


Introduction to Artificial Intelligence

CS 486/686

University of Waterloo

Plan for Today


- Talk about the excitement about AI today and our goals for the course
- Logistics
- What is AI? What will *this* course be covering?
- Rational Agent Paradigm



Microsoft researchers achieve new conversational speech recognition milestone

August 20, 2017 | By [Xuedong Huang](#), Technical Fellow, Speech and Language

DeepFace: Closing the Gap to Human-Level Performance in Face Verification Conference on Computer Vision and Pattern Recognition (CVPR)




Superhuman AI for multiplayer poker

Noam Brown^{1,2*} and Tuomas Sandholm^{1,3,4,5*}

¹Computer Science Department, Carnegie Mellon University Pittsburgh, PA 15213, USA. ²Facebook AI Research New York, NY 10003, USA. ³Strategic Machine, Inc. Pittsburgh, PA 15213, USA. ⁴Strategy Robot, Inc. Pittsburgh, PA 15213, USA. ⁵Optimized Markets, Inc. Pittsburgh, PA 15213, USA.

*Corresponding author. E-mail: noamb@cs.cmu.edu (N.B.); sandholm@cs.cmu.edu (T.S.)

In recent years there have been great strides in artificial intelligence (AI), with games often serving as challenge problems, benchmarks, and milestones for progress. Poker has served for decades as such a challenge problem. Past successes in such benchmarks, including poker, have been limited to two-player games. However, poker in particular is traditionally played with more than two players. Multiplayer games present fundamental additional issues beyond those in two-player games, and multiplayer poker is a recognized AI milestone. In this paper we present Pluribus, an AI that we show is stronger than top human professionals in six-player no-limit Texas hold'em poker, the most popular form of poker played by humans.




Technology

Stephen Hawking warns artificial intelligence could end mankind

By Rory Cellan-Jones

Technology correspondent



(AI)

Elon Musk: artificial intelligence is our biggest existential threat

The AI investor says that humanity risks 'summoning a demon' and calls for more regulatory oversight

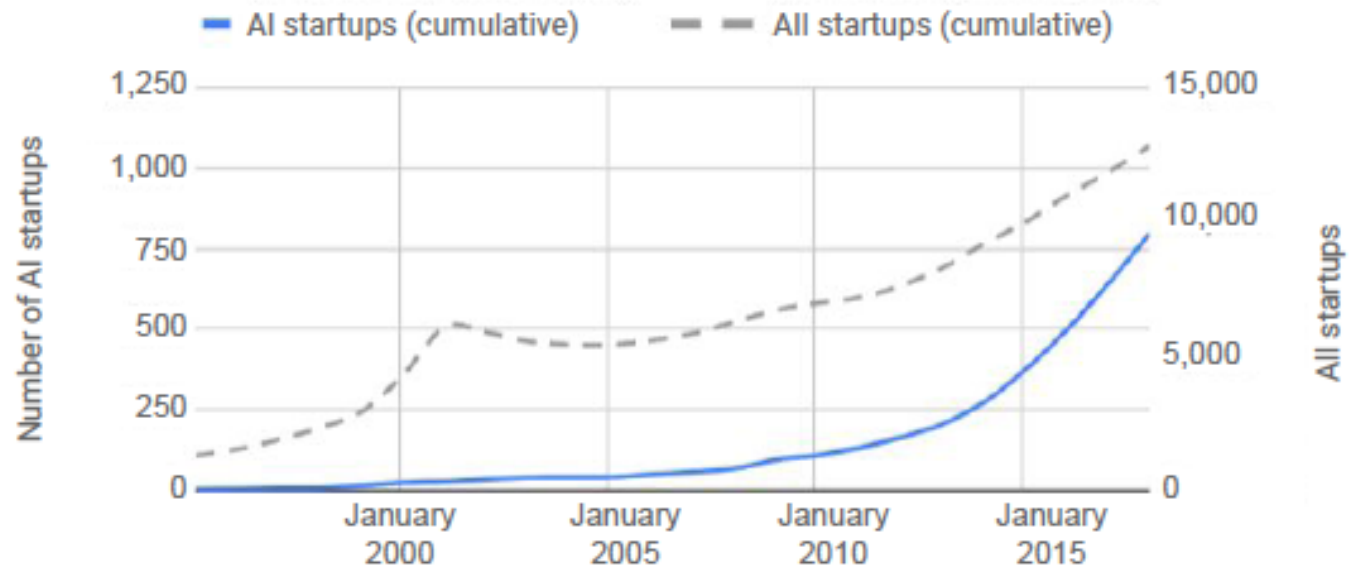
Excitement From Industry

- “An important shift from a mobile first world to an AI first world”
[CEO Sundar Pichai @ Google I/O 2017]
- Microsoft has created an AI and Research group as the 4th engineering division (by 2017 had 8000 people)
- Facebook has FAIR (Facebook AI Research)
- As well as Huawei, Baidu, IBM, Uber, Amazon, Tencent, Apple,.....

