

Artificial Intelligence:

What is it, and why should I care?

Jocelyn Huang



The Big Picture
Methods
Applications

The Big Picture

Overview

What is it?

Some history

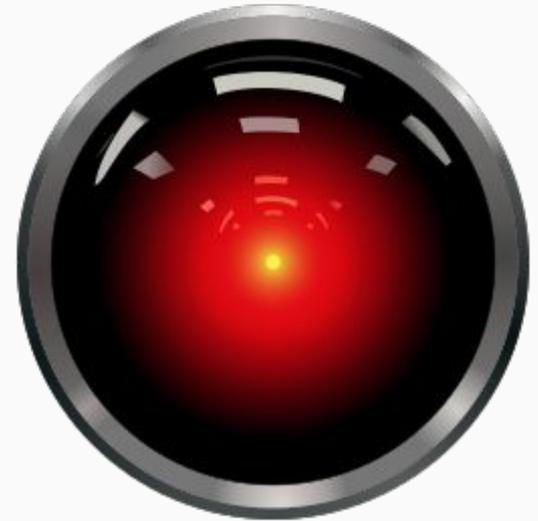
Methods

Applications

What is AI?

We think of...

- Video games/Sci-Fi?
 - Like GLaDOS or HAL, or an NPC
- Facial recognition?
- Siri? Google?
- Machine learning?



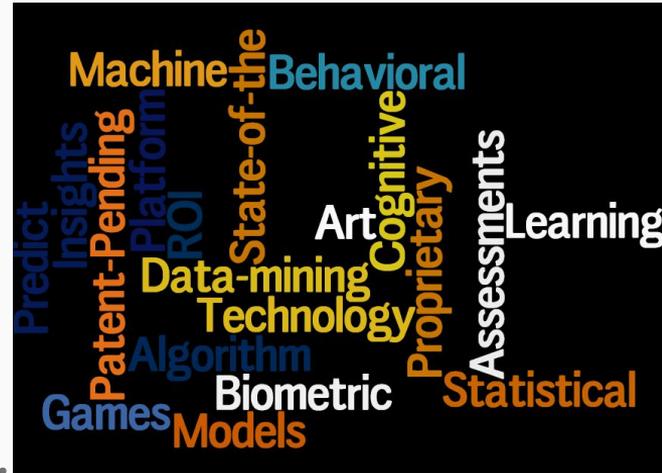
What *is* AI?

- Broad topic that includes/is related to:
 - ML, NLP, big data, robotics, data mining...
 - Look, buzzwords! ----->

- Boundaries blurry

- **Mainly:**

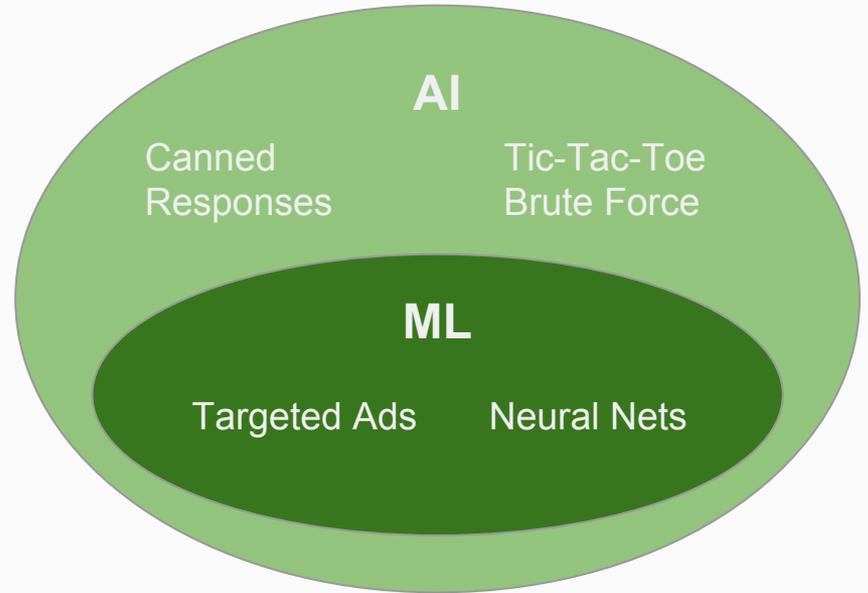
Thinking, problem-solving machines.



