15-780: Introduction and History of AI

J. Zico Kolter and Tuomas Sandholm

January 11, 2016

What is AI?

(Some) history of AI

(Some) current applications of AI

Overview of course

What is “AI”?

Some classic definitions

Buildings computers that ...

Think like humans	**Think rationally**
- cognitive science / neuroscience - e.g., General Problem Solver (Newell and Simon, 1961)	- logic and automated reasoning - but, not all problems can be solved just by reasoning
Act like humans	**Act rationally**
- Turing Test - ELIZA, Julia, Loebner prize	- basis for intelligent agent framework - unclear if this captures the current scope of AI research

The pragmatist's view

AI is that which appears in academic conferences on AI

The pragmatist's view

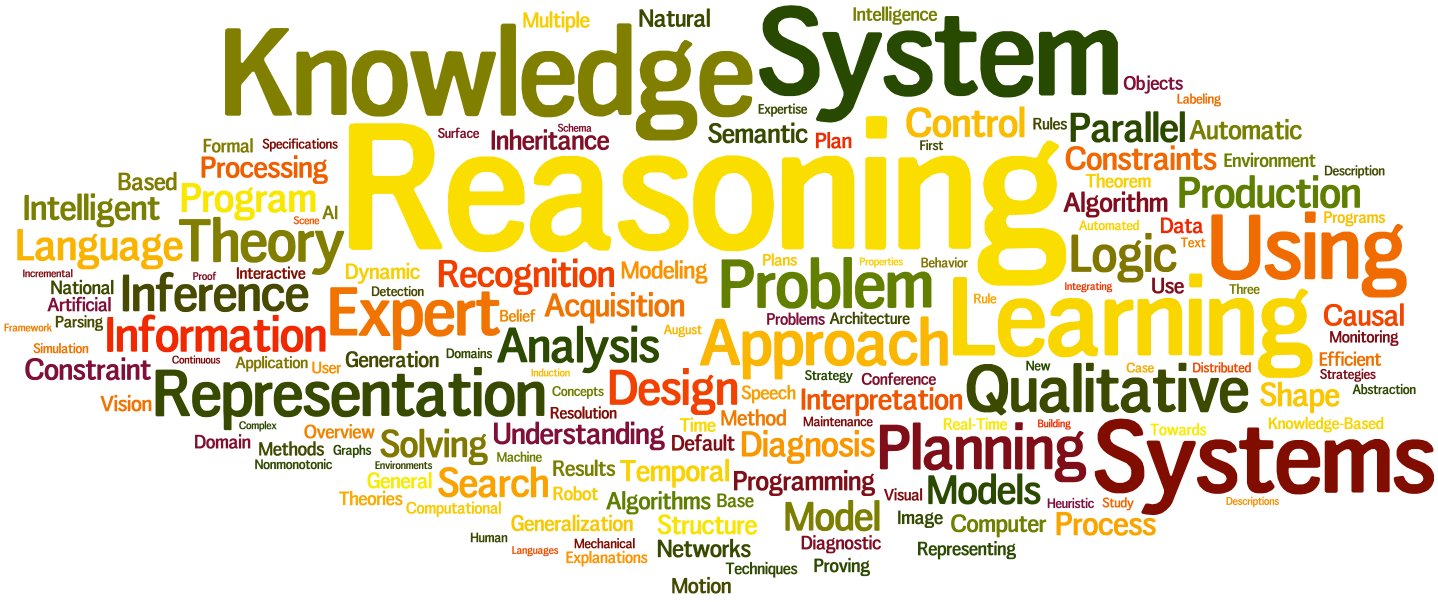

1980s

The pragmatist's view

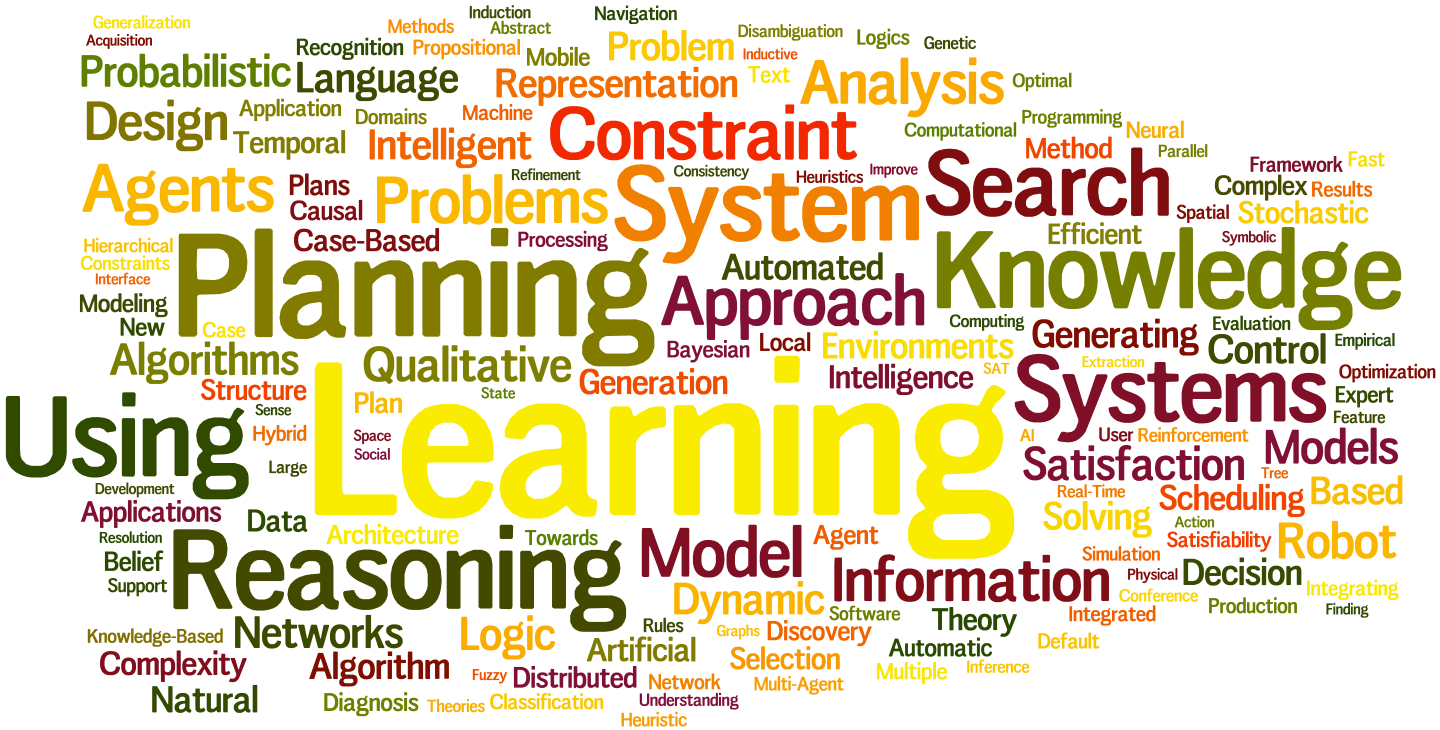

1990s

The pragmatist's view

2000s

The pragmatist's view

2010s

A broader definition

We won't worry too much about definitions of AI, but the following will suffice:

Artificial Intelligence is the development and study of computing systems that address a problem typically associated with some form of intelligence

