Robot brains: AI and Machine Learning

Mike Taylor
CMU 15-110
April 13, 2018

Baymax: Pixar

The Sense-Think-Act Paradigm

- Analogous to functions!
- A **function** takes **inputs** and yields **outputs**
- A robot **thinks** about **sensory** inputs and **acts** on its environment in some way
- The most important principle of AI and Robotics is…?

2

You are now a robot in 15-110

- Why not? Humans study biology. Beep boop.
- You **hate** pop quiz days. Do you think there will be one today?
- Are you 100% sure about that?
- …are you *really* 100% sure? [O _ O'']

3

Part 1: It's all about

PROBABILITY

disney.wikia.com

4

A simple approach first!

- Assume there will only ever be one quiz per day!

- Let's say you've had **16 quizzes** so far…
- …and there have been **50 classes**…
- …and because you are a good robot, you have come to all of them.

- The probability of a quiz today is 32%, or 16/50!

5

Maybe we can find a hint…

- Prof. Cortina likes to have a bottle of water during lectures, but…
  - …sometimes he doesn't bring one if there is a quiz
  - …and sometimes he just forgets

- He doesn't have any water today! Does that mean there's probably a quiz?

6

Important notation

- H – Hypothesis
- E – Evidence
- P(H) – Probability of H
  - Also called the **prior**
- P(H|E) – Probability of H given observation E
  - Also called the **posterior**

We want to know P(H|E)

- H = 'having a quiz'
  P(H) = 0.32
  - because we've had 16 quizzes in 50 classes
- E = 'doesn't have a water bottle'
  P(E) = 0.2
  - You've noticed the professor is four times more likely to have water with him than not
- E | H = 'doesn't have a water bottle given a quiz occurs'
  P(E|H) = 0.5
  - In the last 16 quizzes, the professor has only had water 8 times

Bayesian Inference!

- P(H|E) = P(E | H) * P(H) / P(E)

Bayes' Theorem tells us the probability of an event H conditional on observation E.

We just need to know P(H), P(E), and P(E|H)

9

Bayesian Inference!

P(H) = 0.32

P(E) = 0.2

P(E|H) = 0.5

P(H|E) = 0.5*0.32/0.2
= 0.8 (80% chance)

Hope you studied!

10

Update as we get new data!

If no water and a quiz happened	**If no water and no quiz**
P(H) increases	P(H) decreases
P(E) increases	P(E) increases
P(E|H) increases	P(E|H) stays the same

11

And now we're smarter robots!

We observed some data, made a prediction, and incorporated the result into our model for next time

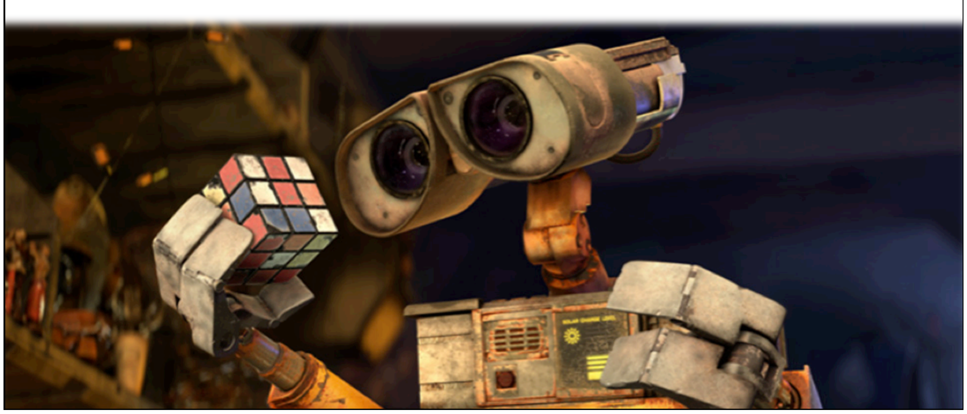

More info…

- There are some flaws…
  - You can't predict an event you haven't seen
  - Small datasets and rare events can yield poor performance
- Training sets can help
- Typical Bayesian applications are similar but more complex
  - Observations and outcomes may not be binary
  - Multiple sources of evidence, i.e. sensors

13

Other examples in the real world

- Predictive text, i.e. on phones
- Simple games, i.e. rock, paper, scissors
- Spam filters and text classification

14

Part 2: How do robots "see?"

COMPUTER VISION

15

Jibo image: PCMag

Which pixels are "Ball?"

