

CS10: The Beauty and Joy of Computing


Artificial Intelligence



Anna Rafferty
(Slides adapted from Dan Garcia)
3 April 2013


What I Do...

- Model human learning using machine learning
- Adaptive instruction and feedback in computer-based educational environments
- E.g., diagnose a student's knowledge by watching her play a game




Lecture Overview

- Definition
- Some AI history
- Tour of areas of AI
- Turing Test and the Chinese room




AI Definition by John McCarthy

- “Getting a computer to do things which, when done by people, are said to involve intelligence”
- Finesses the idea of whether a computer has consciousness, whether they have rights, etc.



A little history...

- Early AI (1956-early 1970s): symbolic reasoning and lots of optimism
- Neural nets (but very simple)



Person: PICK UP A BIG RED BLOCK.
Computer: OK. (does it)
Person: GRASP THE PYRAMID.
Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.
Person: FIND A BLOCK WHICH IS TALLER THAN THE ONE YOU ARE HOLDING AND PUT IT INTO THE BOX.
Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.
Computer: OK. (does it)

<http://hci.stanford.edu/winograd/shrdlu/>

Clicker Question

- Which of these rules is true for all dogs?
 - (A) Has four legs
 - (B) Has fur
 - (C) Barks
 - (D) None of the above



(Image from: http://vision.stanford.edu/resources_links.html)


Revival of AI

- Brittle rules break down with complexity of real world
- Probability and uncertainty
- No “dog rule” – instead: what is the probability that the thing we’re seeing is a dog?



What intelligent things do people do?

- Planning
- (Machine) Learning
- Natural Language Processing
- Motion and manipulation
- Perception
- Creativity
- General Intelligence



en.wikipedia.org/wiki/Artificial_intelligence


Tour of AI Applications

- Questions to keep in mind:
 - How would you evaluate how well a machine performed on the tasks we talk about?
 - How can blending artificial and human intelligence make tasks simpler, even if the AI isn’t perfect?



Planning

- Range of intelligence
 - Low: simple heuristics
 - Medium: pathfinding
 - High: Learns from player
- Dynamic difficulty - adjust to player’s skill
- Allocation of resources
 - E.g., choose what land resources to give to which conservation projects



www.businessweek.com/innovate/content/aug2008/1d20080820_123140.htm
en.wikipedia.org/wiki/Dynamic_game_difficulty_balancing
en.wikipedia.org/wiki/Game_artificial_intelligence
queue.acm.org/detail.cfm?id=971593



Machine Learning

- “A program learns if, after an experience, it performs better”
- What are the right generalizations to make given the data you’ve seen and the task you’re completing?




en.wikipedia.org/wiki/Machine_learning



Machine Learning

- Algorithm Types
 - Supervised learning
 - Give a system input & output training data, and it produces a classifier
 - Unsupervised learning
 - Determine how data is organized or clustered
 - Reinforcement learning
 - No training data, real-time corrections adjust behavior




en.wikipedia.org/wiki/Machine_learning



Example: Deep Learning


- Combines supervised and unsupervised learning: Learn the right *representations* for input -> output



