Lecture 21: AI and Machine Learning

CS 51P

November 25, 2019

What is Artificial Intelligence?

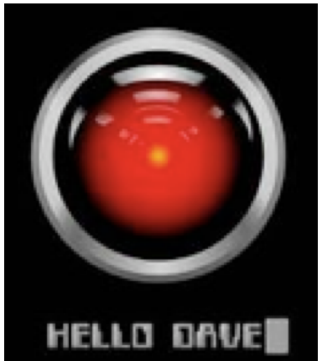

What is Artificial Intelligence?

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	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

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, professors are people, …

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?

Machine Learning is…

Machine learning is about predicting the future based on the past.

-- Hal Daume III

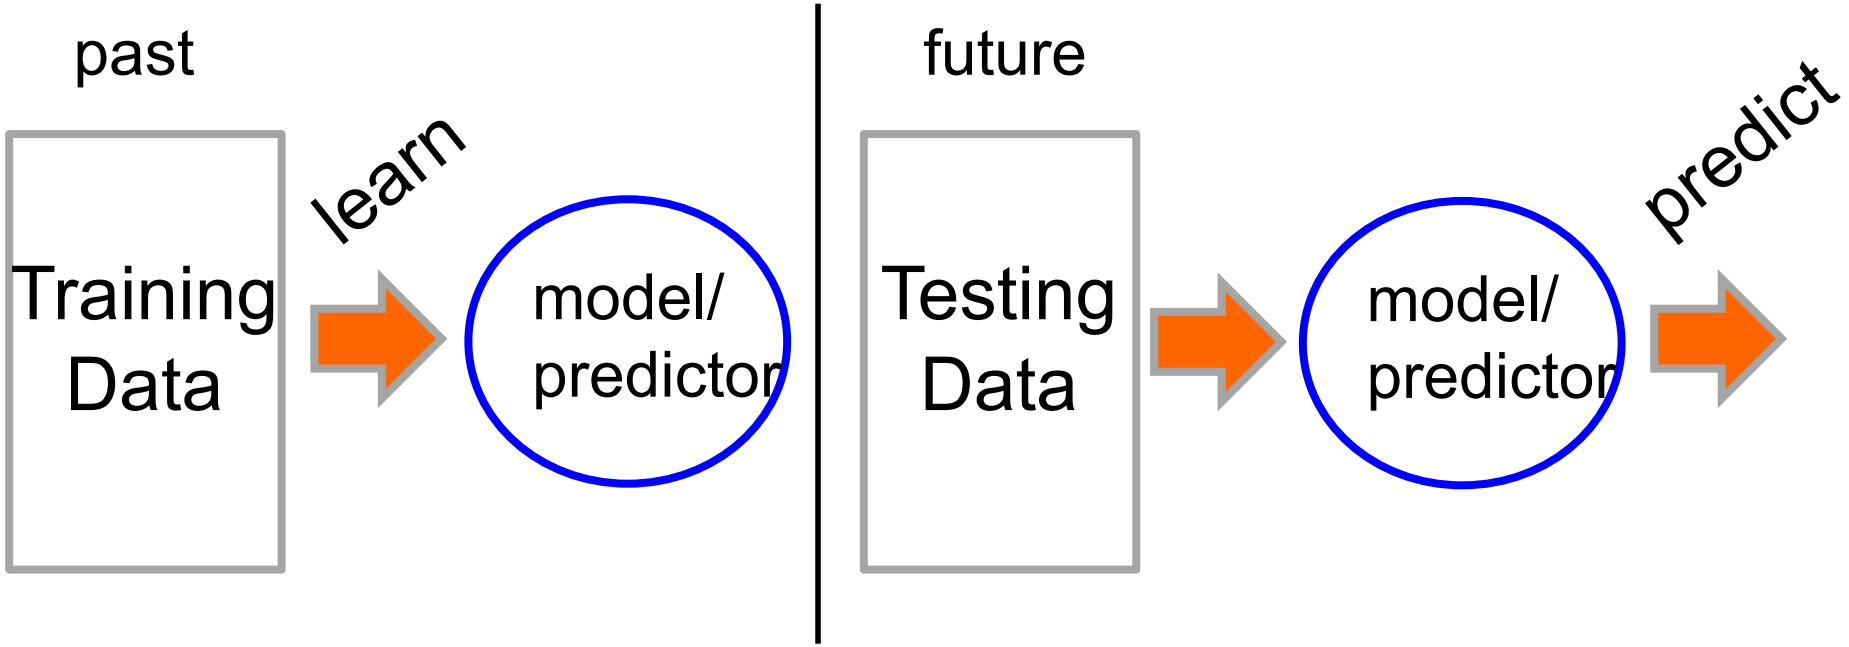

Data

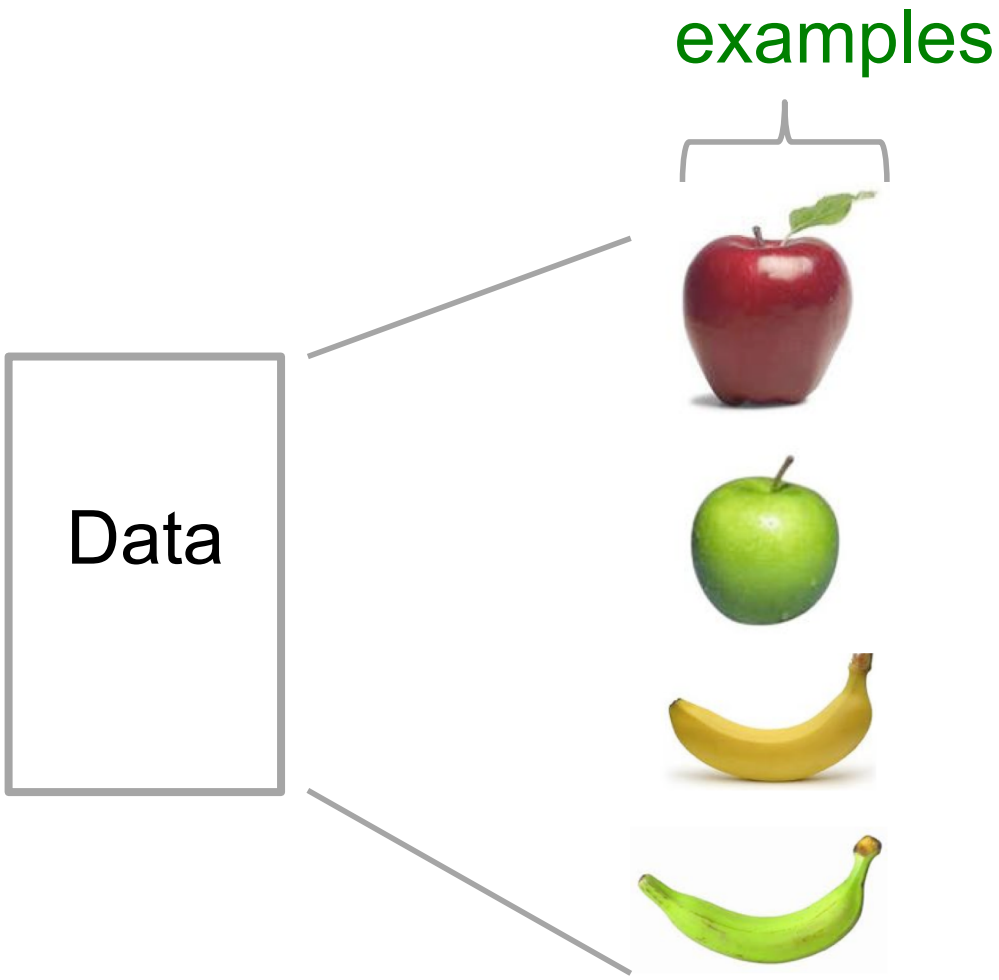

Data

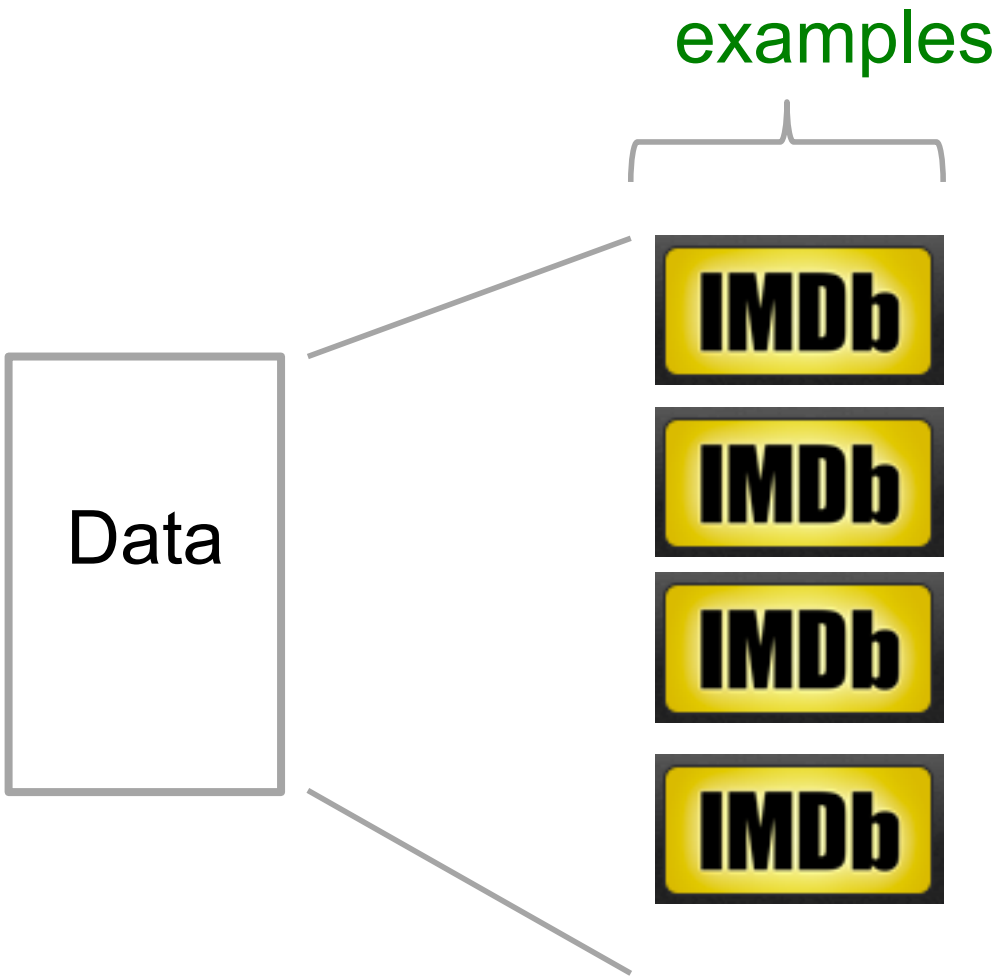

