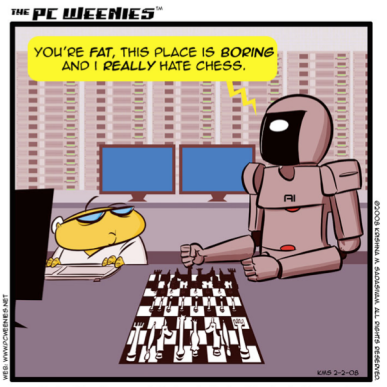

http://www.bbspot.com/comics/PC-Weenies/2008/02/3248.php

Intro to AI

CS51A
David Kauchak
Spring 2019

Adapted from notes from:
Sara Owsley Sood

Admin

Tuesday mentor hours back to 5-7pm

My office hours:

- M: 3-3:50pm
- **T: 11am-12**
- **Th: 11am-12:30**
- **F: 10-11am**

AI is a huge field

What is AI (artificial intelligence)…

AI is a huge field

What is AI…

One definition:

"*Building programs that enable computers to do what humans can do.*"

For example:
read, walk around, drive, play games, solve problems, learn, have conversations…

A broader definition

"Building programs that enable computers to do do *intelligent* things"

	human	vs.	rational
thinking	**Think like a human** Cognitive Modeling		**Think rationally** Logic-based Systems
vs.			
acting	**Act like a human** Turing Test		**Act rationally** Rational Agents

How is AI viewed in popular media?

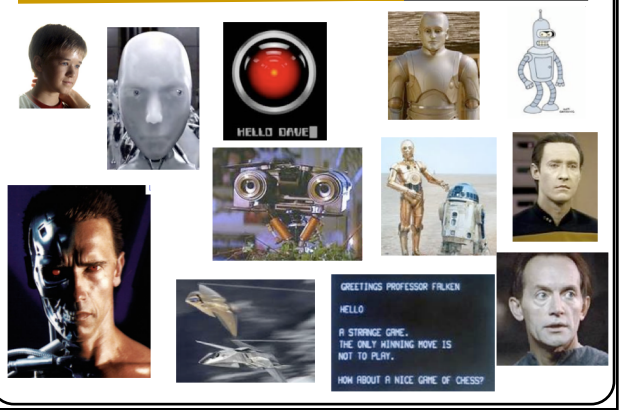

What challenges are there?

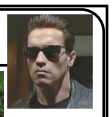

What challenges are there?

Perception
- perceive the environment via sensors

Computer vision (perception via images/video)
- process visual information
- object identification, face recognition, motion tracking

Natural language processing and generation
- speech recognition, language understanding
- language translation, speech generation, summarization

What challenges are there?

Knowledge representation
- encode known information
- water is wet, the sun is hot, Dave is a person, …

Learning
- learn from environment
- What type of feedback? (supervised vs. unsupervised vs. reinforcement vs …)

Reasoning/problem solving
- achieve goals, solve problems
- planning
- How do you make an omelet? I'm carrying an umbrella and it's raining… will I get wet?

Robotics
- How can computers interact with the physical world?

What can we currently do?

What can we currently do?

Understand spoken language?

speech recognition is really good, if:
- restricted vocabulary
- specific speaker with training

Gotten quite good in the last few years and shows up in lots of places:
- Mobile: Siri, Ok Google, etc.
- Home assistants: Alexa, Google Home

What does the spoken language actually mean (language understanding)?
- much harder problem!
- many advances in NLP in small things, but still far away from a general solution

What can we currently do?

Speak?

Understandable, but you wouldn't confuse it for a person

Can do accents, intonations, etc.

Better with restricted vocabulary

Loquendo

- https://www.nuance.com/omni-channel-customer-engagement/voice-and-ivr/text-to-speech.html
- Dealing with facial expression is challenging

Kismet (MIT)

What can we currently do?

Drive a car?

What can we currently do?

Drive a car?

Freeway driving is relatively straightforward

Off-road a bit harder

- see DARPA grand challenges (2004, 2005)

And urban driving is even trickier

- See DARPA urban challenge (2007)
- Google's autonomous vehicle

Hint: there's a connection here

What can we currently do?

Drive a car?

https://www.google.com/selfdrivingcar/

https://www.tesla.com/autopilot

Uber

…

What can we currently do?

Drive a car?

Many driver assist technologies:

- Automatic breaking
- Automatic pedestrian detection
- Lane drift avoidance
- "smart" cruise control
- Blind spot warning
- …

What can we currently do?

Identify emotion?

This is hard!

Some success in text

- movie reviews (assignment 7!)
- blogs
- twitter
- dealing with sarcasm is hard

Some success with faces

- strongly biased by training data
- works best when exaggerated

What can we currently do?

Reasoning?

Success on small sub-problems

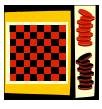

General purpose reasoning is harder

- Wolfram Alpha
- OpenCyc

What can we currently do?

Walk?

Robots have had a variety of locomotion methods

Walking with legs, is challenging

- Differing terrains, stairs, running, ramps, etc.

Recently, a number of successes

- Honda's Asimo
  - https://www.youtube.com/watch?v=SARB9OI_W
- Sony QRIO
  - http://www.youtube.com/watch?v=9vwZ5FQEUFg
- Boston Dynamic's Atlas
  - https://www.youtube.com/watch?v=hSjKoEva5bg

When will I have my robot helper?

What can we currently do?

What can we currently do?



What can we currently do?

Fold a pile of towels?

UC Berkeley towel folding robot:

http://www.youtube.com/watch?v=gy5g33S0Gzo

How do we make computers "intelligent?"

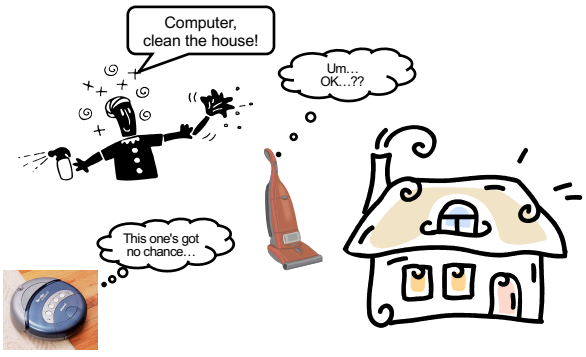

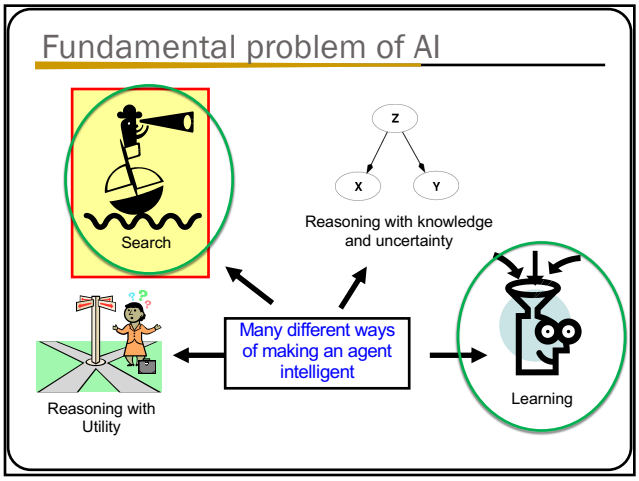
Fundamental problem of AI
Search
Z
X
Y
Reasoning with knowledge and uncertainty
Many different ways of making an agent intelligent
Reasoning with Utility
Learning