Artificial Intelligence vs. Machine Learning?

Okay, let's get this cleared up.

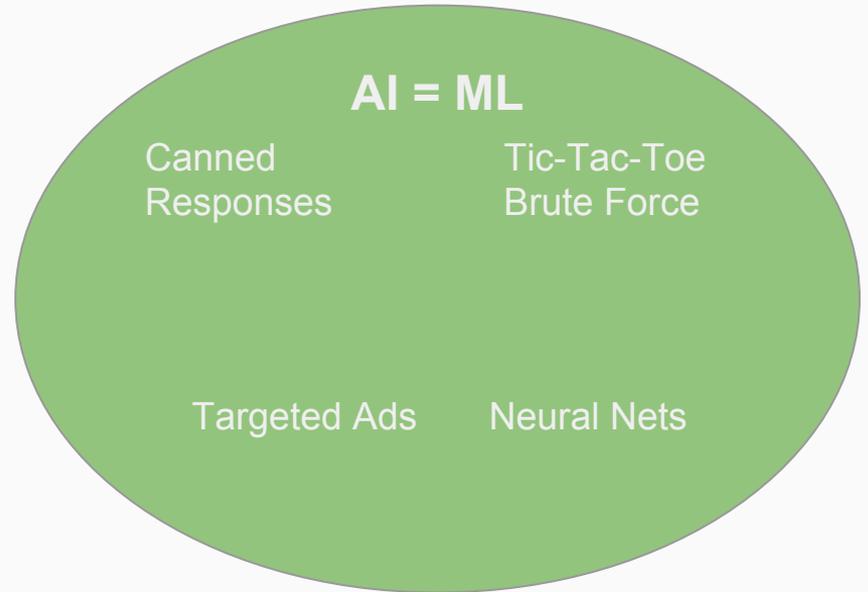
- In the past: $ML \subseteq AI$



Artificial Intelligence vs. Machine Learning?

Okay, let's get this cleared up.

- In the past: $ML \subseteq AI$
- Now: $ML \cong AI$



(Very) Brief history!

The Dream: thinking machines!

(sci-fi, mythology, etc.)

Emergence of artificial intelligence:

Alan Turing: “Can machines think?”

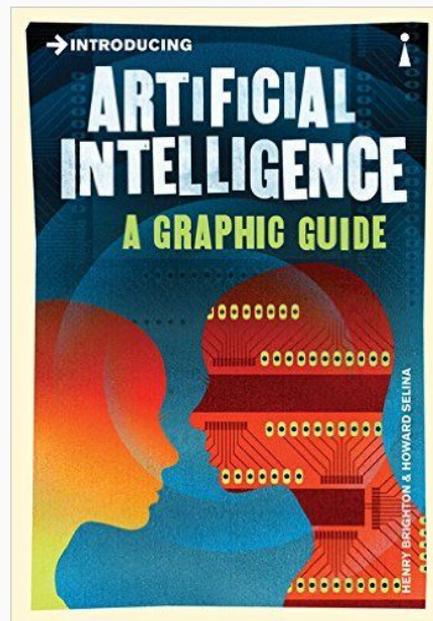
You may have heard of **Newell** and **Simon**, and **J.C. Shaw**

=> Logic Theory Machine and General Problem Solver

More Names: Marvin Minsky, John McCarthy

Emergence of machine learning

For more... there’s a great graphic guide!



Where do we draw the line for “intelligence”?

Can a machine have consciousness? How do we tell?

This is all very philosophical.