Data

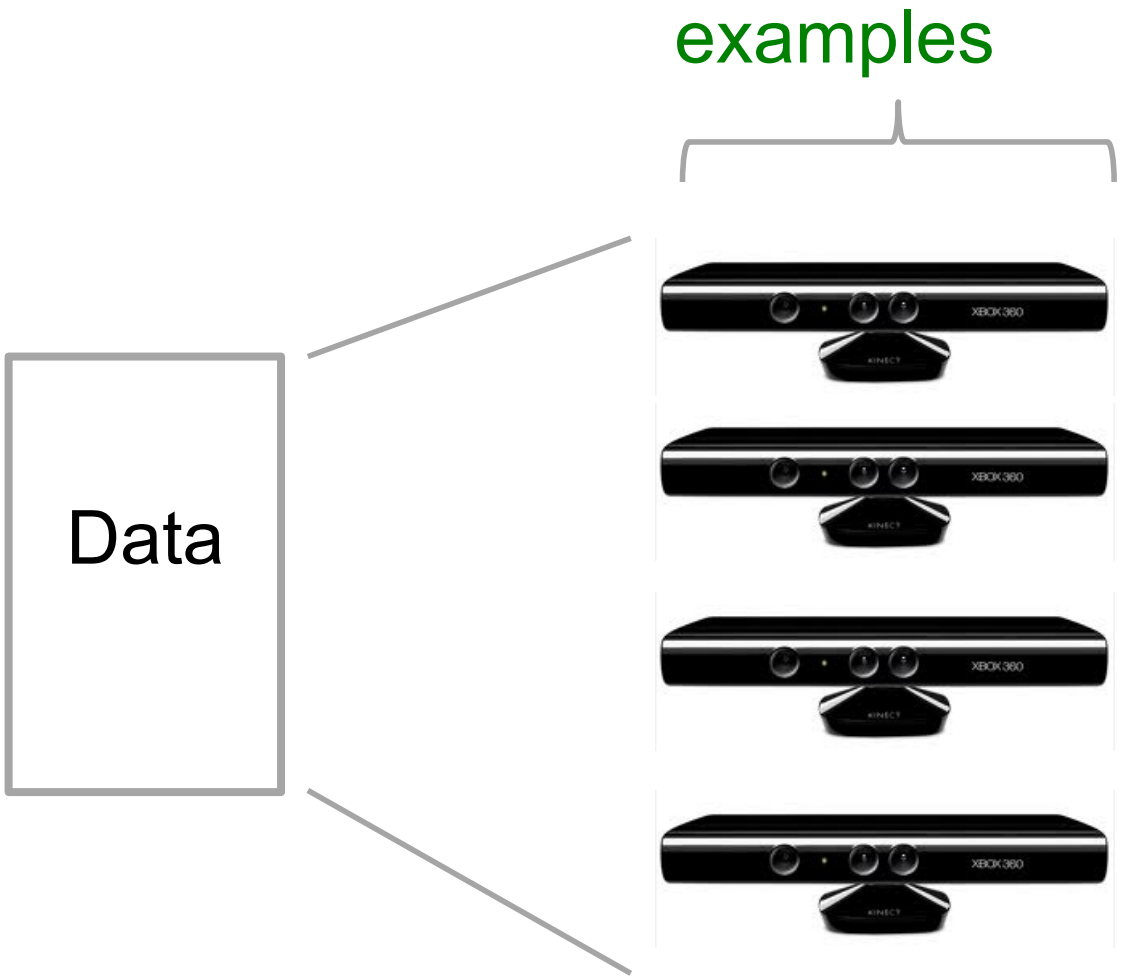

Supervised learning

examples

label

$label_1$

$label_3$

$label_4$

$label_5$

labeled examples

Supervised learning

label

$label_1$

$label_3$

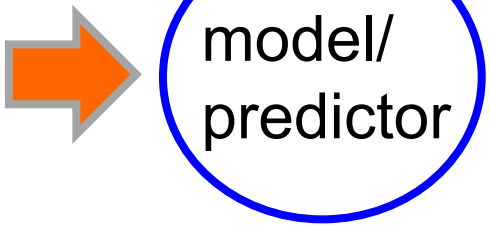

$label_4$

$label_5$

Supervised learning

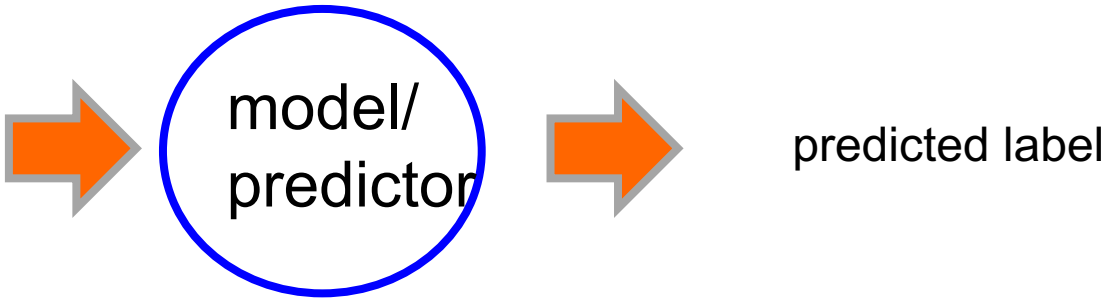

Supervised learning: classification

label

apple

apple

banana

banana

Classification: a finite set of labels

Classification Example

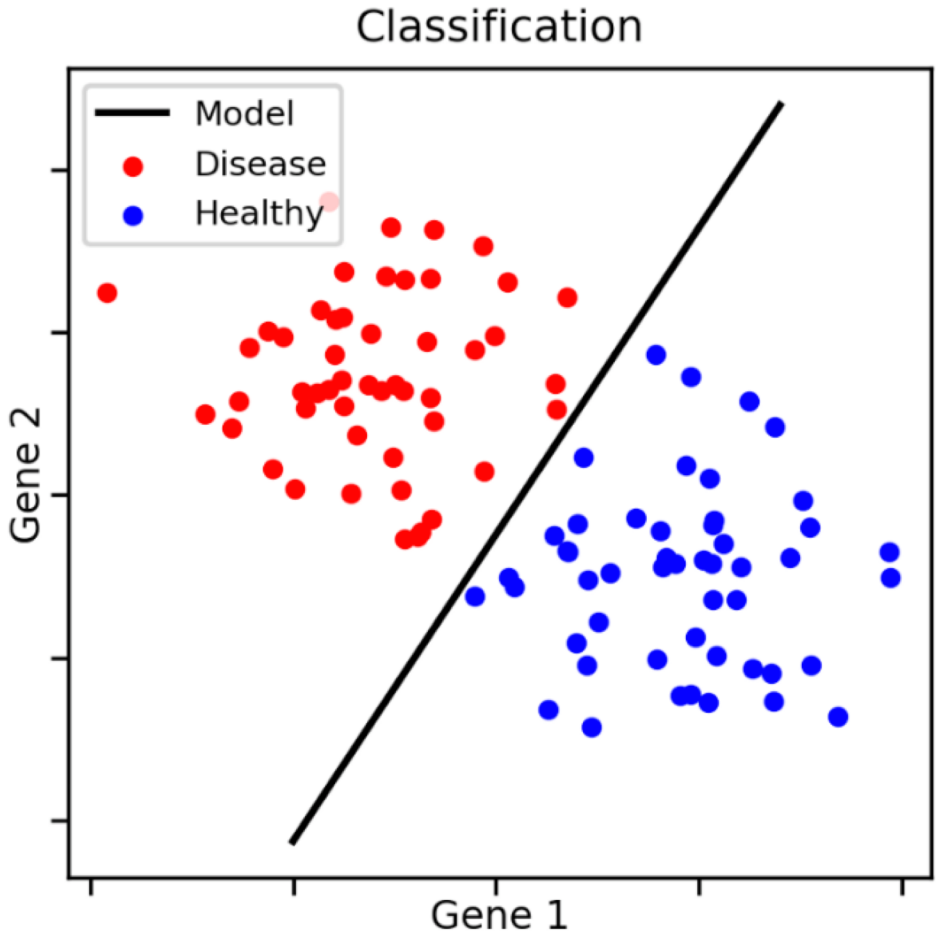

Classification Applications

Face recognition

Character recognition

Spam detection

Medical diagnosis: From symptoms to illnesses