AI Index: Number of AI Startups

AI startups (U.S., January 1995 – January 2018)

Source: Sand Hill Econometrics



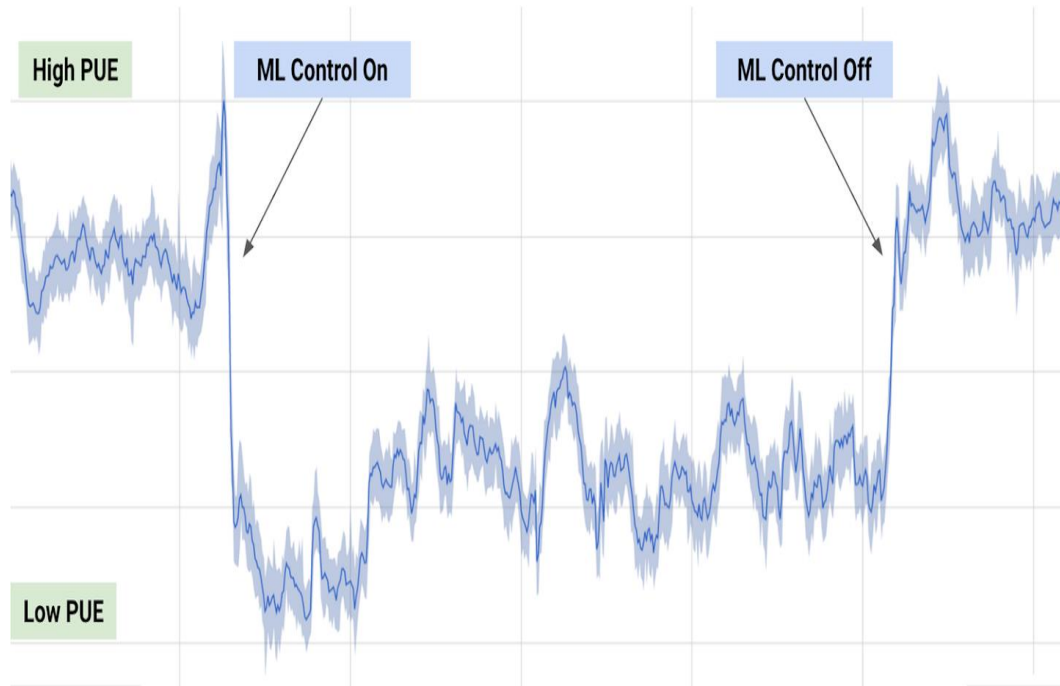
Note: The majority of the AI startups above develop AI systems. A minority use AI as an integral part of business, but do not develop the systems themselves. See appendix for more details.

Governments now have AI Strategies

- Canada: Pan-Canadian AI Strategy
- US: “AI holds the potential to be a major driver of economic growth and social progress” [White House Report, 2016]
- China: Released a domestic strategic plan in 2017 to become world leader in AI by 2030
- Russia: “Whoever becomes the leader in this sphere (AI) will be the ruler of the world” [Putin, 2017]

Goals for the Course

- Understand the principles and techniques for designing and implementing AI tools and system
- Encourage you to think deeply about the broader social and ethical context in which these systems will be deployed

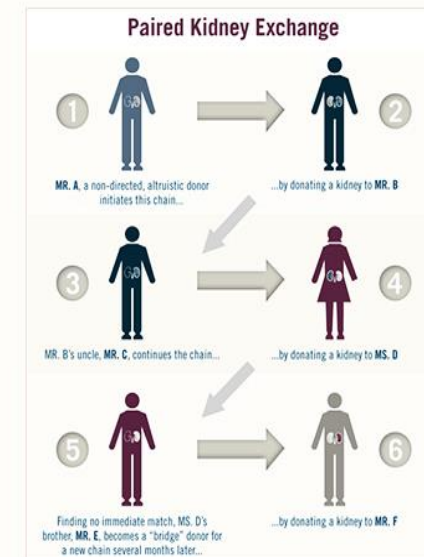


Saving energy by efficiently cooling data centres [DeepMind]



AI for Wildlife Protection [Tambe et al]

Improving Kidney exchanges [Dickerson et al]





Gender was misidentified in up to 1 percent of lighter-skinned males in a set of 385 photos.



Gender was misidentified in up to 7 percent of lighter-skinned females in a set of 296 photos.



Gender was misidentified in up to 12 percent of darker-skinned males in a set of 318 photos.



Gender was misidentified in 35 percent of darker-skinned females in a set of 271 photos.

Photos were selected from among those used in Joy Buolamwini's study.