(Some) history of AI

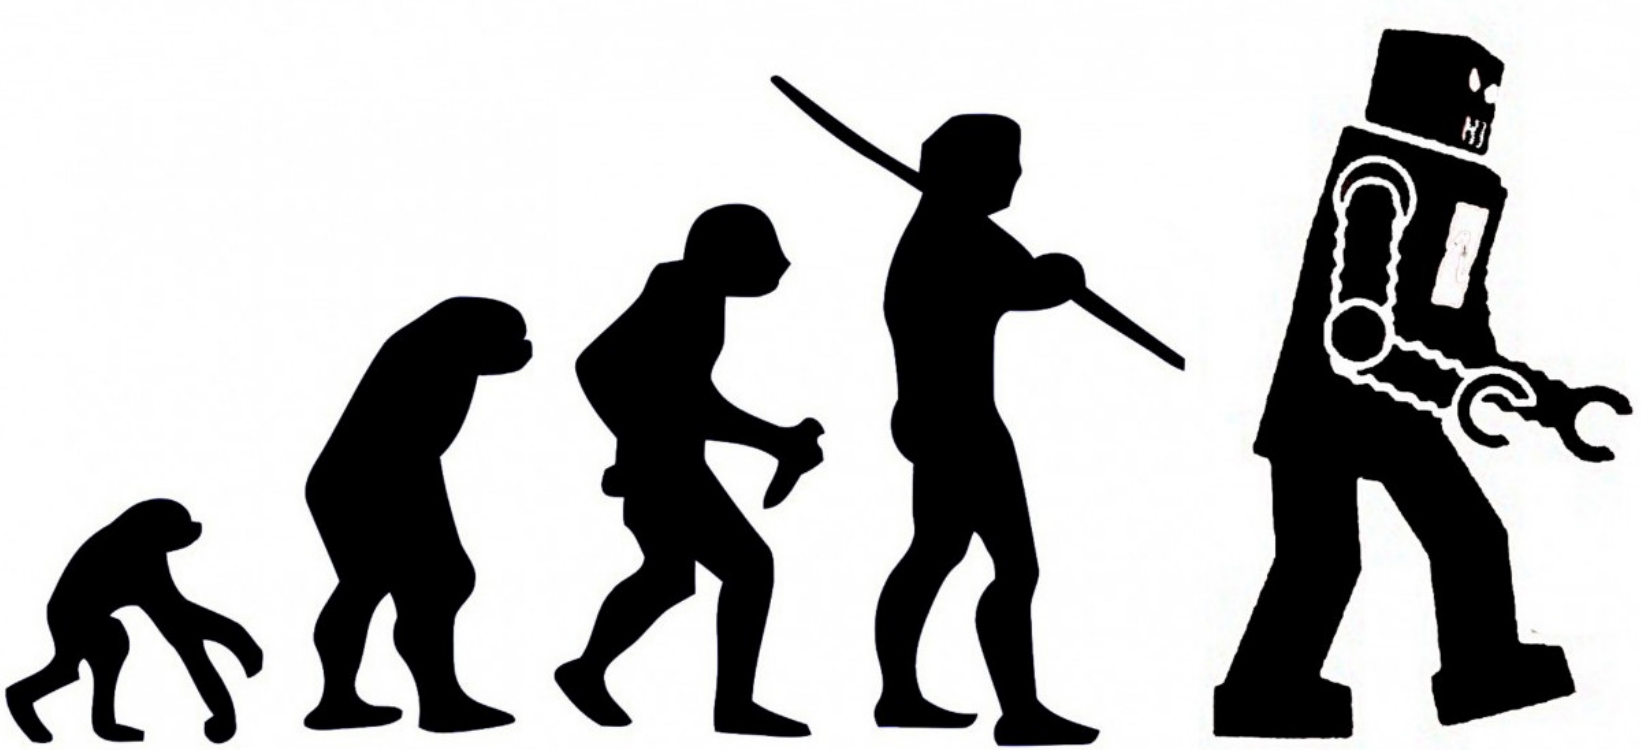

Reading: Russell and Norvig, Chapter 1

Pre-history (400 B.C. –)

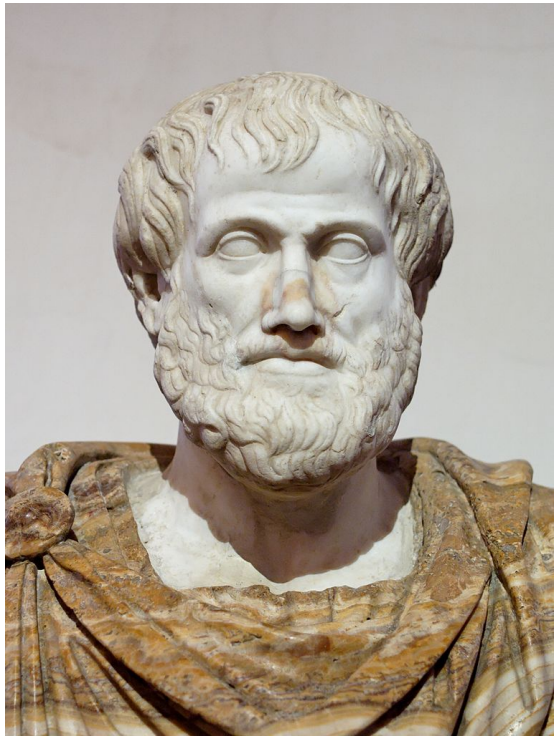

Philosophy: mind/body dualism, materialism

Mathematics: logic, probability, decision theory, game theory

Cognitive psychology

Computer engineering

Birth of AI (1943–1956)