Biometrics: Recognition/authentication using physical and/or behavioral characteristics: Face, iris, signature, etc

...

Supervised learning: regression

	label
	-4.5
	10.1
	3.2
	4.3

Regression: label is real-valued

Regression Example

Price of a used car

x : car attributes
(e.g. mileage)

y : price

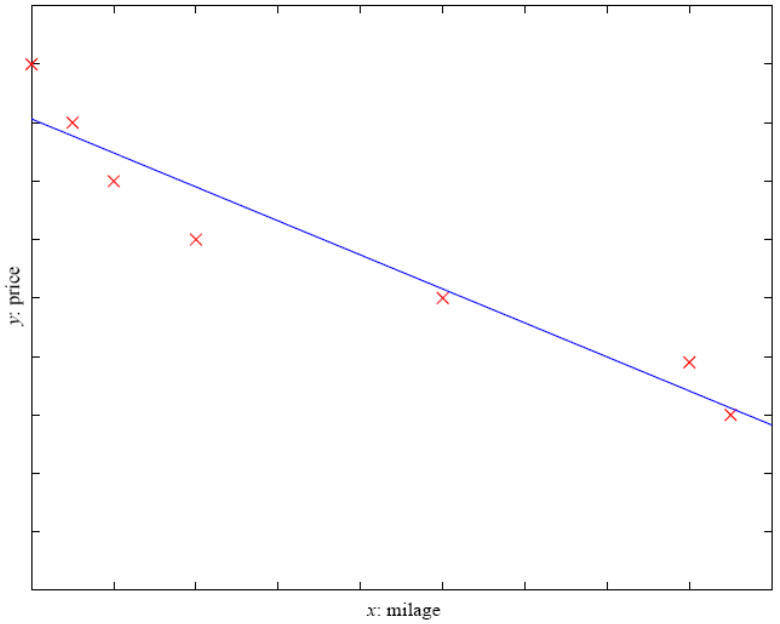

Regression Applications

Economics/Finance: predict the value of a stock

Epidemiology

Car/plane navigation: angle of the steering wheel, acceleration, …

Temporal trends: weather over time

…

Supervised learning: ranking

label

1

4

2

3

Ranking: label is a ranking

Ranking example

Given a query and
a set of web pages,