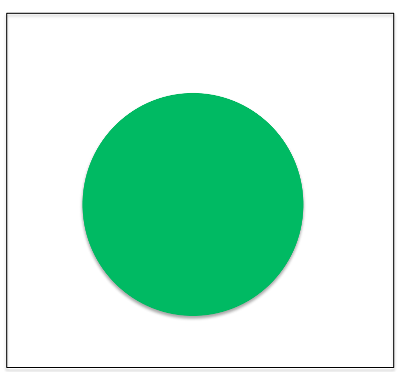

16

Which pixels are "Ball?"

```
000000000000000000000000000000
000000000000000000000000000000
000000000000111111000000000000
000000001111111111111000000000
000000111111111111111100000000
000001111111111111111111000000
000001111111111111111111000000
000001111111111111111111000000
000000111111111111111100000000
000000001111111111111000000000
000000000000111110000000000000
000000000000000000000000000000
```

```
000000000000000000000000000000
000000000000000000000000000000
000000000000111111000000000000
000000001111111111111000000000
000000111111111111111100000000
000001111111111111111111000000
000001111111111111111111000000
000001111111111111111111000000
000000111111111111111100000000
000000001111111111111000000000
000000000000111110000000000000
000000000000000000000000000000
```

Remember, images can be lists of lists!
Or with color… lists of lists of lists ad nauseum…

Which pixels are "Ball?"

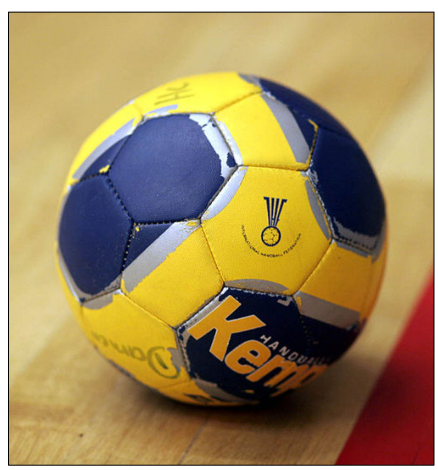

https://en.wikipedia.org/wiki/Ball

Where's Wally?!

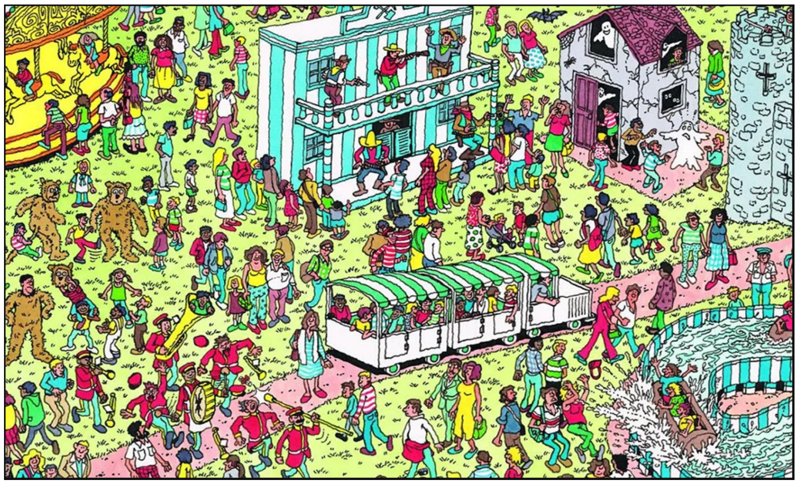

https://www.theguardian.com/technology/2015/feb/06/wheres-wally-wheres-waldo-machine-learning-algorithm

More information than just color! Gradients, for example:

