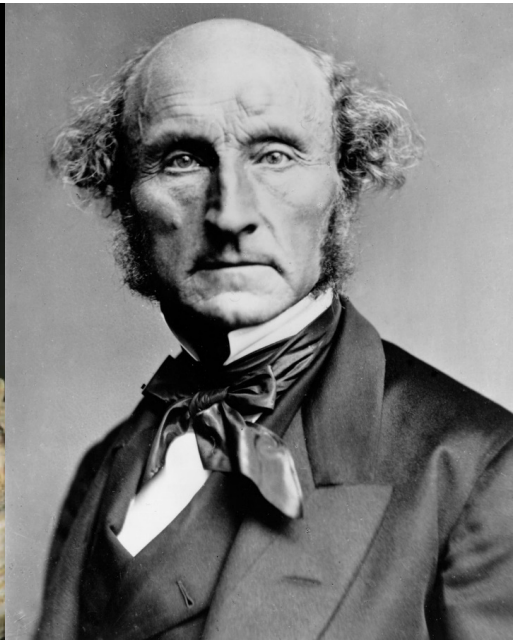
John Stuart Mill

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- A brief history (chapter 1)
- Environments in which an AI can operate (chapter 2.3)
- Outline of the rest of this course

Course Intro: Syllabus

- Web page: https://courses.engr.illinois.edu/ece448/sp2021/
- Grading
- Exams & Collabs
- Homework
- Graduate Research Seminar
- Apps
- Textbook

Grading: 3-credit

- Exams: 36%
- Homework (MPs): 50%
- Collabs: 14%

Grading: 4-credit

- Exams: 30%
- Homework (MPs): 40%
- Collabs: 10%
- Research seminar: 20%

Exams

- Midterms 3/5 and 4/9, Final during finals week
- Synchronous: Take them during the class period
- Open-book, open-notes, open-internet
- Timed: practice doing problems quickly.
- Collab worksheets will be composed of past exam problems.

Collabs

- Scheduled in week 1; mandatory from week 2. Most will be M, W, or F, 1-2pm.
- Participation: 7% of your grade (be there, help your teammates. Collaborate to solve all the problems on the worksheet.)
- Worksheet: 7% of your grade (turn it in to gradescope, every problem answered; answers need not be correct. Correct answers will be distributed after the submission deadline.)

Homework

- 6 machine problems autograded on gradescope.com.
- Language: python. Learn it: https://docs.python.org/3/tutorial/
- MP1 (search): released Friday 1/29, due Wednesday 2/10 by 11:59pm.
- Late policy:
  - Late submissions are accepted any time until the end of finals week.
  - If your submission is late, your autograder score is multiplied by

  $0.5 + 0.5e^{-t/168}$, t =autograder lateness score, in hours
  - Lateness without penalty is possible if you are in the hospital or in legal custody for most of the last day.

Graduate Research Seminar

- Meet weekly, Mondays 1-2pm, starting week 2
- Basic format:
  - Each student gives a 3-slide presentation, each week
  - Presentation templates distributed in LaTeX, to help you know what to present. Presentation in LaTeX not required but strongly recommended.
- Topics:
  - Week 1: learning LaTeX
  - Weeks 2-5: Research paper #1
  - Weeks 6-9: Research paper #2
  - Weeks 10-13: Research paper #3 (could be one of yours, if you wish)
- Grading:
  - Participation (present, ask questions): 9%
  - Be session chair, typically 3X in the semester: 1%
  - Submit your slides after presentation, to gradescope: 10%

Apps

- All homework and exams graded and submitted at: https://www.gradescope.com/
- Exams: https://compass2g.illinois.edu.
- Q/A forum: https://piazza.com
- Videos of lectures: https://mediaspace.Illinois.edu