1943 – McCulloch and Pitts: simple neural network

1950 – Turing test

1955-56 – Newell and Simon: Logic Theorist

1956 – Dartmouth conference, organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, Claude Shannon

Early successes (1950s–1960s)

1952 – Arthur Samuel develops checkers program, learns via self-play

1958 – McCarthy LISP, advice taker, time sharing

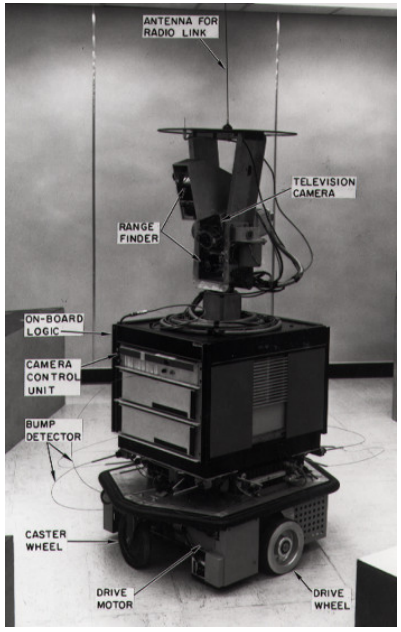

1958 – Rosenblatt's Perceptron algorithm learns to recognize letters

1968-72 – Shakey the robot

1971-74 – Blocksworld planning and reasoning domain

First “AI Winter” (Later 1970s)

Many early promises of AI fall short

1969 – Minsky and Pappert’s “Perceptrons” book shows that single-layer neural network cannot represent XOR function

1973 – Lighthill report effectively ends AI funding in U.K.

1970s – DARPA cuts funding for several AI projects

Expert systems & business (1970s–1980s)

Move towards encoding domain expert knowledge as logical rules

1971-74 – Feigenbaum's DENDRAL (molecular structure prediction) and MYCIN (medical diagnoses)

1981 – Japan's "fifth generation" computer project, intelligent computers running Prolog

1982 – R1, expert system for configuring computer orders, deployed at DEC

Second “AI Winter” (Late 1980s–Early 1990s)

As with past AI methods, expert systems seem to fail to deliver on promises

Complexity of expert systems made them difficult to develop/maintain

1987 – DARPA again cuts AI funding for AI expert systems

1991 – Japan’s 5th generation project did not meet goals

Splittering of AI (1980s–2000s)

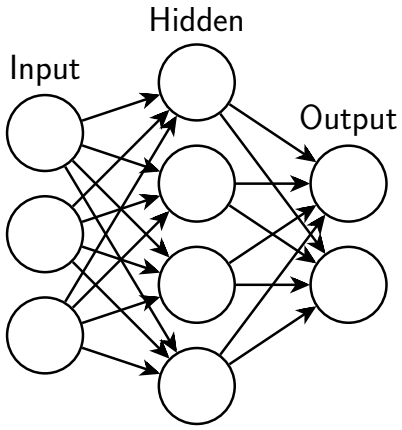

Much of AI focus shifts to subfields: machine learning, multiagent systems, computer vision, natural language processing, robotics, etc; but several hugely important developments in these areas:

1982 – Backpropagation for training neural networks popularized by Rumelhart, Hopfield (amongst other)

1988 – Judea Pearl's work on Bayesian networks

1995 – NavLab5 automobile drives across country, steering itself 98% of the time

Focus on applications (1990s-2010s)

Meanwhile, AI (sometimes under the guise of a subfield), achieved some notable milestones

1997 – Deep Blue beats Gary Kasparov

2001-2010 – $60 billion involved in combinatorial sourcing auctions

2005,2007 – Stanford and CMU respectively win DARPA grand challenge in autonomous driving

2011 – IBM's Watson defeats human opponents on Jeopordy

Reemergence of “AI” (2010s–??)

