Security
Machine Learning
Deep Learning

Announcement

- Trouble with HW 4? Please come to the TA office hours!
- How to increase your grade?
  - Bonus projects
  - Worst homework and worst quiz dropped
  - Class participation: 5 points
- Quiz 5 is going to be on Python string processing and Diffie-Hellman key exchange

Machine Security Continued

Passwords

- Passwords allow access to machines or system settings
- Passwords are encrypted and stored in a computer file
- When you log in, your typed-in password is encrypted, compared against encrypted stored password
- Nobody knows your plaintext password

Set password: ChangeMe

Encrypted: Edr4^7dW

Login: ChangeMe

OS compares Edr4^7dW

to the stored Edr4^7dW

Only you know your plaintext

Password Cracking

Sometimes, people gain access to encrypted password file

- Can spend days/weeks trying out passwords to see if they match
  - Modern systems can try 3.5 billion passwords/sec
- Use dictionaries of common passwords to speed search
- How many possibilities if random?
  - 4 letters (26 x 26 x 26 x 26) = 456, 976
  - 8 letters = 200 billion
- Adding uppercase and numbers increases possibilities
  - 4 characters (80^4) = 40,960,000
  - 8 (80^8) = 1.6 quadrillion

Social Engineering

Gaining access to secure areas/passwords through social means

- Physical access is often enough to gain admin privileges to machines
- Pretexting: learn enough about someone to gain initial access, learn more, repeat
- Baiting: leave malicious software in a location where it might be picked up and installed
- Quid Pro Quo: randomly call offering tech support. Have users type commands that install malware.

Phishing

Attempt to gain sensitive info

- Pretend to be from trusted site: use logos, obscure web page locations
- Use panic to push people to act: account is compromised/suspended, etc.
- Works because:
  - Will match some small percentage of people who use the site/service
  - People with ongoing problem so the need appears legitimate

Web Server Hiding

Many ways of hiding malicious web servers:

- Incorrect link in a message
- Use close approximations

- http://privatebanking.mybank.com.**ch**
- http://mybank.**privatebanking**.com
- http://privatebanking.**mybonk**.com or even http://privatebanking.**mybánk.**com
- http://privatebanking.mybank.**hackproof.com**

Malware: malicious software

Mal – Latin for bad or evil

Malware is any program that is designed to harm a computer

2004: average time for a new computer to get infected was 4 minutes

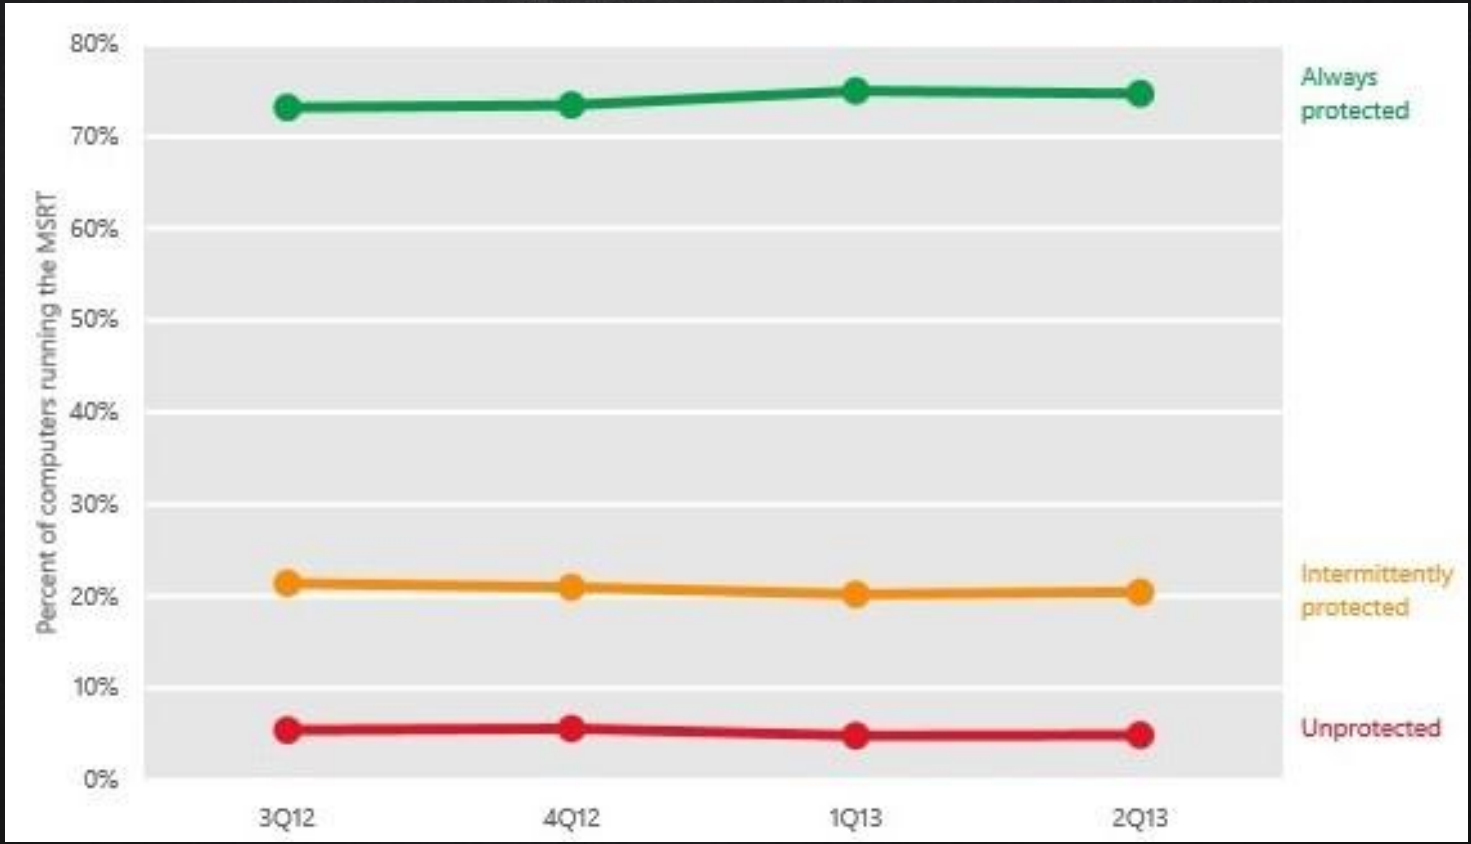

https://blogs.technet.microsoft.com/mmpc/2013/10/29/new-infection-rate-data-for-unprotected-computers/