Textbook

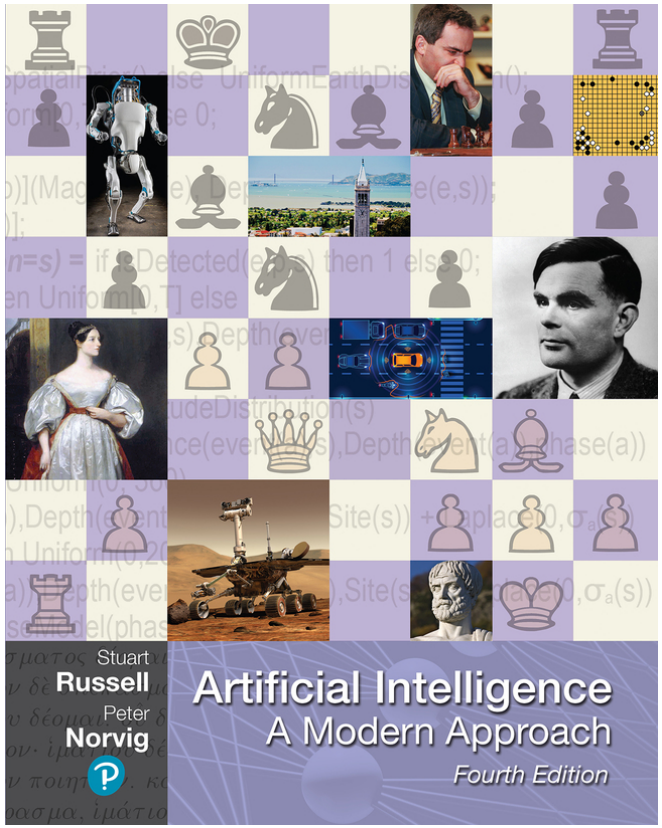

Artificial Intelligence, A Modern Approach: Fourth Edition by Russell & Norvig

- Pretty good listing of the topics covered in this course
- In-depth treatment of knowledge-based/expert-system AI; introduces probabilistic and learning-based methods
- Sample problems and readings will be specified when applicable

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- Artificial Intelligence: a brief history (chapter 2)
- Environments in which an AI can operate (chapter 3)
- Outline of the rest of this course

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

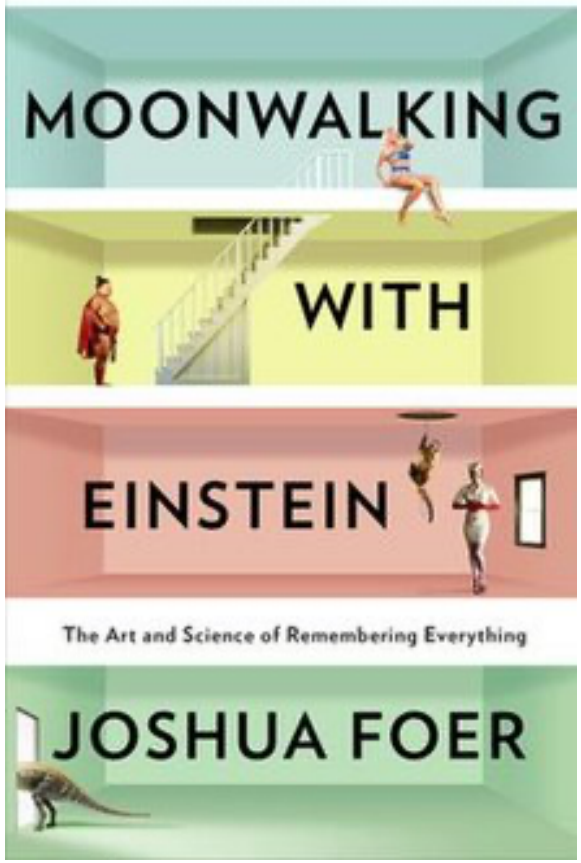

- It is possible to learn new strategies for learning.
- By doing so, it is possible to increase your score on tests of general intelligence.