rank them according
to relevance

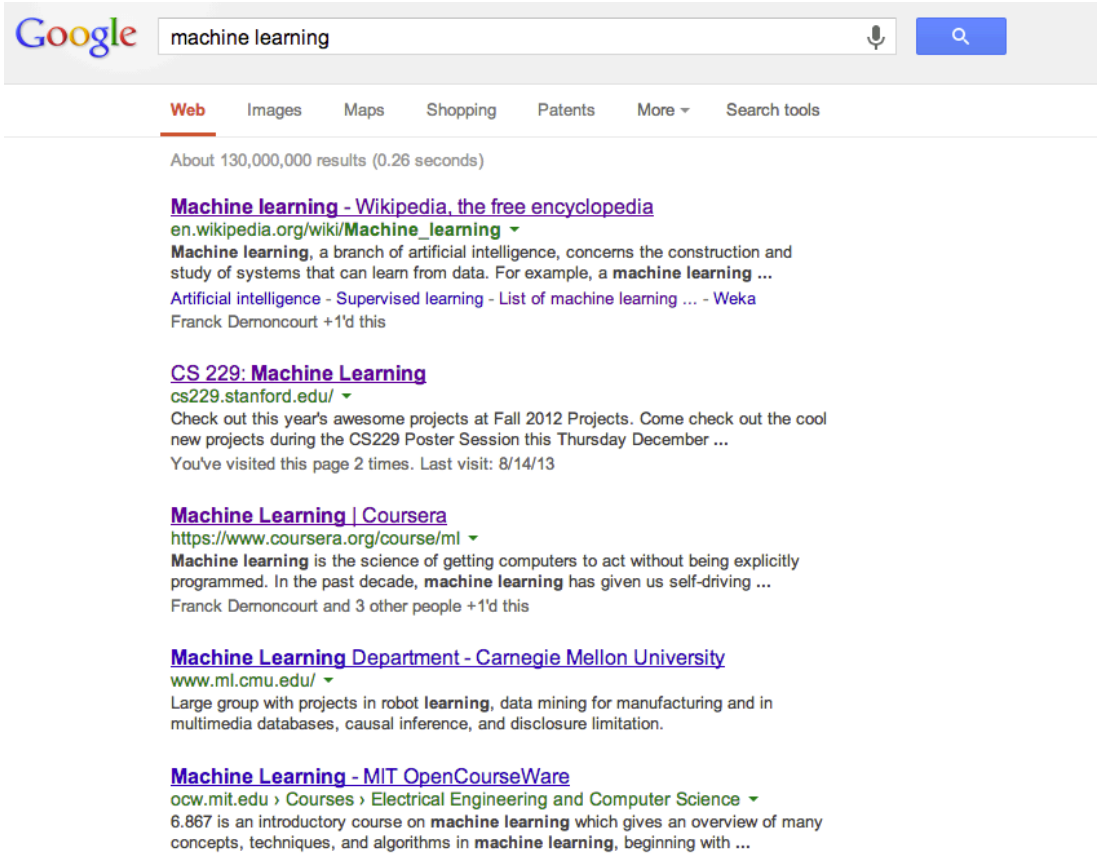

Ranking Applications

User preference, e.g. Netflix “My List” -- movie queue ranking

iTunes

search

…

Unsupervised learning





Unsupervised learning: given data, i.e. examples, but no labels

Learn clusters/groups

customer segmentation (i.e. grouping)

image compression

bioinformatics: learn motifs

…

Reinforcement learning

left, right, straight, left, left, left, straight	**GOOD**
left, straight, straight, left, right, straight, straight	**BAD**

left, right, straight, left, left, left, straight	**18.5**
left, straight, straight, left, right, straight, straight	**-3**