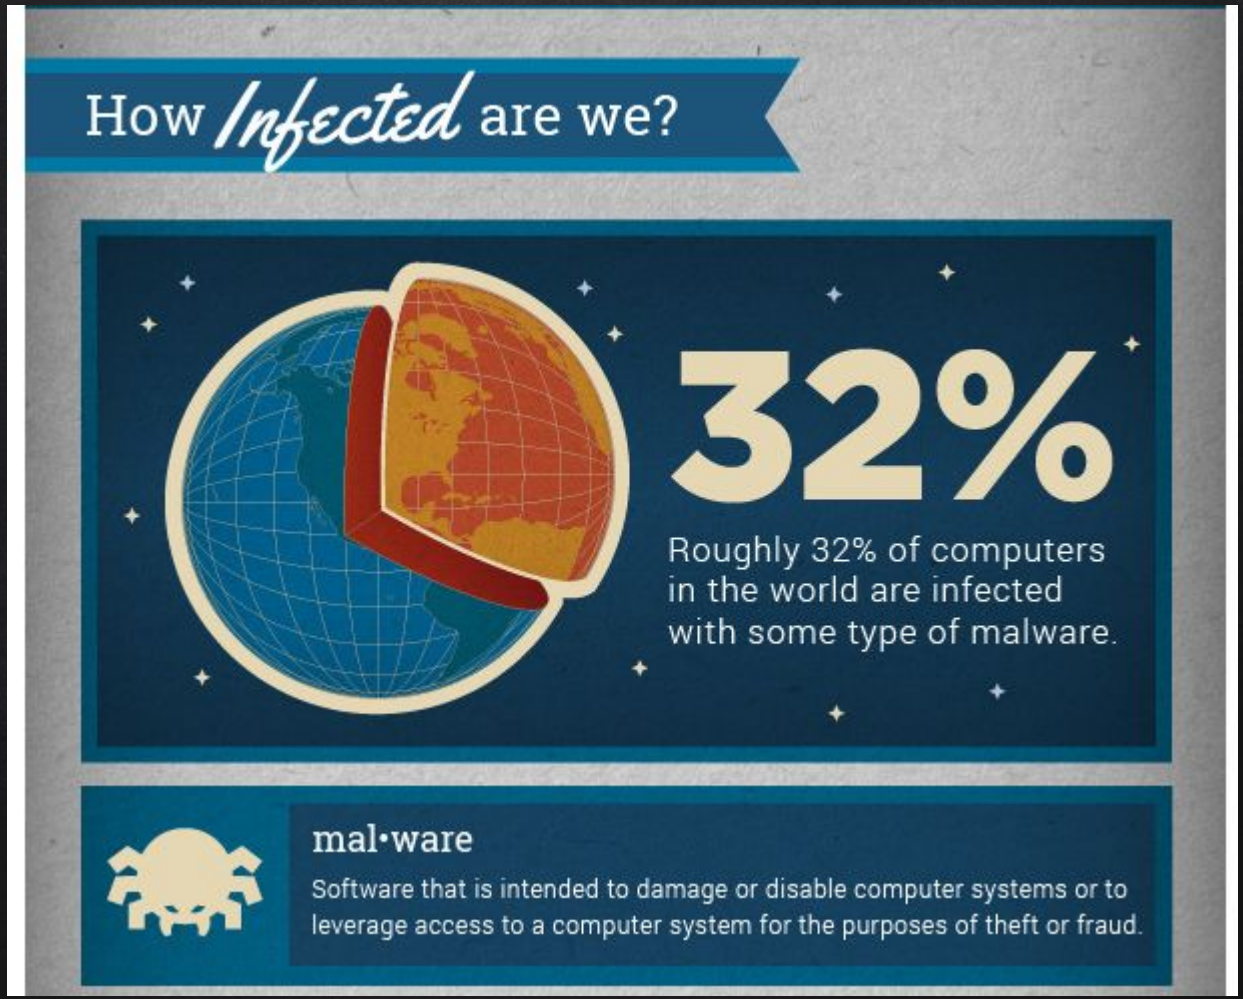

http://anti-virus-software-review.toptenreviews.com/how-infected-are-we.html

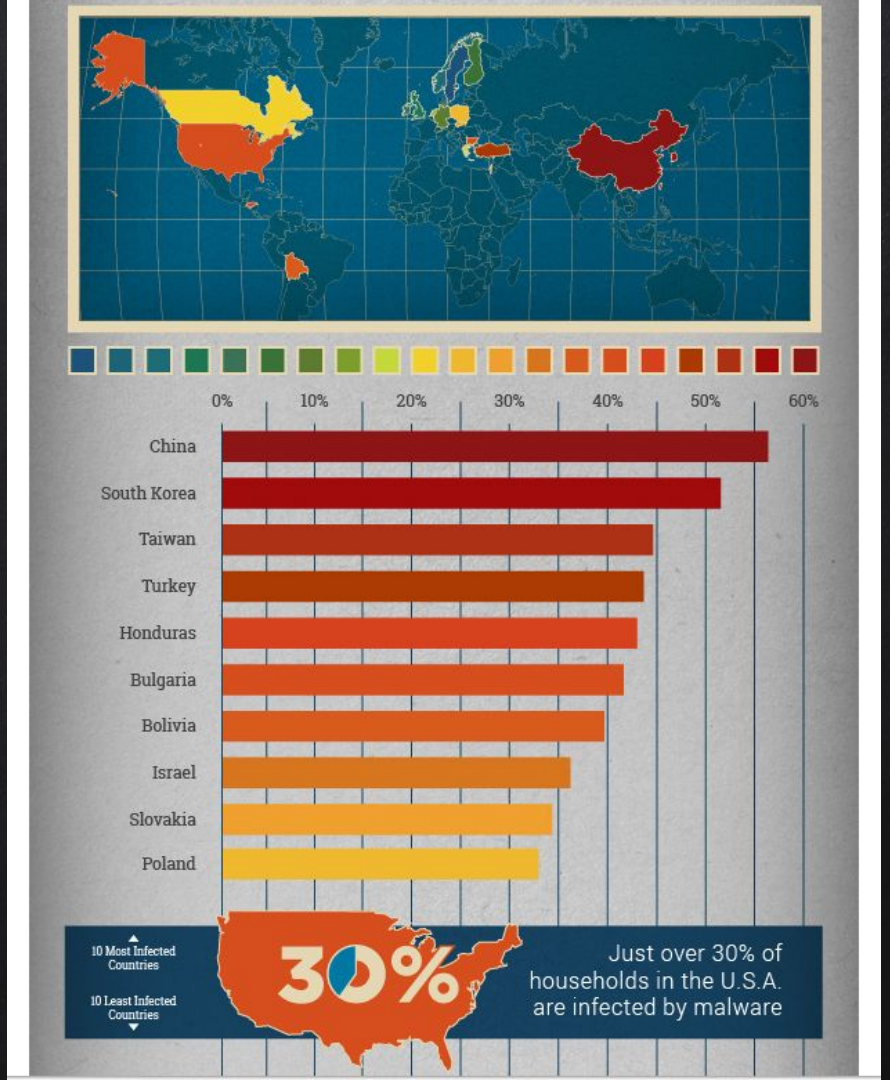

http://anti-virus-software-review.toptenreviews.com/how-infected-are-we.html

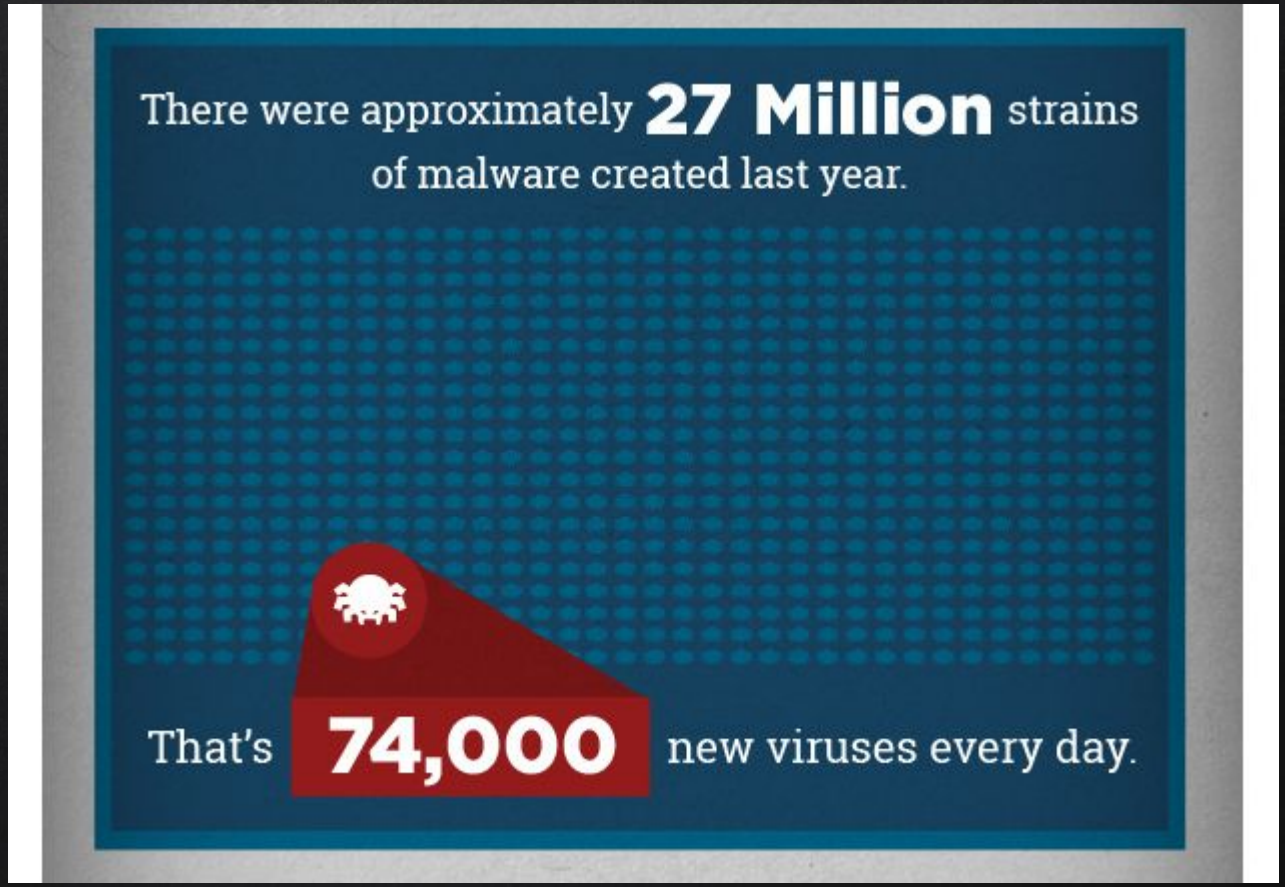

http://anti-virus-software-review.toptenreviews.com/how-infected-are-we.html

The most *Prolific* virus of all time

Conficker (also known as Downup, Downadup and Kido) is a computer worm that targets flaws in the Windows operating system to spread across system networks while forming its own network of auto-acting malware. It is known to be unusually difficult to counter. The Conficker infected millions of computers across 200 countries including everything from home personal computers to business and government networks. It is the largest known computer worm infection on record.