- Searle’s Chinese Room (1980)
-> “Hollow Shell” argument
- Turing Test (1950)

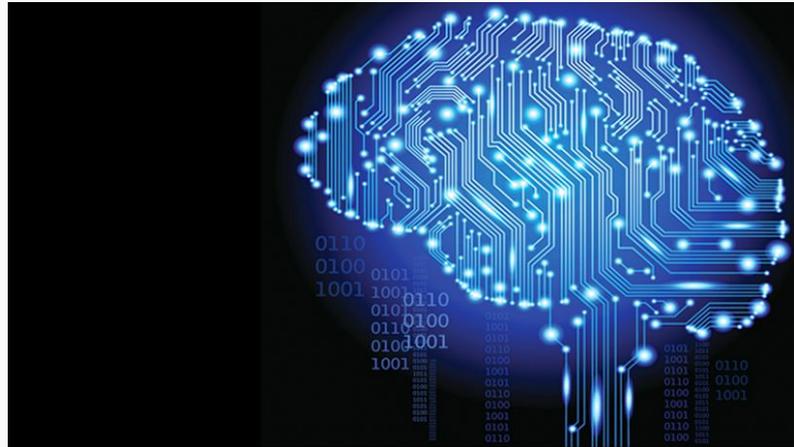
=> Existential Crisis?!



But in the meantime...

...We're computer scientists and we just like programs that do cool things!

So let's get right to it.



The Big Picture

Methods

Categories

Techniques

Applications

A quick overview

Goals:

- General Intelligence
- Social Intelligence
 - Affective Computing
- Computer Vision
- Natural Language Processing (NLP)
- Machine Learning
 - Supervised/Unsupervised Learning

Social Intelligence

Also called Affective Computing

Goals:

- Figuring out emotion
- Imitating emotion

FloBi, Pepper



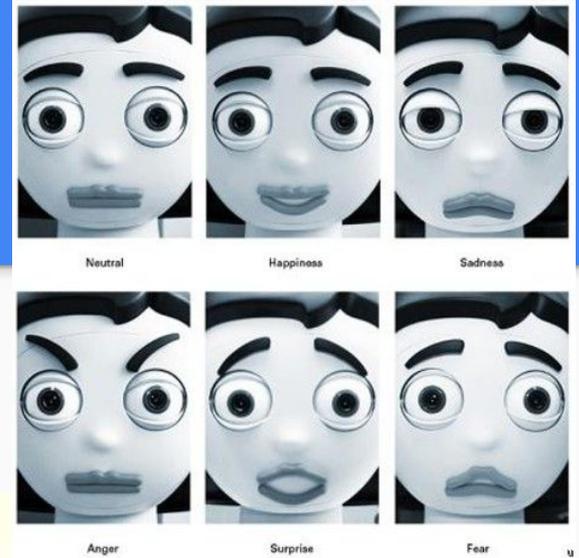
Social Intelligence

Also called Affective Computing

Goals:

- Figuring out emotion
- Imitating emotion

FloBi, Pepper

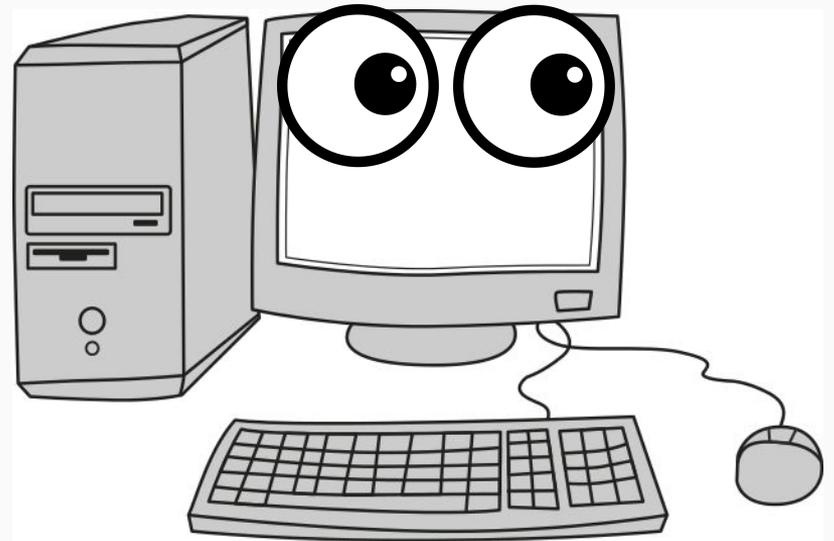


Computer Vision

Let's have machines look at things!