Given a ***sequence*** of examples/states and a ***reward*** after completing that sequence, learn to predict the action to take in for an individual example/state

Reinforcement learning example

Backgammon

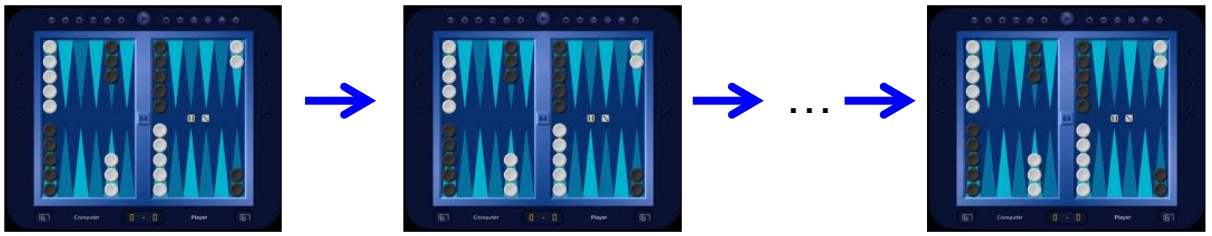

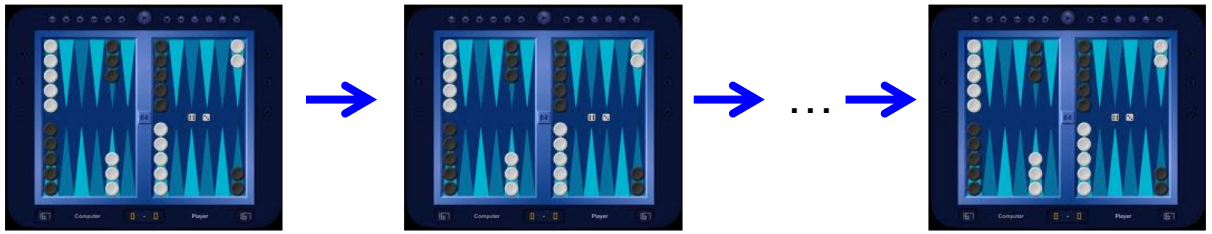

Given sequences of moves and whether or not the player won at the end, learn to make good moves

Other learning variations

What data is available:

- Supervised, unsupervised, reinforcement learning
- semi-supervised, active learning, …

How are we getting the data:

- online vs. offline learning

Type of model:

- generative vs. discriminative
- parametric vs. non-parametric

Representing examples

examples

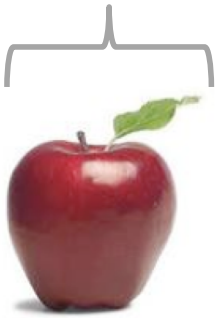

What is an example?
How is it represented?

Features

examples

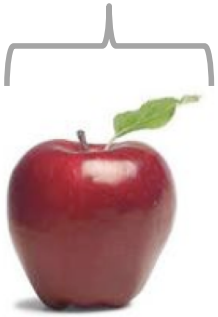

features

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

$f_1, f_2, f_3, \ldots, f_n$

How our algorithms actually "view" the data

Features are the questions we can ask about the examples

Features

examples

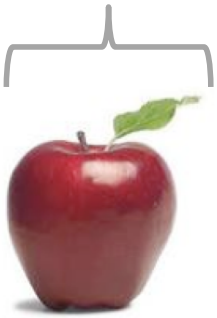

features

red, round, leaf, 3oz, …

green, round, no leaf, 4oz, …

yellow, curved, no leaf, 8oz, …

green, curved, no leaf, 7oz, …

How our algorithms actually "view" the data

Features are the questions we can ask about the examples

Classification revisited

examples	label
red, round, leaf, 3oz, …	apple
green, round, no leaf, 4oz, …	apple
yellow, curved, no leaf, 8oz, …	banana
green, curved, no leaf, 7oz, …	banana

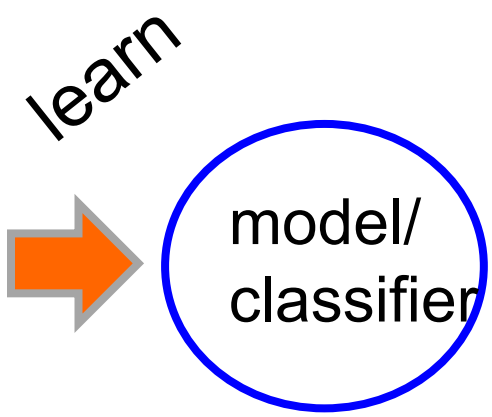

During learning/training/induction, learn a model of what distinguishes apples and bananas *based on the features*