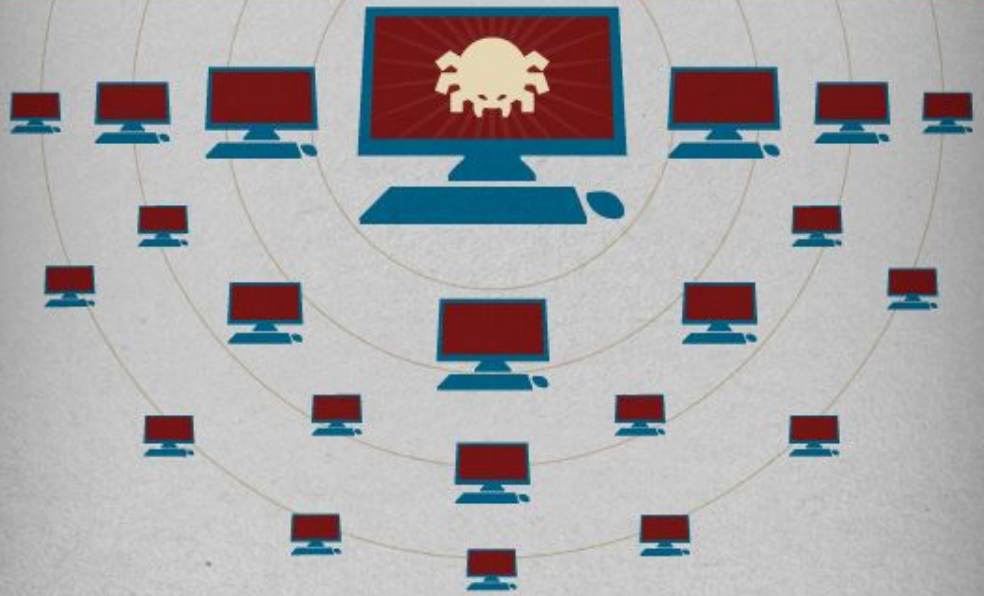

http://anti-virus-software-review.toptenreviews.com/how-infected-are-we.html

Computer Virus

Sometimes used as a term for any malware.

A virus is a computer program that can copy itself to infect another machine:

- Often copied along with a host file
- Modern viruses often use macro languages in Excel and Word

Trojan Horse

- A trojan horse is a software program that appears useful but contains malware: some look like anti-virus programs
- Does not replicate itself: depends on users

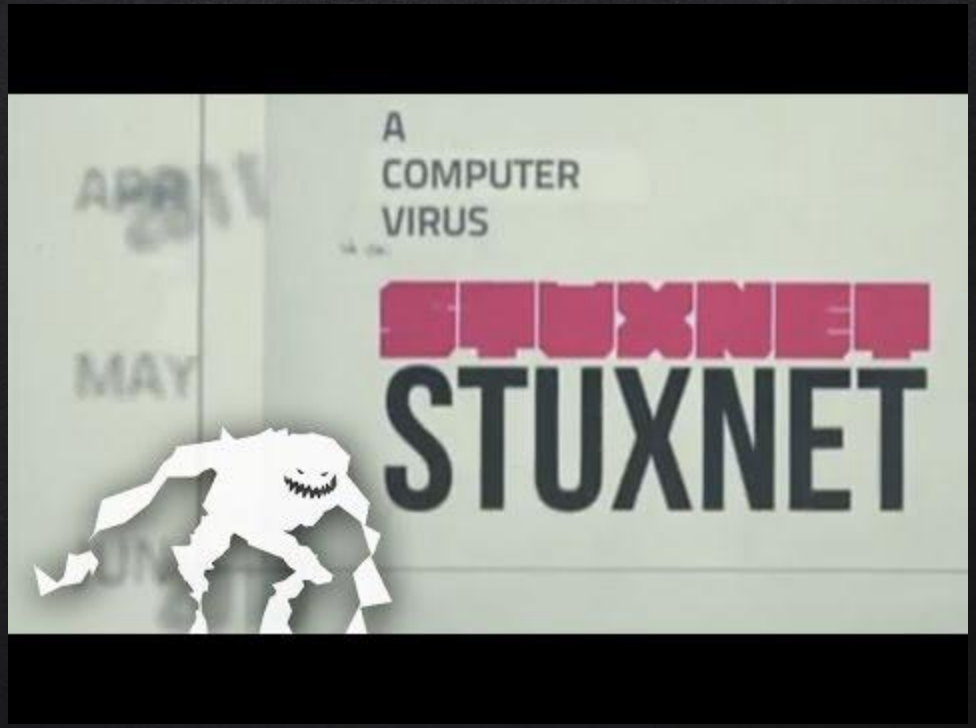

https://www.youtube.com/watch?v=7g0pi4J8auQ

Worm

- A worm is a type of virus capable of replicating w/o human help or host file
- Might carry more dangerous programs
  - Crypto extortion: encrypt your hard drive and extort money to decrypt it
- Stuxnet
  - Computer worm found in 2010
  - Attacks industrial equipment
    - Some believe it was written to target Iranian nuclear enrichment capabilities

Rootkits and backdoors

- Backdoors compromise computer to allow further access
  - Bypasses normal authentication
  - Some large-systems have backdoors installed by original programmers
  - Speculations that compilers could install backdoors by recognizing code
- Rootkits bypass normal login and also hide malicious activities
  - Sony music installed rootkits on computers in an effort to thwart piracy (2005)