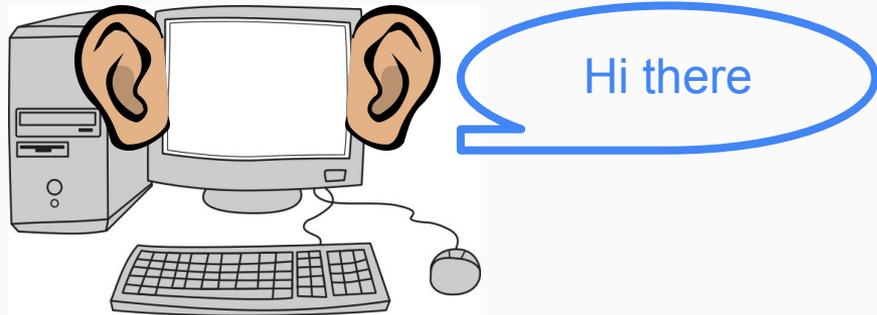
Usages:

- Facial recognition
- Image categorization
- Inspection and autonomous control
- Navigation
 - Self-driving cars
 - Robotics



Natural Language Processing (NLP)

- Programs typically need formatted input
- What about articles? Product reviews? Facebook posts?
 - “My name is Jocelyn Huang, and I’m an undergraduate CS student at Carnegie Mellon.”
 - Name: Jocelyn Huang .. Major: Computer Science .. School: Carnegie Mellon University
- Human interaction



NLP

Sentiment Analysis

Deals with emotions

→ A user's attitude towards a topic

Hurdles:

- Subtle differences:
 - “They wouldn't let my dog stay here” vs. “I wouldn't let my dog stay here”
- Sarcasm
 - “Oh, great, they lost my luggage”
 - Cornell students: *sentiment shift*

NLP

Semantics

What are the humans saying?

- Summarization
- Translation
- Understanding and Response

This is difficult.

Context-dependent:

“People like Harrison Ford act.”

Machines lack common sense:

“Police help dog bite victim.”

“Stolen painting found by tree.”

Machine Learning (ML)

Heavily statistical:

- Training Data...
- to Test Data...
- and Verification and Minimizing Error

Supervised vs. Unsupervised Learning



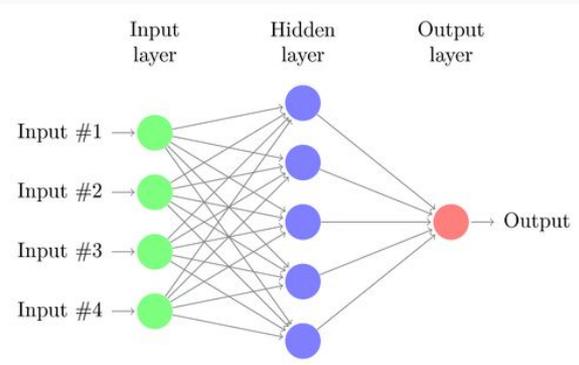
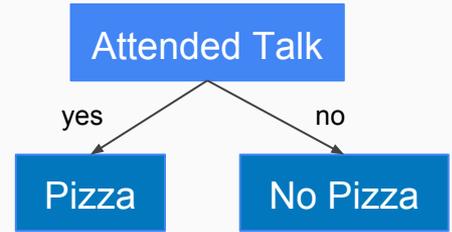
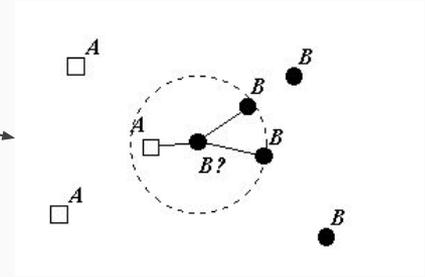
ML