Classification revisited

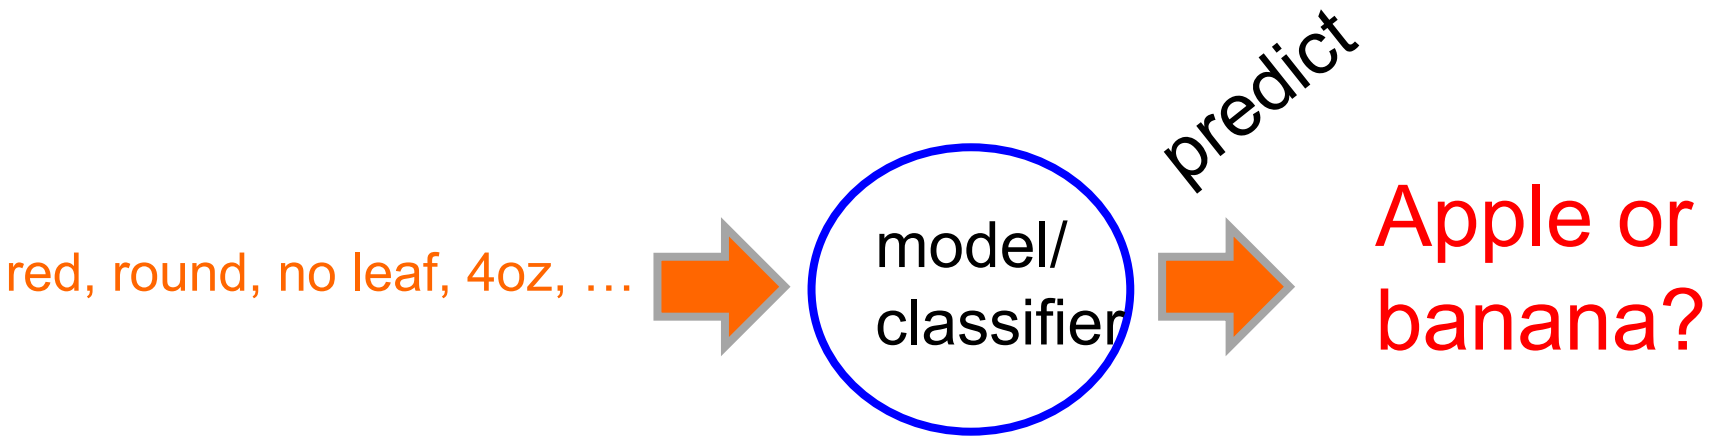

The model can then classify a new example *based on the features*

Classification revisited

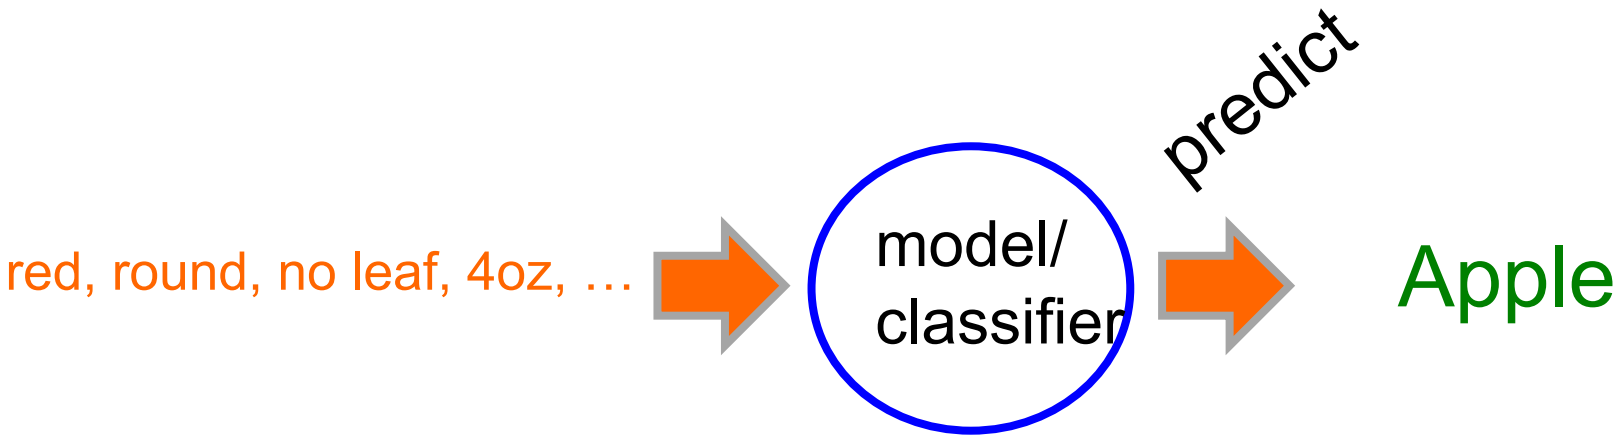

Why?

The model can then classify a new example *based on the features*

Classification revisited

Training data		Test set
examples	label	
red, round, leaf, 3oz, …	apple	
green, round, no leaf, 4oz, …	apple	red, round, no leaf, 4oz, … ?
yellow, curved, no leaf, 4oz, …	banana	
green, curved, no leaf, 5oz, …	banana	

Classification revisited

Training data		Test set
examples	label	
red, round, leaf, 3oz, …	apple	
green, round, no leaf, 4oz, …	apple	red, round, no leaf, 4oz, … ?
yellow, curved, no leaf, 4oz, …	banana	
green, curved, no leaf, 5oz, …	banana	

Learning is about ***generalizing*** from the training data

models

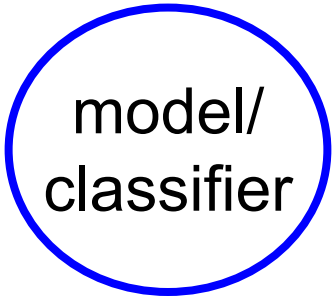

We have many, many different options for the model

They have different characteristics and perform differently (accuracy, speed, etc.)