Intelligence is just the beginning

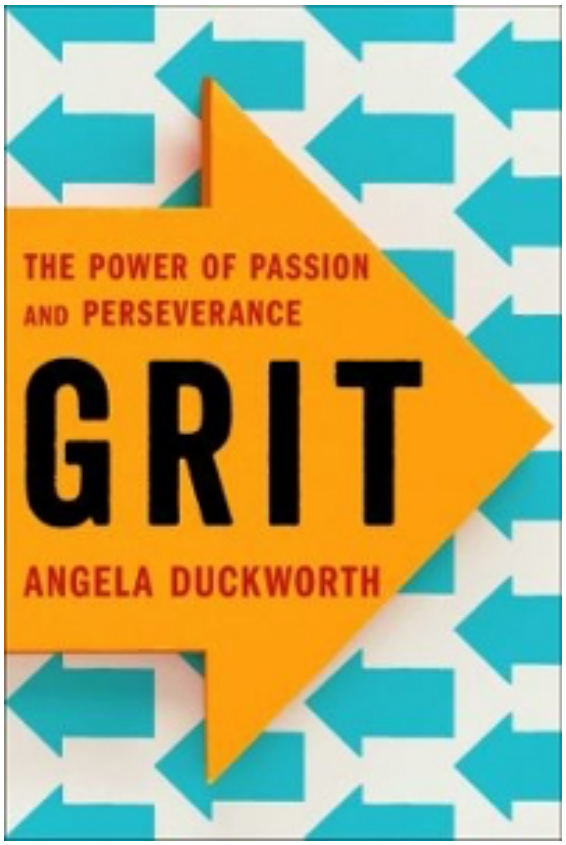

Grit ("passion and perseverance for long-term goals") is uncorrelated with general intelligence and is a better predictor of long-term success.

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- A brief history (chapter 2)
- Environments in which an AI can operate (chapter 3)
- Outline of the rest of this course

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 plan (a.k.a. rationality)
- Ability to communicate in natural language (a.k.a. abstraction)
- Ability to reason about morality (a.k.a. humanity)

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- A brief history (chapter 2)
- Environments in which an AI can operate (chapter 3)
- Outline of the rest of this course

A brief history

- 1943: Artificial Intelligence imitates neuroscience (ability to learn)
- 1956: Artificial Intelligence is about symbolic logic (ability to plan)
- 1984: Artificial Intelligence is about probability (ability to communicate in natural language)
- 1988: Artificial Intelligence is about the understanding of causality (ability to reason about moral judgments)