“AI” seems to be a buzzword again

Google, Facebook, Twitter, etc, all have large AI labs, labeled as such

2012 – Deep neural network wins image classification contest

2013 – DeepMind shows computer learning to play Atari games

Current applications of AI

AI is all around us

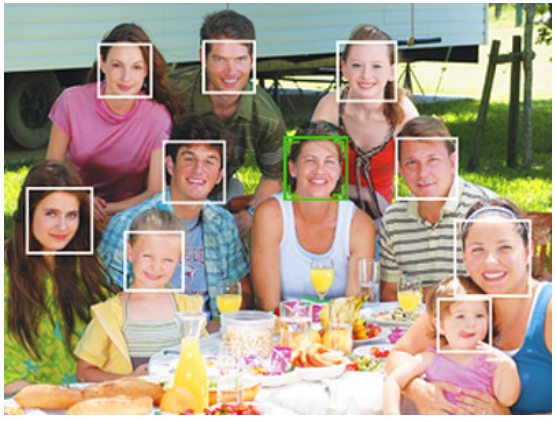

Face detection

Personal assistants

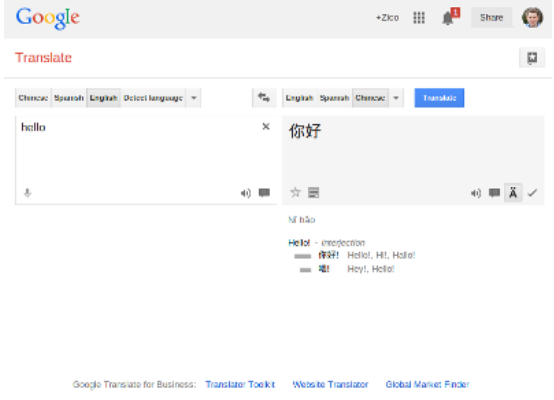

Machine translation

Logistics planning

Kidney exchange

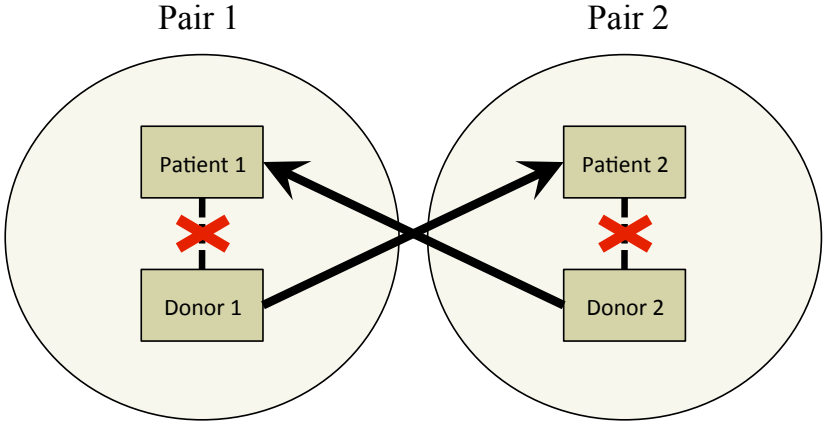

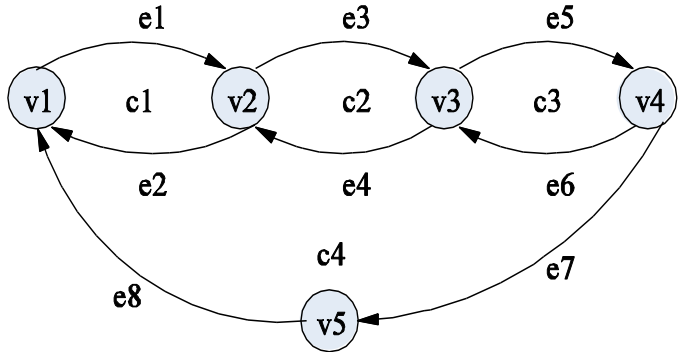

Autonomous ... “driving”

Automatic image captioning

"man in black shirt is playing guitar."

"construction worker in orange safety vest is working on road."

"two young girls are playing with lego toy."

Figure from (Karpathy and Fei-Fei, 2015)