A Few Techniques:
A quick rundown

{H,H,T,H,T,T,T,T,H}

$$\Rightarrow \hat{p} = \#H / (\#T + \#H) = 4/9$$

- K-Nearest Neighbors
- Decision Trees
- Naive Bayes
- Neural Nets

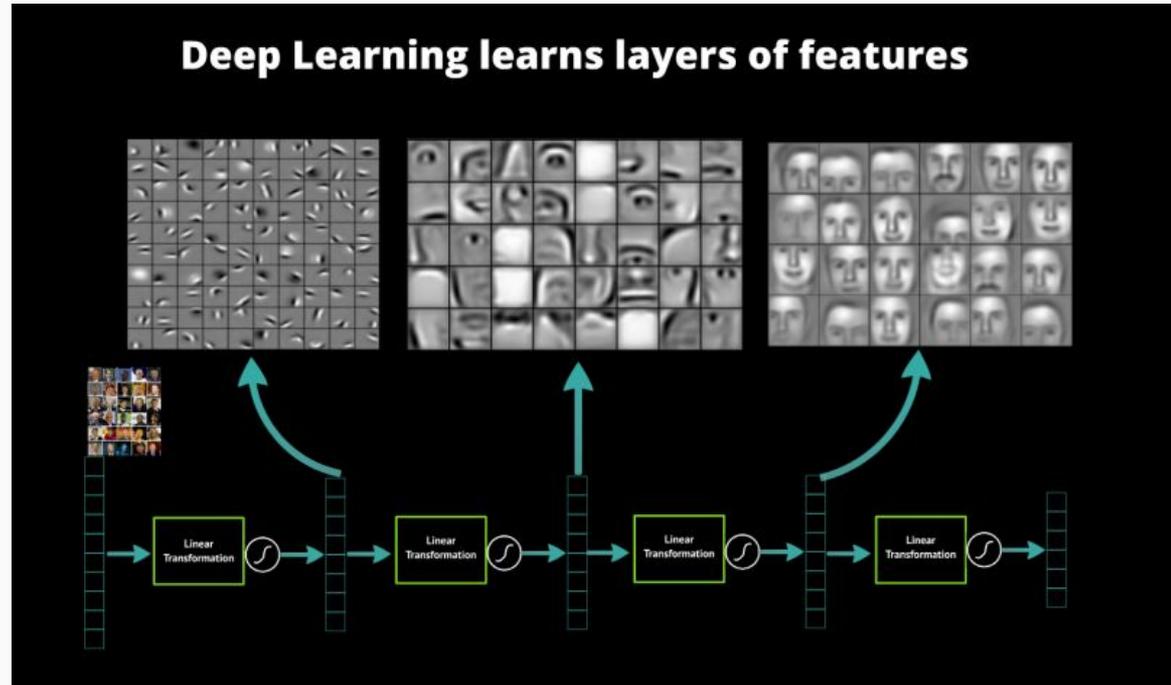


ML

Deep Learning

A neural net with more than one hidden layer

They're surprisingly powerful...



The Big Picture

Methods

Applications

A big list!

Specific examples

When things go wrong

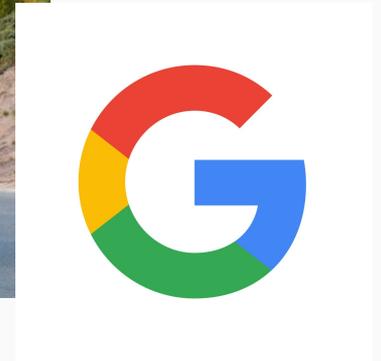
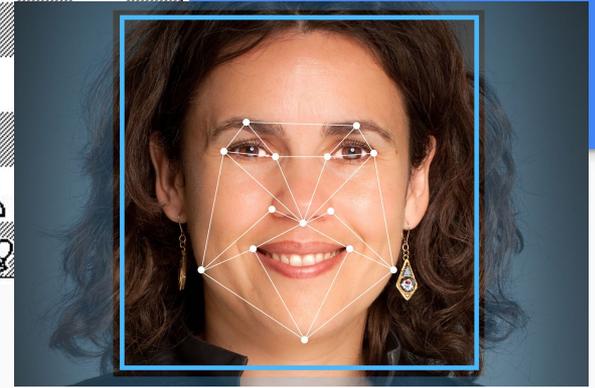
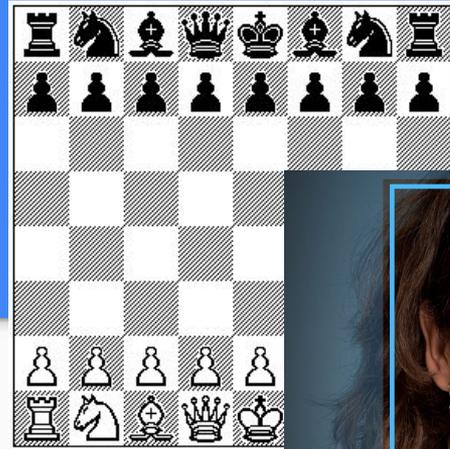
Applications

Okay, this is cool and all, but why do we care?