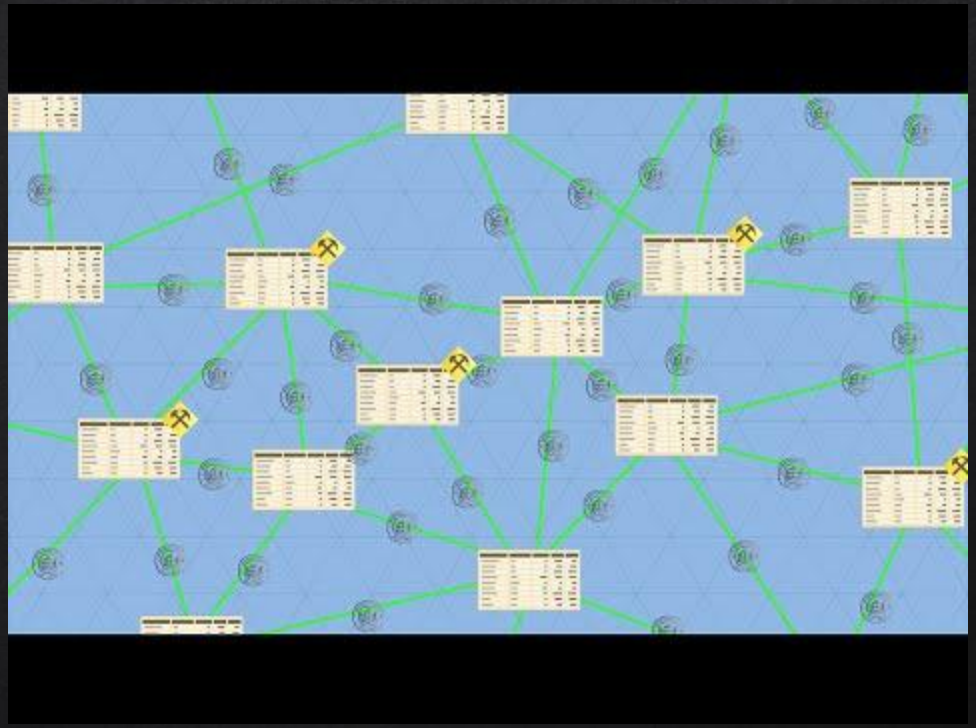

https://www.youtube.com/watch?v=YIVAIuSL9SU&feature=youtu.be

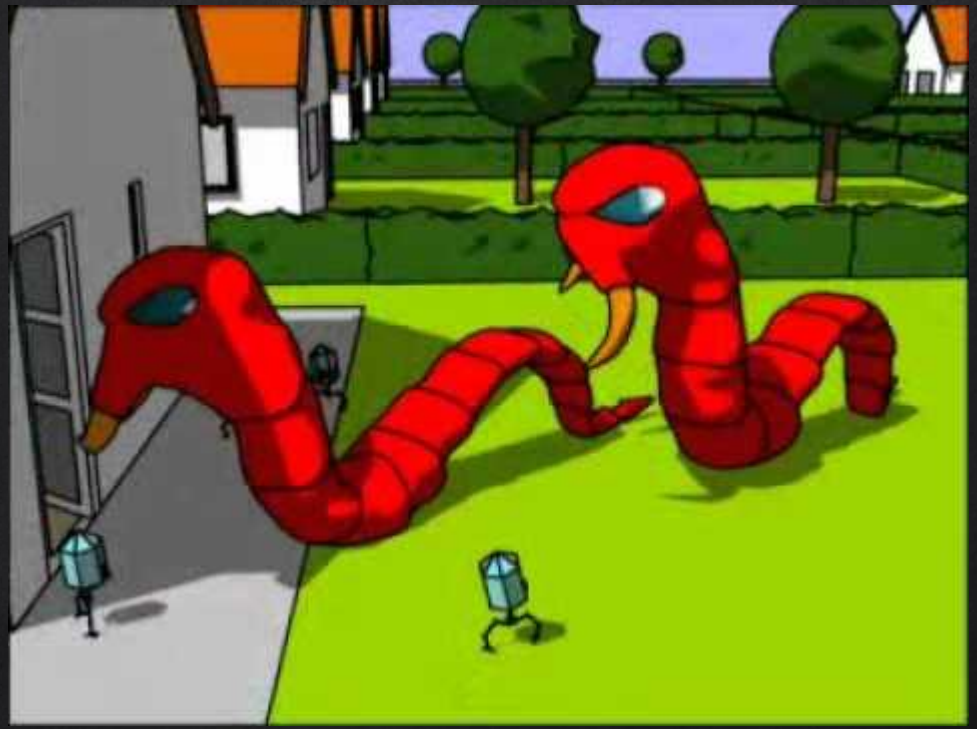

https://www.youtube.com/watch?v=c34QwtYI40g&ebc=ANyPxKrQcDOYHVwXUsY8vxWe26XxmUcREmLRlWPiU24nVe6341B7q_jVcxTopAjRrMbW_rsDCf7vDp79sxJLR5yVvuVvIP__nw

Protection

- Firewall
  - Prevents unauthorized communications
  - Keeps viruses from spreading
- Anti-virus
  - Searches files, messages from known viruses
- Web browsers
  - Maintain lists of phishing and malware sites
- Spam filters
  - Look for word patterns

Ongoing Battle

- Attacks and protection get more sophisticated
- Keeping a machine updated is important
  - Other programs have vulnerabilities as well: Acrobat, browsers
- Use strong passwords
- Don't use a personal machine password for some shopping site
- Backup data
  - Do a clean OS install