1943: McCulloch-Pitts Neuron Model

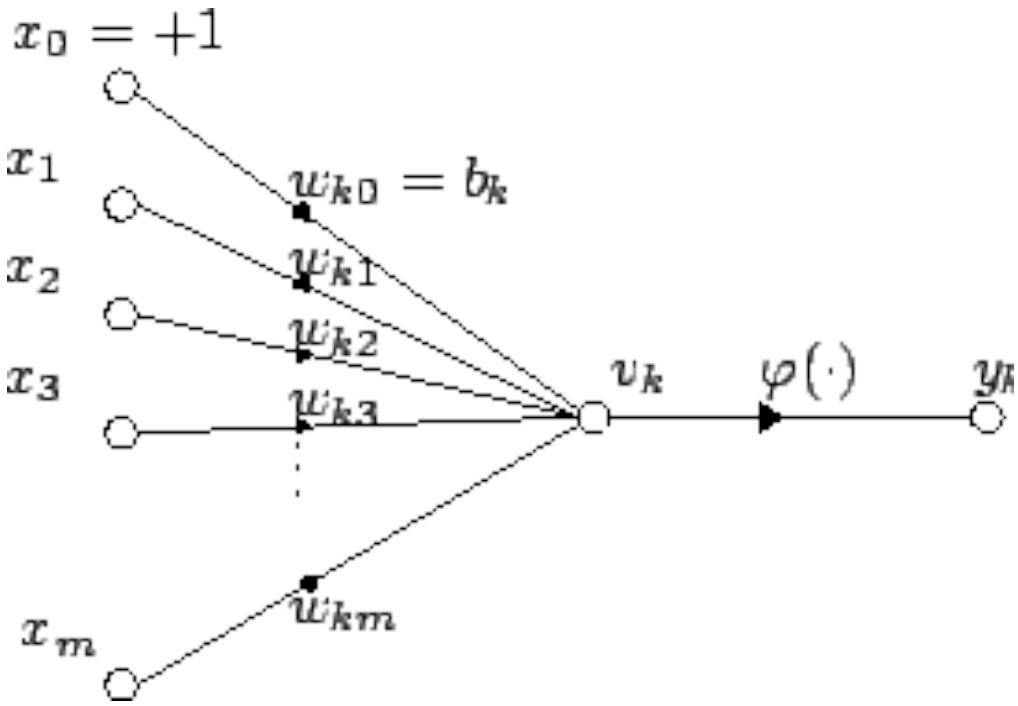

Attribution: Plarroy, https://commons.wikimedia.org/wiki/File:Artificial_neuron.png

1956: “Artificial Intelligence”

The term was invented in (John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon, “A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence,” August 1955):

“We propose that a 2-month, 10-man study of artificial intelligence be carried out during the summer … An attempt will be made to find how to make machines

1. use language,
2. form abstractions and concepts,
3. solve kinds of problems now reserved for humans
4. improve themselves.”

1956: Logic Theorist

- **<u>Reasoning as search</u>**:
  - Root of the search tree: initial hypothesis
  - Branch: a deduction based on the rules of logic
  - Goal state: the theorem to be proven
- Proved 38 of the first 52 theorems in chapter 2 of the *Principia Mathematica*.