What can machine learning even do?

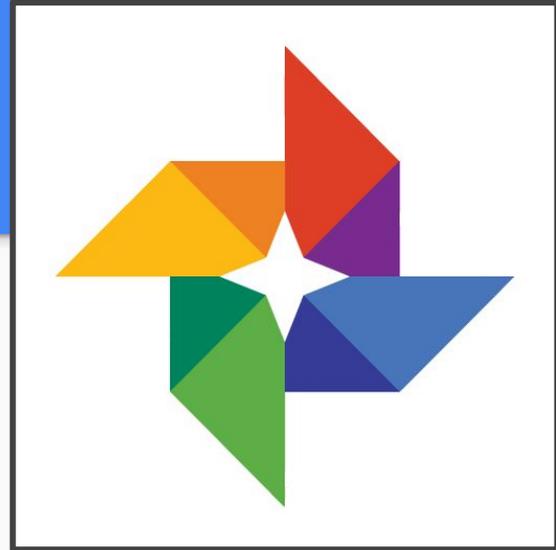
Applications

- Game playing (we'll get back to this)
- Facial recognition
- IBM Watson (more on this later)
- Google search (maybe you've heard of this? ;)
- Google ads
- Google spam filters
- Google photos
- Google Brain
- Pretty much all of Google (yeah.)
- Stocks: automated investment trading
- Self-driving cars
- The list goes on...



Google Photos

- Released Summer 2016
- What's it got?
 - Machine vision (image recognition)
 - Machine learning (classification)
- The Cool:
 - Classification -> Search
 - Face recognition and classification
 - Location recognition (landmarks, languages, even architecture)



When Google Photos goes wrong

- The Bad:
 - Google Photos being politically incorrect

Google Mistakenly Tags Black People as 'Gorillas,' Showing Limits of Algorithms

ARTICLE

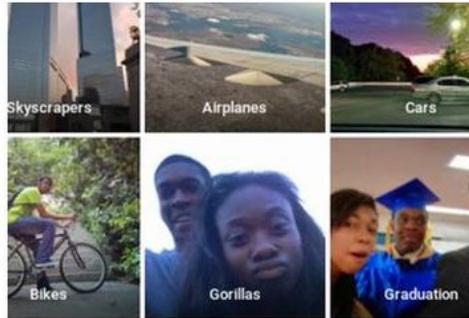
COMMENTS

ALGORITHMS ARTIFICIAL INTELLIGENCE COMPUTER VISION GOOGLE PHOTOS

Email Print

Facebook 4.3k Twitter

By ALISTAIR BARR [CONNECT](#)



Black programmer Jacky Alciné said on [Twitter](#) [TWTR-6.52%](#) that the new [Google](#) [GOOGL-1.39%](#) Photos app had tagged photos of him and a friend as gorillas. —Jacky Alciné and Twitter

Google is a leader in artificial intelligence and machine learning. But the company's computers still have a lot to learn, judging by a major blunder by its Photos app this week.

The app tagged two black people as "Gorillas," according to Jacky Alciné, a Web developer who spotted the error and tweeted a photo of it.

"Google Photos, y'all f**ked up. My friend's not a gorilla," [he wrote on Twitter](#).

Google apologized and said it's tweaking its algorithms to fix the problem.

Games!

- Deep Blue vs. Kasparov (1996 loss, 1997 win)
 - Arguably the most famous
- Recently, Go
 - Google (DeepMind's AlphaGo)
 - Complex! 10^{170} configurations



LearnFun & PlayFun

- A CMU alum's creation!
 - Tom Murphy (tom7)
- Submitted for SIGBOVIK 2013
- Learns and plays (video) games!
 - Deduce objective function from recording of inputs
 - Maximize.
- Pretty impressive:
 - <https://youtu.be/Q-WgQcnessA?t=4m2s>
 - <https://youtu.be/Q-WgQcnessA?t=13m14s>