Machine Learning Explained

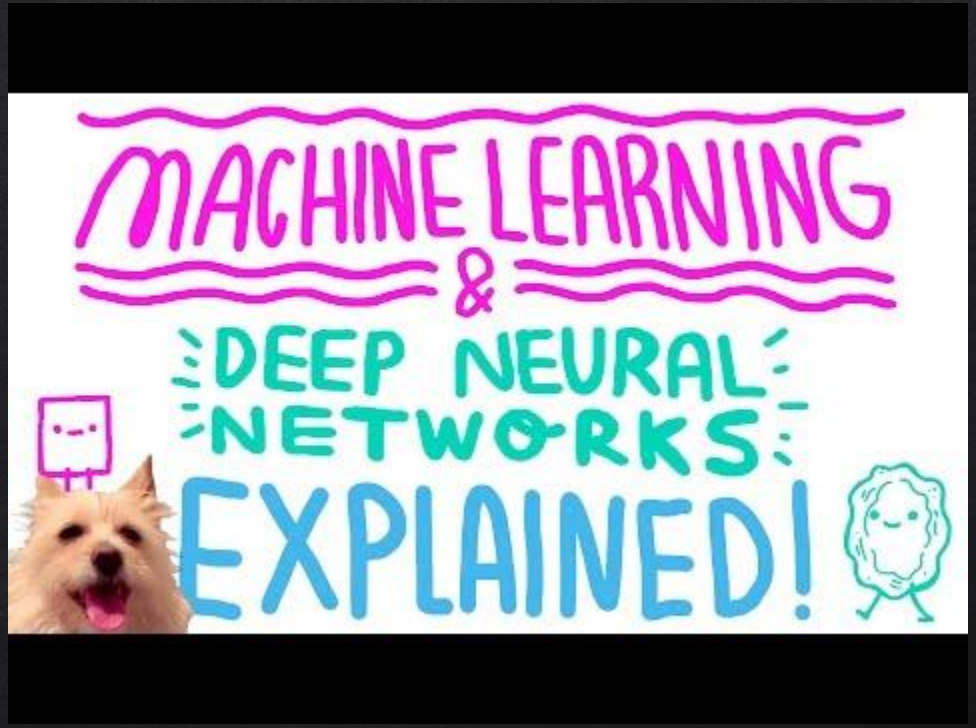

https://www.youtube.com/watch?v=bHvf7Tagt18

Machine Learning

Face recognition, etc....

A.I. Review

- Make computers do human skills
- Turning test: compare responses of human and computer
- Knowledge approaches
  - Semantic networks
  - Search trees
  - Expert Systems
- Language understanding
  - Problem of language ambiguity

Problems with A.I. Approach

- A.I. tried to classify everything:
  - strict categories
- Real world is messier

- It flies
  - It is a bird
  - Therefore it has feathers
- What about
  - Flying fish
  - Bats
  - Airplanes
  - Ostriches

Machine Learning

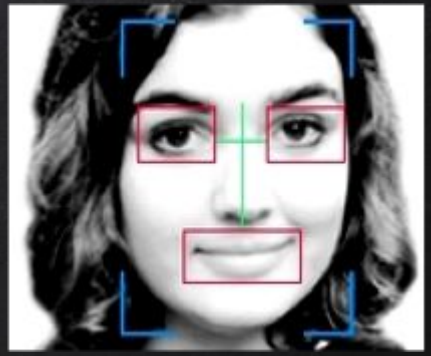

Machine learning uses probability and statistics

- Looks for patterns
  - Facial recognition
- Classification
- Learn based on empirical data
  - Humans learn from real-life experiences
  - Training
  - Generalization

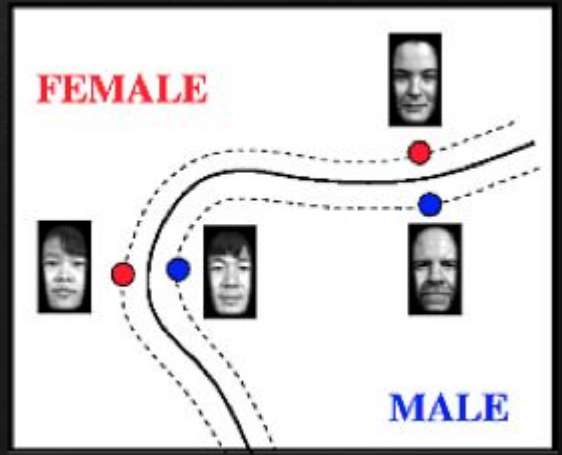

Classification

- Decide what category something falls into
  - Male / female
  - Healthy / diseased
  - Buy / hold / sell
- Train against data classified by an expert
- Ask questions about unknown objects
- Problem
  - Identify fish in packing plant
  - Sea bass v.s. Salmon

Goal

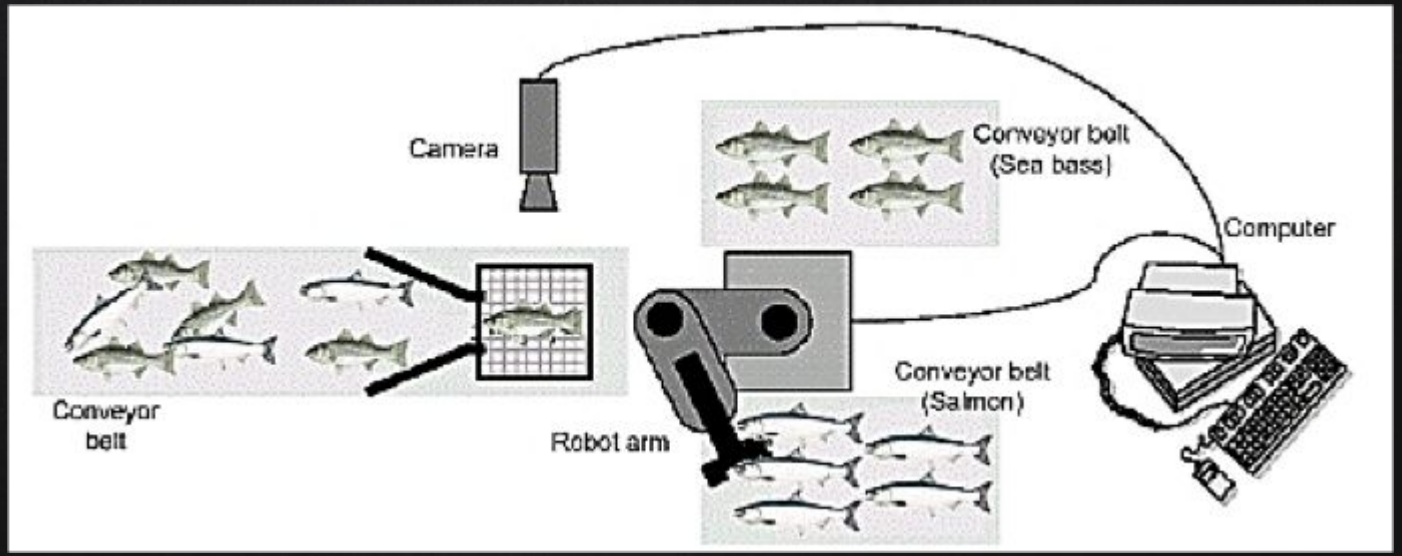
Camera
Conveyor belt
(Sea bass)
Computer
Conveyor belt
(Salmon)
Conveyor
belt
Robot arm