Benefiting from Big Data



Translation



Computer vision

Netflix Prize

Leaderboard

Rank	Team Name	Best Score	% Improvement	Last Update Time
1	Chiranjeev	0.816	0.00%	2008-02-22 11:54
2	Chiranjeev	0.816	0.00%	2008-02-22 11:54
3	Chiranjeev	0.816	0.00%	2008-02-22 11:54
4	Chiranjeev	0.816	0.00%	2008-02-22 11:54
5	Chiranjeev	0.816	0.00%	2008-02-22 11:54
6	Chiranjeev	0.816	0.00%	2008-02-22 11:54
7	Chiranjeev	0.816	0.00%	2008-02-22 11:54
8	Chiranjeev	0.816	0.00%	2008-02-22 11:54
9	Chiranjeev	0.816	0.00%	2008-02-22 11:54
10	Chiranjeev	0.816	0.00%	2008-02-22 11:54
11	Chiranjeev	0.816	0.00%	2008-02-22 11:54
12	Chiranjeev	0.816	0.00%	2008-02-22 11:54
13	Chiranjeev	0.816	0.00%	2008-02-22 11:54
14	Chiranjeev	0.816	0.00%	2008-02-22 11:54
15	Chiranjeev	0.816	0.00%	2008-02-22 11:54
16	Chiranjeev	0.816	0.00%	2008-02-22 11:54
17	Chiranjeev	0.816	0.00%	2008-02-22 11:54
18	Chiranjeev	0.816	0.00%	2008-02-22 11:54
19	Chiranjeev	0.816	0.00%	2008-02-22 11:54
20	Chiranjeev	0.816	0.00%	2008-02-22 11:54

Recommendation

Vision

- Tasks related to understanding images/camera input



Figure/ground segmentation



Pedestrian detection




phoning

Action recognition

(Some images from Berkeley vision group)

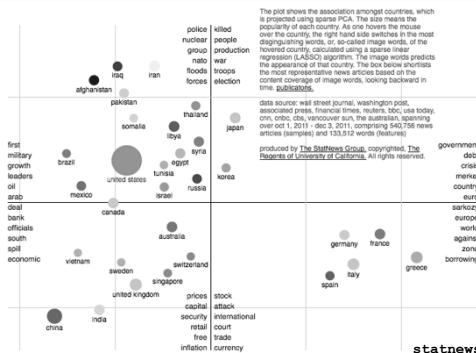
Natural Language Processing

- Known as "AI-complete" problem
 - (Often) requires extensive knowledge of world
- Statistical NLP
 - Correcting/guessing text
 - Suggesting news stories
 - Finding articles that are similar to one another
 - Translate or paraphrase texts



en.wikipedia.org/wiki/Natural_language_processing

Unsupervised Learning Example



Robotics

- For many, the coolest and scariest part of AI
- Combines fields of AI/CS
 - Speech recognition
 - Synthetic voice
 - Machine vision
 - Planning
 - HCI




TOPIO, the ping-pong playing robot




Assistive robots



Surgical robots



Autonomous helicopter



UC Berkeley's towel-folder

en.wikipedia.org/wiki/Robotics

Recap

- All of these applications are tough because they require:
 - Knowing about context
 - Uncertainty about input
 - Intensive computations
- But AI has been relatively successful at making progress (and in some cases, better than people!)




Clicker Question

- What would a “truly intelligent” AI system look like?
 - (A) Behaves in an optimal or rational manner
 - (B) Behaves similarly to people – when it makes errors, those errors are similar to people’s errors
 - (C) Carries out the same type of processing (mental representations) people do – i.e., thinks like people

Turing Test for Intelligence

- In 1950, Turing defined a test of whether a machine could “think”
- “A human judge engages in a natural language conversation with one human and one machine, each of which tries to appear human. If judge can’t tell, machine passes the Turing test”
- John Searle argued against the test via the Chinese room experiment, in which someone carries on a conversation by looking up phrases in a book. Does that person understand Chinese?



en.wikipedia.org/wiki/Turing_test

Clicker Question

- How would you respond to Searle’s Chinese room experiment?
 - (A) The system as a whole understands Chinese
 - (B) The man doesn’t understand Chinese, but if he had a way to connect with the outside world (rather than just receiving strings of symbols), he could understand Chinese
 - (C) We must be missing something about “understanding” since the argument implies that brains, which are collections of neurons, cannot understand

Summary

- AI systems excel in things computers are good at
 - Big data (using web to parse language)
 - Constrained worlds (chess, math)
- It’s getting better at...
 - Language understanding
 - Real-time robotics
- Lots more applications that I didn’t have time to talk about!
- CS188: Artificial Intelligence
 - One of the most popular courses on campus!
- CogSci131: Computational Models of Cognition



Thanks! Feel free to email me with questions at rafferty@cs.berkeley.edu