When PlayFun goes wrong

- Pit jumping: <https://youtu.be/Q-WgQcnessA?t=42s>

IBM Watson



- Everyone's favorite *Jeopardy!* champion
- What's it got?
 - Natural Language Processing
 - Rule Based AI & Machine Learning
- The Cool:
 - 200 M pages of content (4TB)
 - Robotic finger
 - 4-Step Decision Process
 - Amazing NLP that beat the best in *Jeopardy!*



IBM Watson

Let's take a look:

<https://www.youtube.com/watch?v=P18EdAKuC1U>

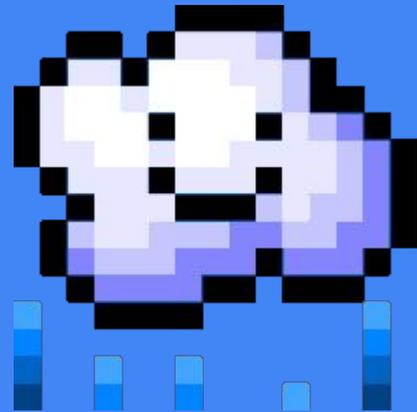
When Watson goes wrong

- The Bad:
 - Watson's Final Jeopardy slip-up
<https://www.youtube.com/watch?v=C5Xnxjq63Zg>
→ But why?
 - Less weight on category
 - "US Cities" wasn't in the clue itself
 - Cities called "Toronto" in both the US and Canada

Where's Watson Now?

- Cooking (Chef Watson)
<https://www.ibmchefwatson.com/>
- Crime solving (Discovery Advisor--Data Analytics) <http://www.ibm.com/smarterplanet/us/en/ibmwatson/discovery-advisor.html>
- Healthcare Help (Q&A)
<http://www.ibm.com/smarterplanet/us/en/ibmwatson/health/>

So it's not all sunshine and
flowers...



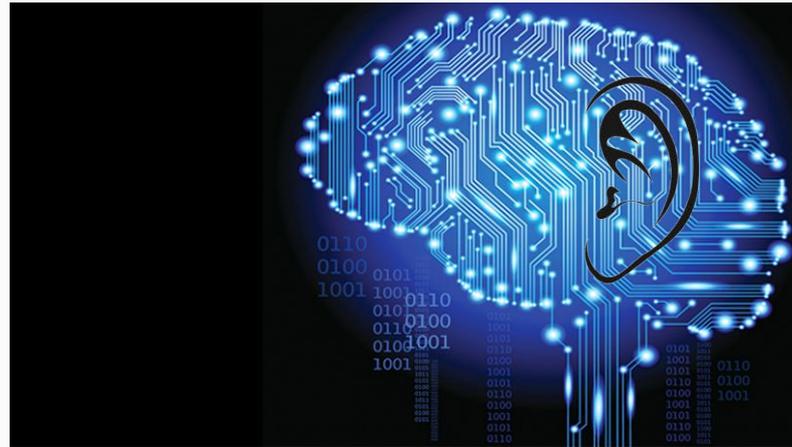
Still worth it...?

- Recommendation software
- Medical diagnoses
- Game playing
- Facial recognition
- IBM Watson
- Google
- Autocorrect
- Stocks: automated investment trading
- Self-driving cars
- Handwriting recognition
- ...
- And so on.

You like these things,
don't you?

Thanks for listening!

Any questions?



Further Reading

Methods:

- Sarcasm detection: <http://arstechnica.com/information-technology/2016/01/snark-attack-cornell-students-teach-software-to-detect-sarcasm/>
- ML: <http://blog.echen.me/2011/04/27/choosing-a-machine-learning-classifier/>
- Deep Learning: <https://www.datarobot.com/blog/a-primer-on-deep-learning/>

Further Reading

Applications:

- IBM Watson:

<http://www.bobblum.com/ESSAYS/COMPSCI/Watson.html>

<https://www.youtube.com/watch?v=DywO4zksfXw>

Slip-up: <http://asmarterplanet.com/blog/2011/02/watson-on-jeopardy-day-two-the-confusion-over-an-airport-clue.html>

- PlayFun:

<http://www.cs.cmu.edu/~tom7/mario/>

<http://www.cs.cmu.edu/~tom7/mario/mario.pdf>