Energy disaggregation

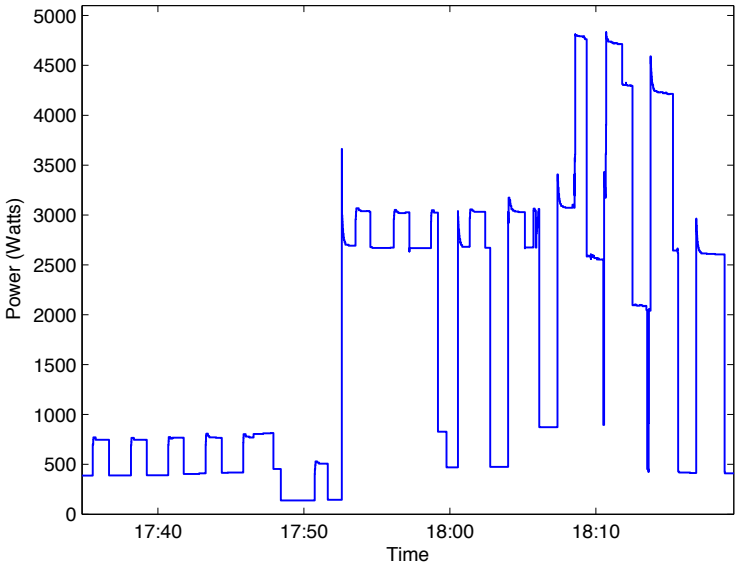

Observed signal

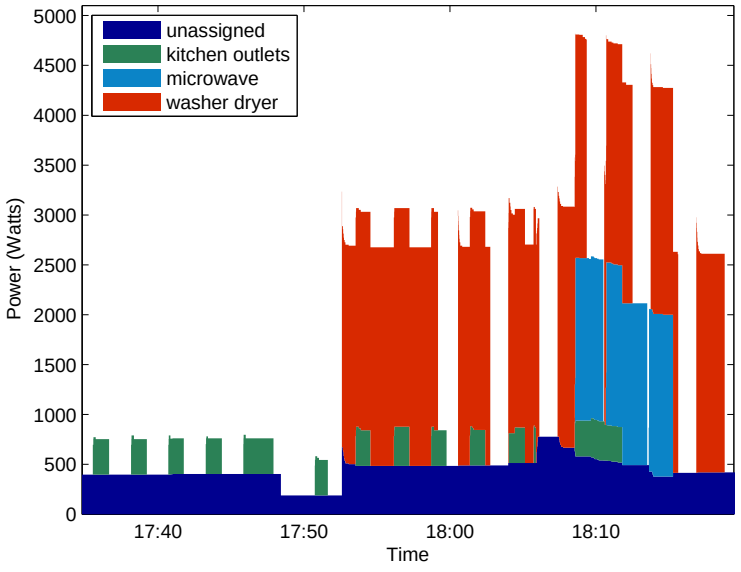

Predicted breakdown

Poker

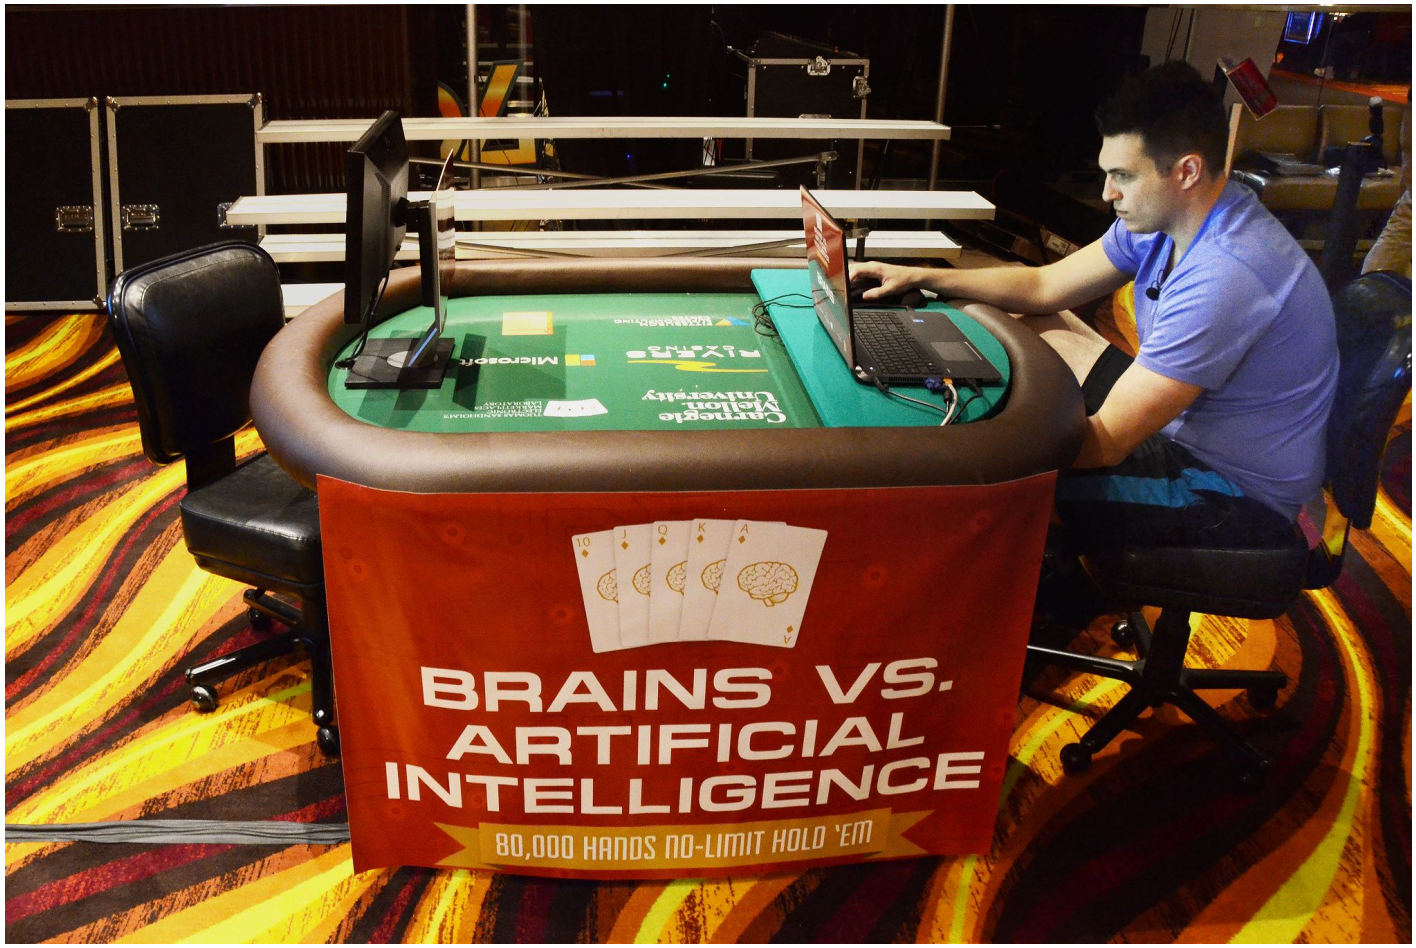
BRAINS VS.
ARTIFICIAL
INTELLIGENCE
80,000 HANDS NO-LIMIT HOLD 'EM

What should an AI course include?

Two philosophies:

1. Cover all topics in AI at a high level, about one lecture per topic

2. Cover only a few topics, in depth

This course does a bit of both, but probably more of the second

Topics covered in course

Topics covered: Uninformed and informed search, linear programming, mixed integer programming, probabilistic models and inference, reinforcement learning, continuous optimization, machine learning, deep learning, computational game theory

Complete schedule on course web page

Heavier focus on search, MIP, optimization, deep learning, computational game theory

Course materials

Schedule and lecture slides posted to website:

`http://www.cs.cmu.edu/~15780`

Textbook: Russell and Norvig, "Artificial Intelligence: A Modern Approach" (only Ch 1-4,6)

Piazza for class discussion

Video lectures available on Blackboard

Assignments and grading

Seven homeworks (2-3 questions each, including one programming problem)

- Homework submission done Autolab

- 8 late days for homeworks over semester, can use at most 3 for one homework

Midterm, final project (presentations during final)

Grading: 50% homeworks, 20% midterm, 25% project, 5% class participation

Class projects

A chance to explore an applied, theoretical, or algorithmic aspect of AI in more detail

To be done in groups of 2-3

Poster session during final exam period

Final writeup of 5 pages

(More material to be posted on website as dates arrive)

Instructors

Zico Kolter

Tuomas Sandholdm

TAs: Christian Kroer, Wennie Tabib, Daniel Guo, Guillermo Cidre

Recommended background

No formal pre-requisites

But, substantial programming background is required (assignments will be in Python)

Additional background in data structures and algorithms, linear algebra, probability will all be helpful, but not required

Honor code

Strict honor code with severe punishment for violators

CMU's academic integrity policy is here:
`http://www.cmu.edu/academic-integrity/`

You may discuss assignments with other students as you work through them, but writeups must be done alone

No downloading / copying of code or other answers is allowed

If you use a string of at least 5 words from some source, you must cite the source

Some parting thoughts

“Computers in the future may have only 1,000 vacuum tubes and weigh only 1.5 tons.” – Popular Mechanics, 1949

“Machines will be capable, within twenty years, of doing any work a man can do.” – Herbert Simon, 1965