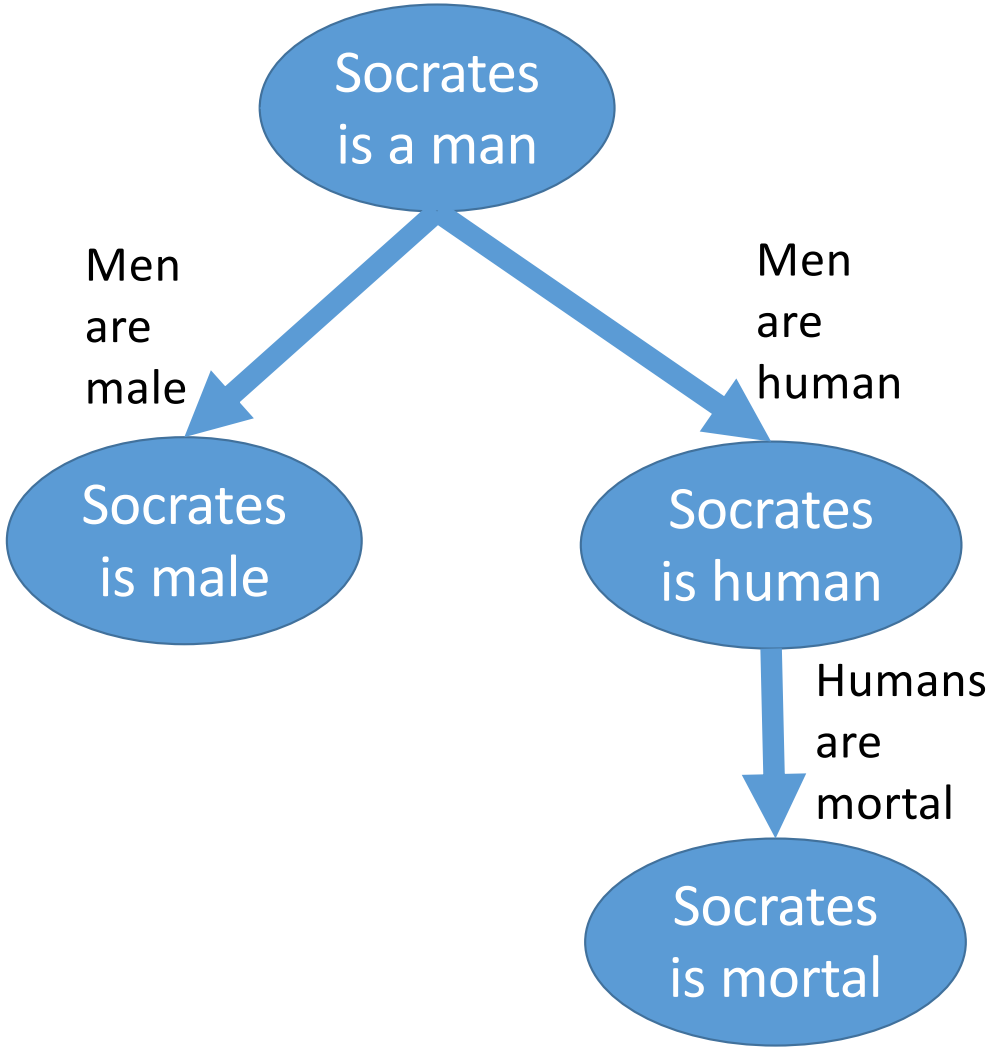

The ALPAC Report of 1966

"They concluded, in a famous 1966 report, that machine translation was more expensive, less accurate and slower than human translation."

Photo: Eldon Lyttle, https://commons.wikimedia.org/wiki/File:Computer-translation_Briefing_for_Gerald_Ford.jpg


AI Winter

1984: Dragon Dictate 1.0 speech recognizer

Explicitly minimized the probability of error using a new technology called a ”hidden Markov model.”

1988: Probabilistic reasoning becomes causal reasoning

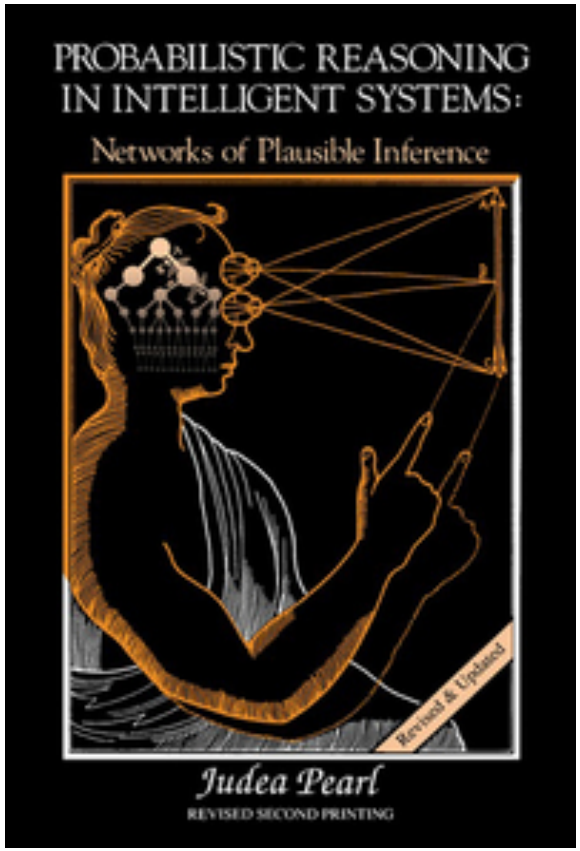

- Judea Pearl's "Bayesian Networks" show computer scientists how to extend the main ideas of the HMM to every other problem in AI.
- Introduced a mathematics for reasoning about causality, and therefore, about moral judgments.