Probabilistic modeling

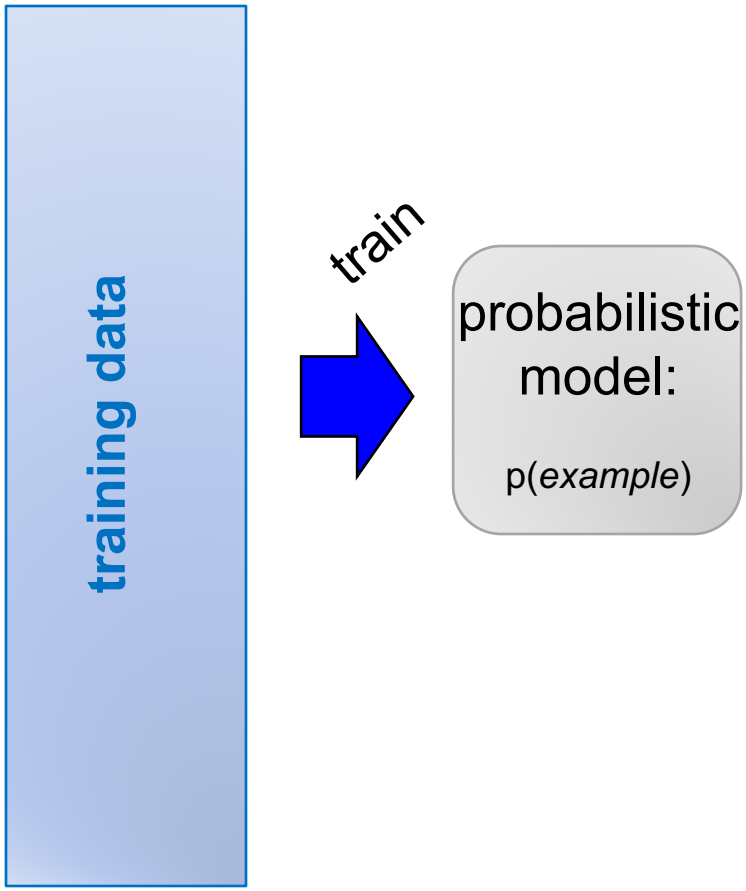

Model the data with a probabilistic model which tells us how likely a given data example is