Source: Joy Buolamwini, M.I.T. Media Lab

In modern artificial intelligence, data rules. A.I. software is only as

Distance/Angle	Subtle Poster	Subtle Poster Right Turn	Camouflage Graffiti	Camouflage Art (LISA-CNN)	Camouflage Art (GTSRB-CNN)
5' 0°					
5' 15°					
10' 0°					
10' 30°					
40' 0°					
Targeted-Attack Success	100%	73.33%	66.67%	100%	80%

Bias in Machine Translation

Fairness in Decision Making

California just replaced cash bail with algorithms

By Dawn Greenberg · September 4, 2018



[BUSINESS NEWS](#)

OCTOBER 9, 2018 / 11:12 PM / A YEAR AGO

Amazon scraps secret AI recruiting tool that showed bias against women

[Jeffrey Dastin](#)

8 MIN READ

Logistics: Who We Are

Kate Larson

- Professor in the AI Lab, Cheriton School of CS
- Research interests: Multiagent systems, Decision making under uncertainty, Reinforcement learning
- Also affiliated with DeepMind as a research scientist

Mathieu Doucet

- Associate Professor in Philosophy
- Research interests: Ethics, Moral Psychology, Political Philosophy, Bioethics

Great team of TA's

- Ehsan Ganjidoost, Alexandre Parmentier, Charupriya Sharma, Daniel Tamming, Colin Vandenhof, Ronghao Yang

Logistics

Lectures

- Mondays and Wednesdays 8:30-9:50 (E2 1736) and 11:30-12:50 (PHY 235)

Tutorials

- Fridays 8:30-9:20 (RCH 207), 9:30-10:20 (RCH 308), 10:30-11:20 (RCH 305)
 - Starting next Friday and running ~every second Friday. Check calendar on Learn for the schedule!

Logistics

Course Material

- Learn
- www.cs.uwaterloo.ca/~klarson/teaching/F19-cs468

Newsgroup

- We will use Piazza. Instructions for signing up are in the syllabus

Texts (Optional)

- Artificial Intelligence: A Modern Approach by S. Russell and P. Norvig
- Artificial Intelligence: Foundations of Computational Agents, D. Poole and A. Mackworth (available online)

Evaluation

CS 486

- 4 technical assignments: 28%
- 3 tutorial assignments: 15%
- Midterm Exam: 12%
- Final Exam: 45%
- Project (Optional): Up to 5 bonus marks

CS 686

- 4 technical assignments: 24%
- 3 tutorial assignments: 15%
- Midterm exam: 10%
- Final exam: 38%
- Project: 13%

What is AI?

What is AI?

Definitions tend to differ along two dimensions:

- Reasoning vs behaviour
- Fidelity to humans vs rationality

Systems that **think** like **humans**

Systems that **think** **rationally**

Systems that **act** like **humans**

Systems that **act** **rationally**


Systems that think like humans

Cognitive modelling approach

- How do humans think?
 - Brain imaging (MRI)
 - Introspection
- Cognitive science

System that act like humans

- Turing Test: Introduced by Alan Turing in the 1950 paper “Computing machinery and intelligence”



- Predicted that by 2000 a computer would have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI
- Suggested the major components of AI
 - Knowledge, Reasoning, Language Understanding, Learning

Rationality

Rationality

- An abstract ideal of intelligence
- “Doing the right thing” given what is known
 - Does not necessarily require “thinking”, but thinking often helps

My favorite definition

“Artificial Intelligence is that activity devoted to making machines intelligent, AND intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”

Nils J. Nilsson

For This Course

We are interested in intelligent systems where

a system is intelligent if and only if it acts rationally

Questions to think about


- Why do we care about behavior instead of thought processes and reasoning?
- Why do we measure success against rationality instead of against humans?

Topics We Will Cover

- Search
 - Heuristic search, CSPs, Local Search, Adversarial Search
- Supervised Learning
 - Principles of Learning, Decision Trees, Neural Networks
- Reasoning Under Uncertainty
 - Graphical Models (Bayesian Networks), Hidden Markov Models
- Learning Under Uncertainty
 - Statistical Learning, EM
- Decision Making Under Uncertainty
 - Markov Decision Processes, Reinforcement Learning, Decision Making in Multiagent Systems (Game Theory)

Back to Rationality: Rational Agent Paradigm

- A rational agent is an entity that perceives and acts:
 - Function from percepts to actions




- Performance measures
 - Goal achievement, resource consumption,...
- Caveat: Computational limitations and environmental constraints means we do not have perfect rationality

Task Environments

To design a rational agent the task environment must be defined

- Performance measures
- Environment
- Actuators
- Sensors

Performance Measures Can Be Tricky



Percepts: [Location, Dirty or Clean]

Actions: Right, Left, Vacuum, NoOp, Dump

Function: ([A,Clean],Right), ([A, Dirty], Vacuum), ([B, Dirty], Vacuum), ([B, Clean], Left)...

Properties of the Task Environment

Properties of the task environment determine what AI approach is most suitable

- **Fully Observable vs Partially Observable:**
 - Does the agent's sensors give it access to the complete state of the environment?
- **Deterministic vs Stochastic**
 - Is the next state completely determined by the current state and action executed?
- **Episodic vs Dynamic**
 - Does the current decision or action influence future decisions or actions?
- **Discrete vs Continuous**
 - How are states, time, actions modelled?
- **Static vs Dynamic**
 - Is the environment changing as the agent is planning what to do next?
- **Single Agent vs Multiagent**