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- A brief history (chapter 2)
- Environments in which an AI can operate (chapter 3)
- Outline of the rest of this course

Environments in which an AI may operate

- Ability to sense your surroundings:
  - Fully observable vs. Partially observable
- Computational limitations on your ability to plan:
  - Discrete vs. Continuous
  - Static vs. Dynamic
  - Deterministic vs. Stochastic
- Limitations on your ability to control the environment:
  - Episodic vs. Sequential
  - 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?

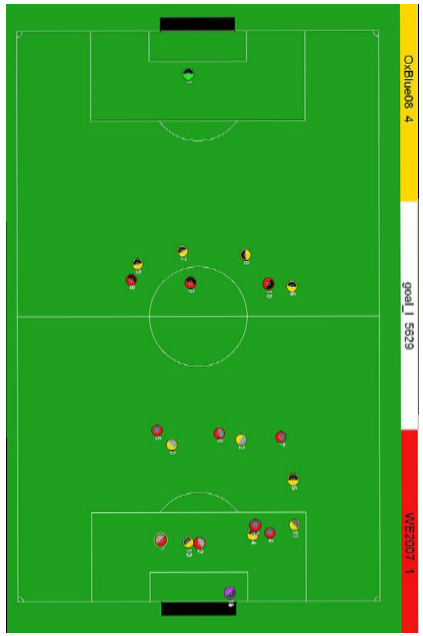

vs.

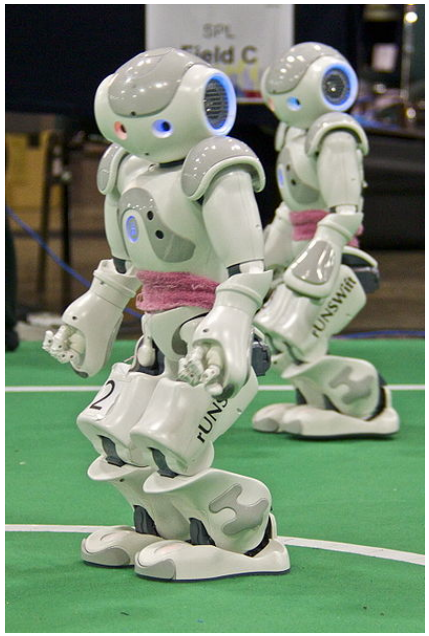

Source: L. Zettlemoyer

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.

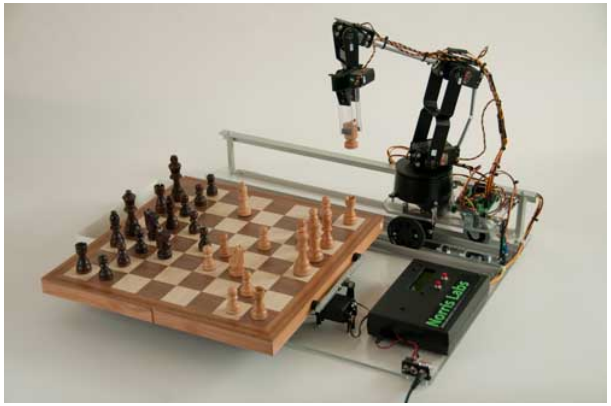

Static vs. dynamic

- Is the world changing while the agent is thinking?