Probabilistic models

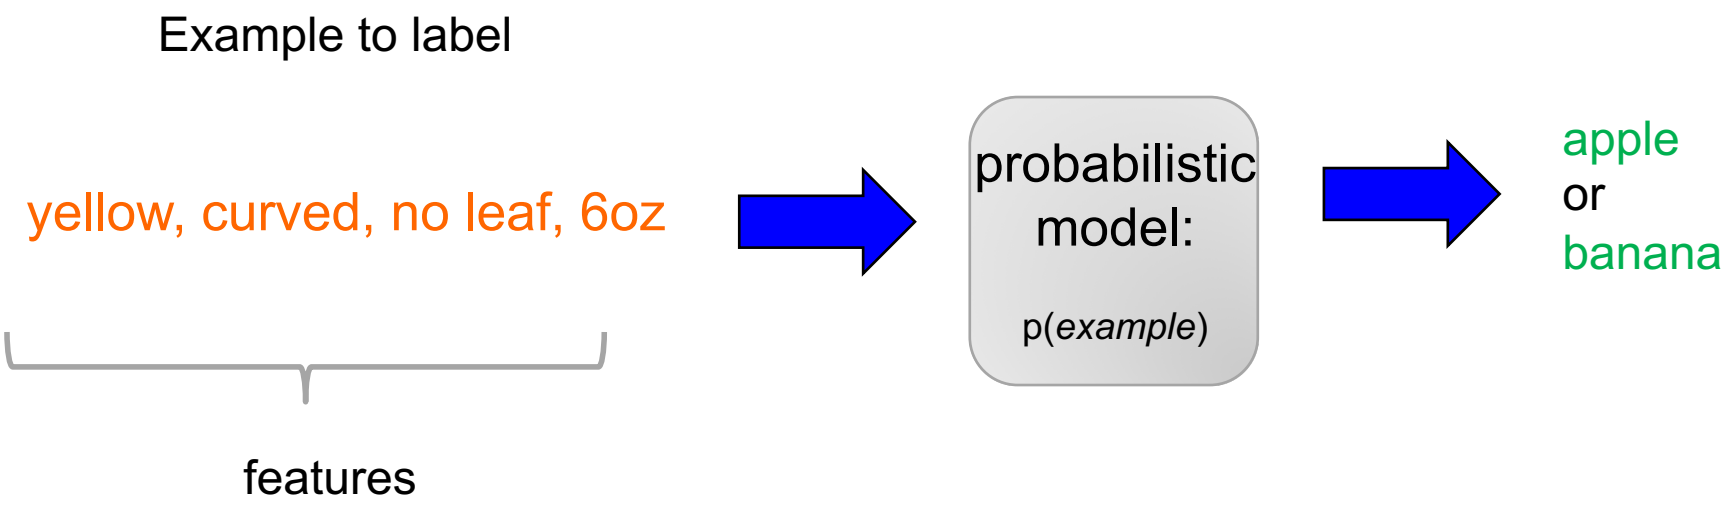

Probabilistic models

For each label, ask for the probability

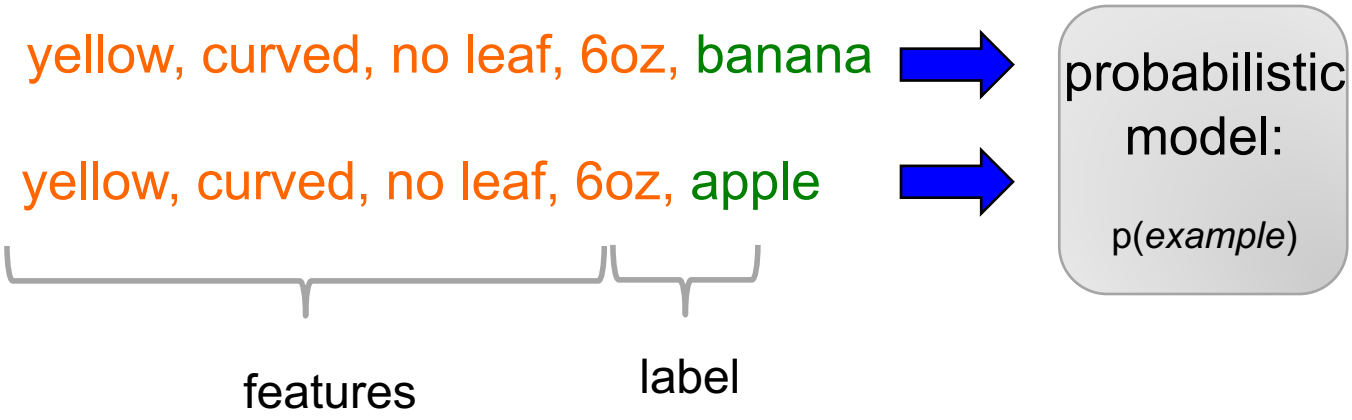

Probabilistic models

Pick the label with the highest probability

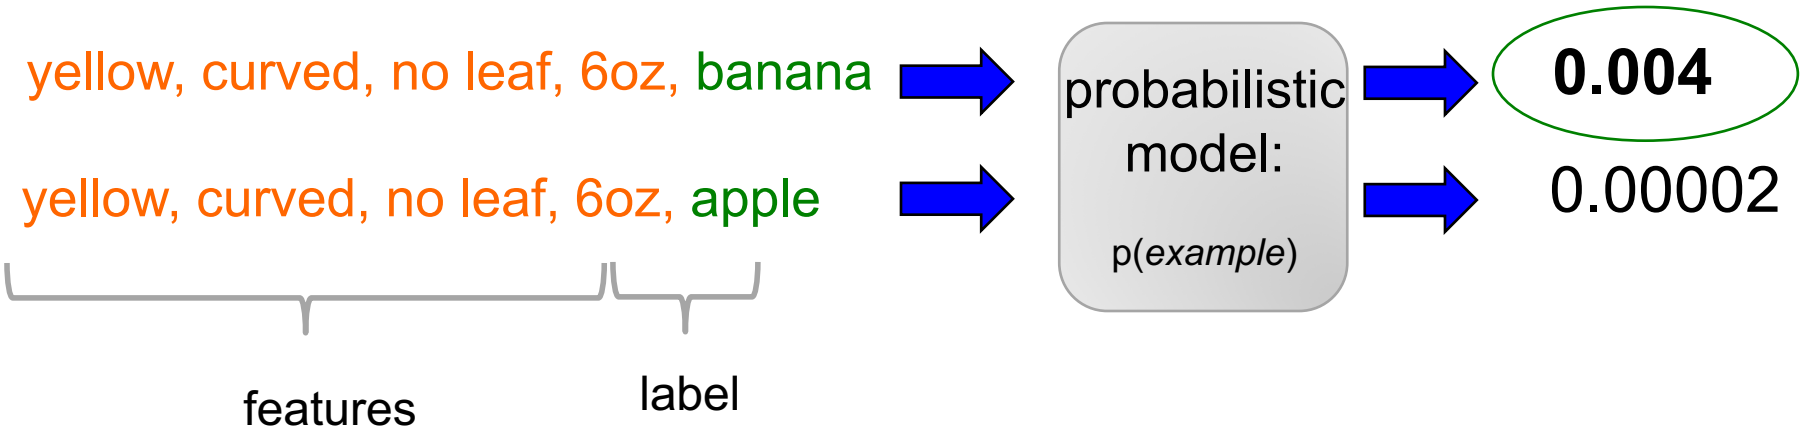

What can we currently do?

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

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

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?

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

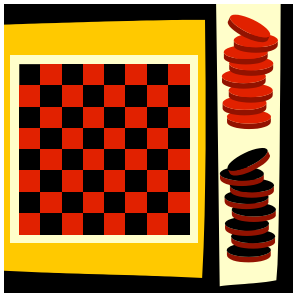

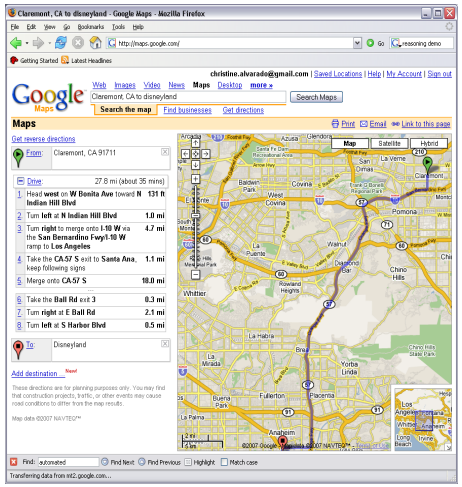

General purpose reasoning is harder

- 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_Wz4
- Sony QRIO
  - http://www.youtube.com/watch?v=9vwZ5FQEUFg
- Boston Dynamic's Atlas
  - https://www.youtube.com/watch?v=hSjKoEva5bg

Fundamental problem of AI

Search

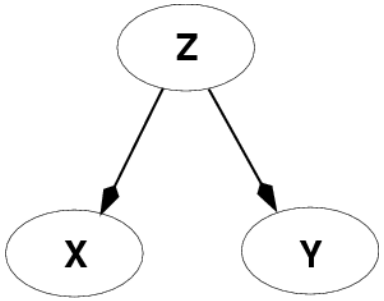

Reasoning with knowledge and uncertainty

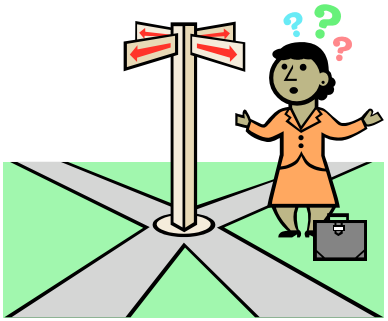

Reasoning with Utility

Many different ways of making an agent intelligent

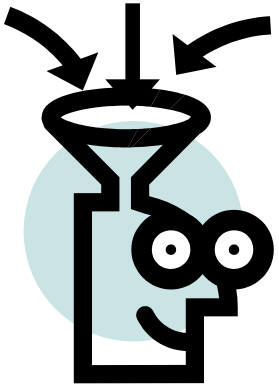

Learning

AI and Ethics

Discriminating algorithms: 5 times AI showed prejudice

Artificial intelligence is supposed to make life easier for us all – but it is also prone to amplify sexist and racist biases from the real world

TECHNOLOGY 12 April 2018 , updated 27 April 2018

By Daniel Cossins

We Teach A.I. Systems Everything, Including Our Biases

By Cade Metz

Nov. 11, 2019

How We Analyzed the COMPAS Recidivism Algorithm

by Jeff Larson, Surya Mattu, Lauren Kirchner and Julia Angwin

May 23, 2016

Amazon Created a Hiring Tool Using A.I. It Immediately Started Discriminating Against Women.

By JORDAN WEISSMANN

OCT 10. 2018 • 4:52 PM

Face-recognition software is perfect – if you're a white man

TECHNOLOGY 13 February 2018

By Timothy Revell