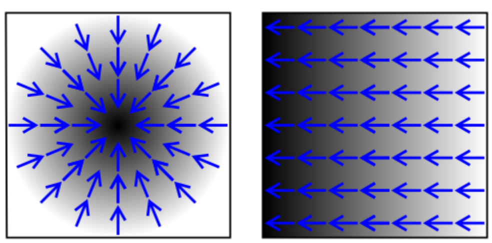

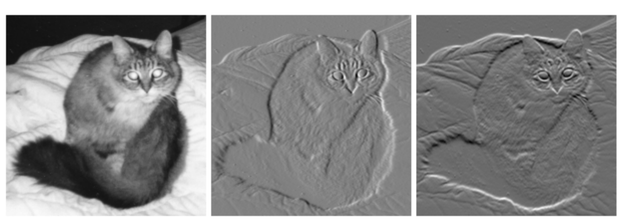

https://en.wikipedia.org/wiki/Image_gradient

Edge detection

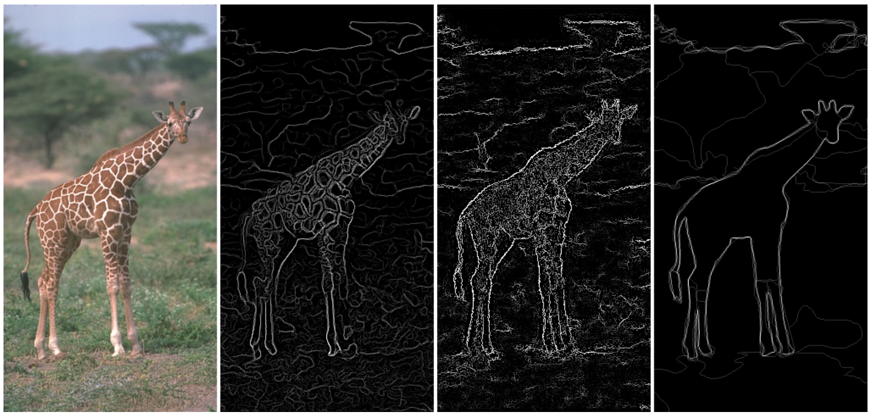

http://cs.brown.edu/courses/cs143/2013/proj5/ 21

Computer vision tasks use probability!

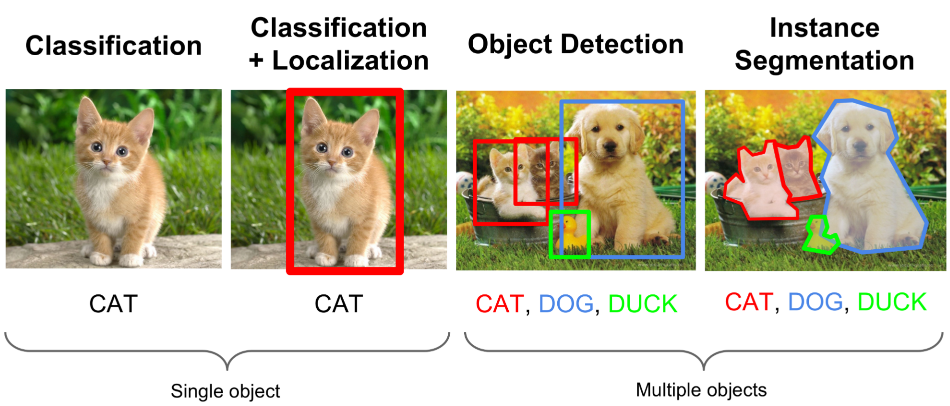

From http://cs224d.stanford.edu/index.html 22

Part 3: Where am I? Where am I going?

ROBOT MOTION

Cosmo image: Apple.com

A robot in an empty parking lot

- I want my robot to drive in a square
  - Drive forward y seconds
  - Turn left x seconds
  - Repeat three times*
- Will I actually get a square?

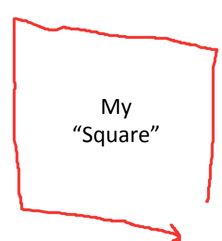

- Try drawing a big square with your eyes closed!

Uncertainty in motion

- Motion model: Given a command to move a certain way, where *might* the robot end up?

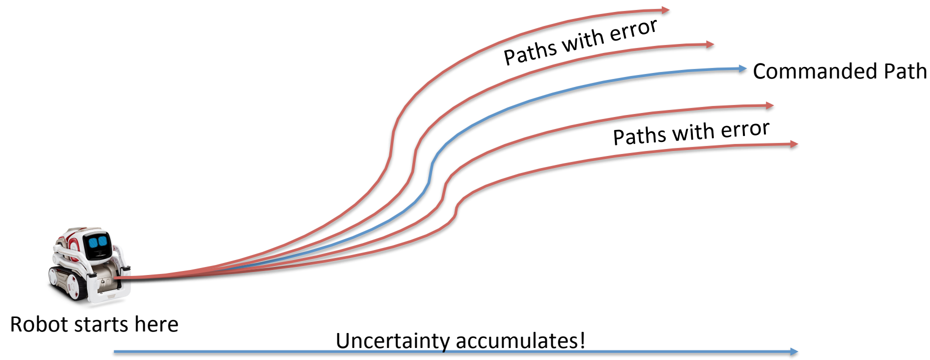

25

Uncertainty in sensing

- Sensor model: Given a reading from one of our sensors, how likely is it to be accurate?
- Mapping: Use sensor data to model the environment
- Localization: Given the sensor model and motion model and a **map**, where is the robot most likely to be?

26

Artificial Intelligence in a Nutshell

Sense – Think – Act

Machine learning is applied probability

Models are imperfect; always model uncertainty

27