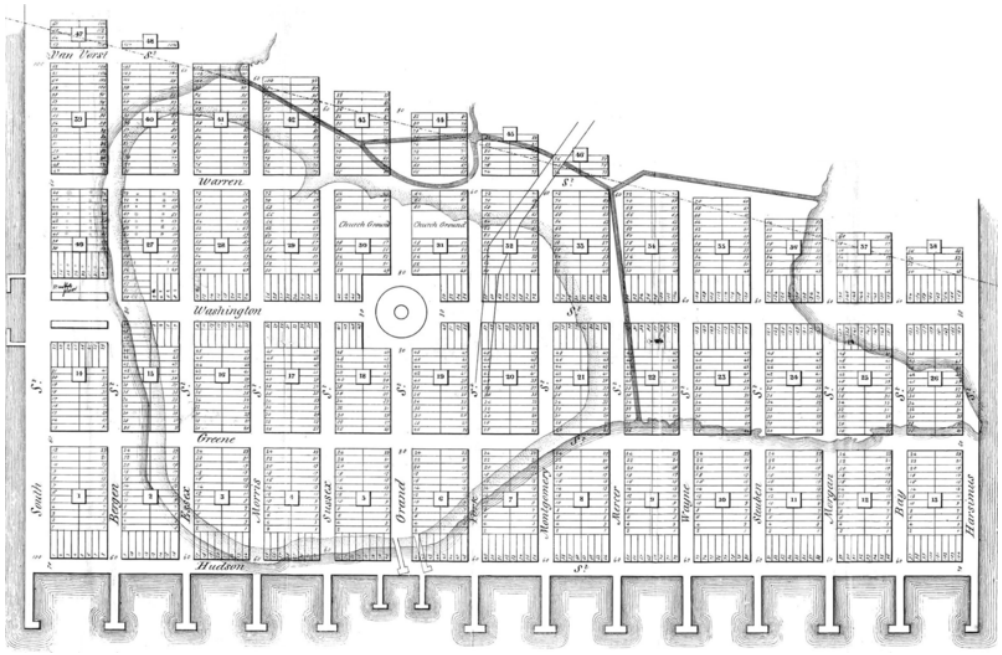

vs.

GFDL Image by Minesweeper, 2005

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.

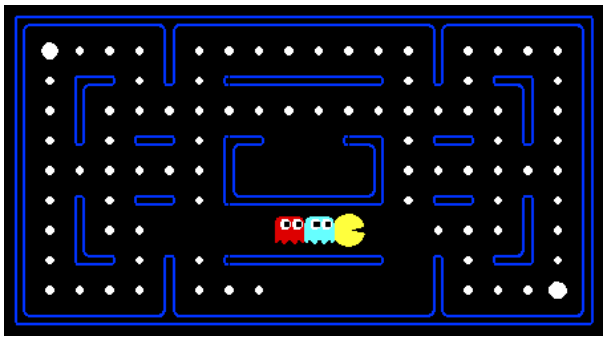

Single-agent vs. multiagent

- Is an agent operating by itself in the environment?

vs.

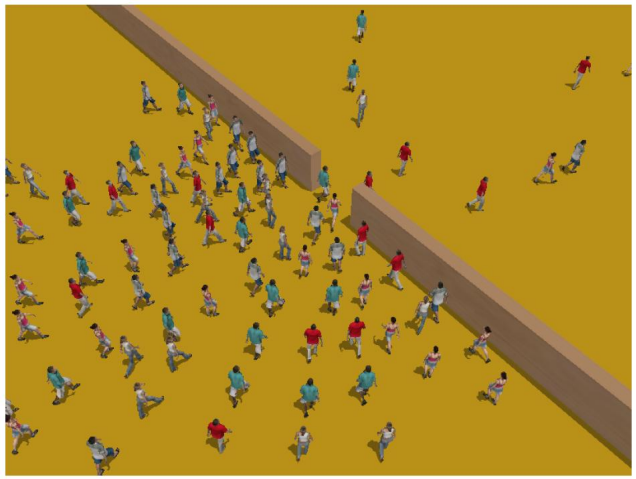

Known vs. unknown

- 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

Monopoly, Parker Brothers 1935; photo CC-BY-NC Fir0002/Flagstaffotos

vs.

Myst, © Cyan Worlds, 1993

Outline

- Syllabus: Grades, Exams, Collabs, Homework, Seminar, and Apps
- Intelligence
- Artificial Intelligence (chapter 1)
- A brief history (chapter 2)
- Environments in which an AI can operate (chapter 3)
- Map of the rest of the course

Map of the rest of the course