Procedure

- Human expert looks at data and does classification
  - Creates training set based on length and color

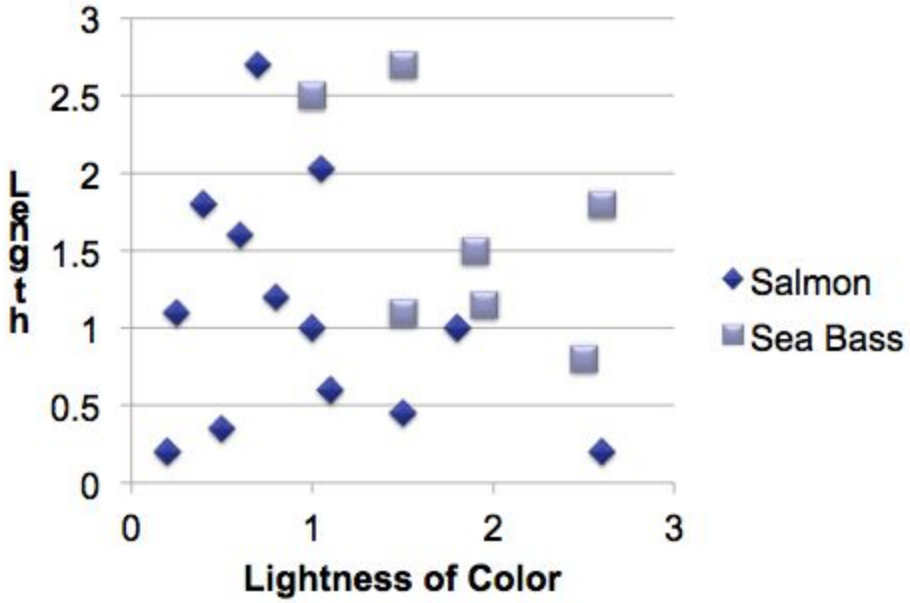

Classification of New Item

- Measure characteristics of new fish
  - Creates a new point on chart

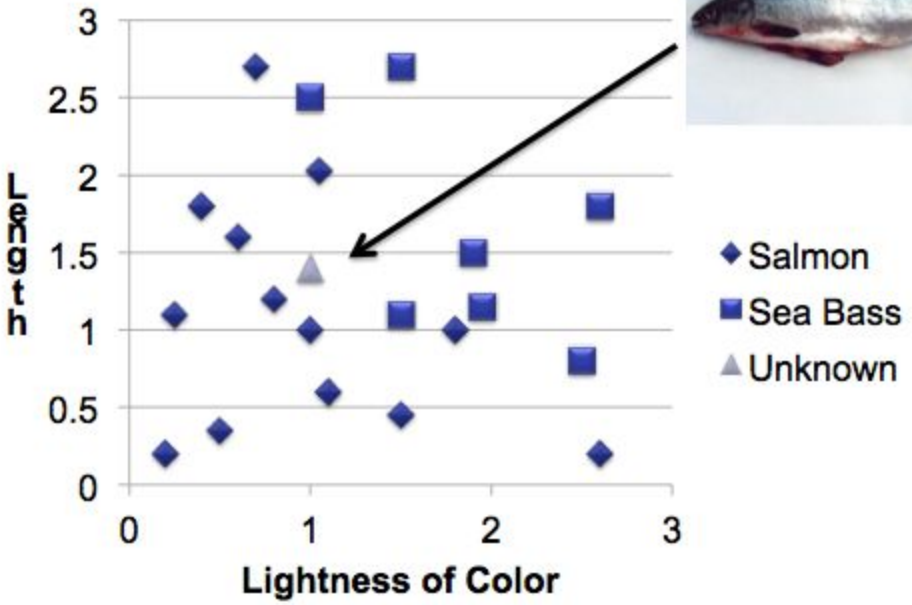

Classification of New Item

- Find closest known point to it
  - Assume it is the same

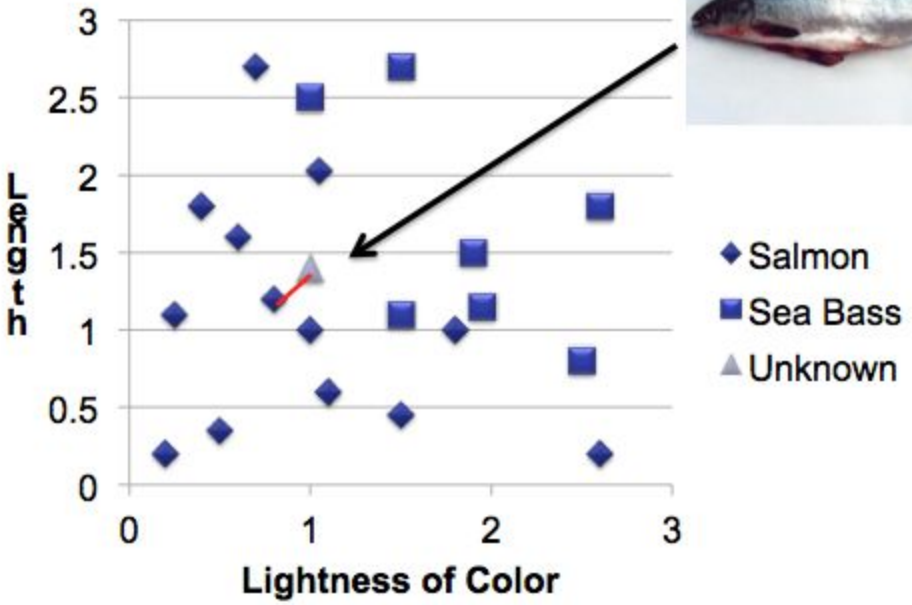

Classification of New Item

- Closest point can lead to errors
  - Classify as sea bass

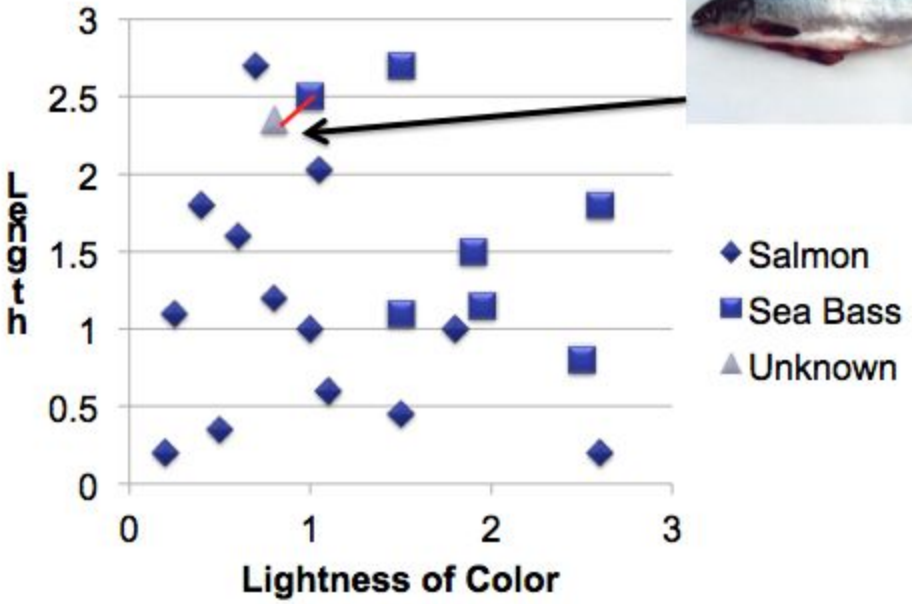

Classification of New Item

- K-nearest neighbor (kNN) can be more robust
  - Find the first few closest
  - Vote

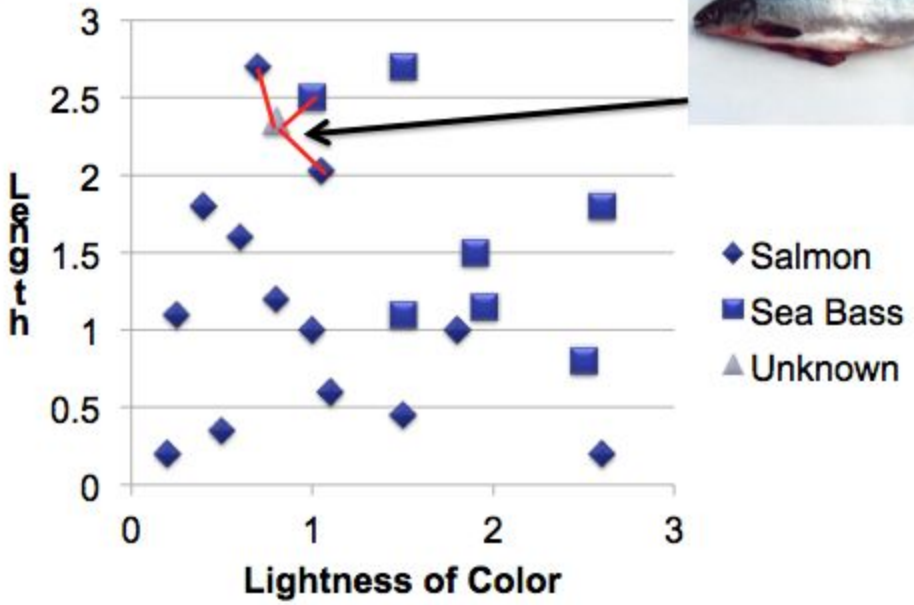

Applying KNN Classifiers

Need to know the distance to examples in the classified set

- What does this mean for more complex examples?
  - Handwriting
  - Faces

Applying KNN Classifiers

- Look at pixels to see how much is different
- Big data sets provide closer matches

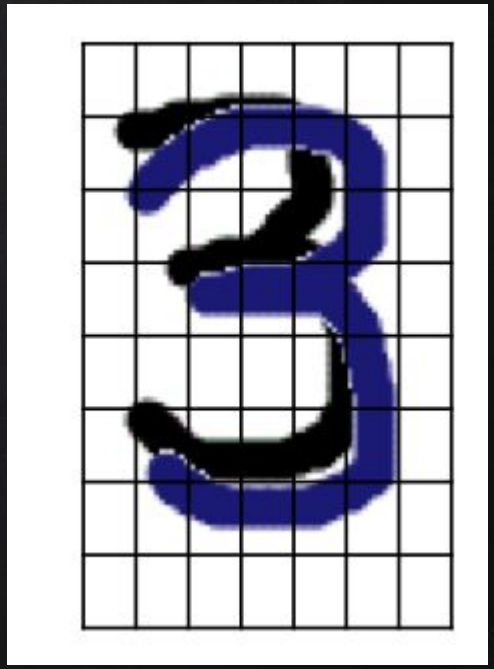

Helmet detection example:
https://www.youtube.com/watch?v=i9jxJzQsZlU

Machine Learning

Modeling Brains...

Deep Learning

Buzzword for, rebranding of neural networks

Modeling Brains

What are everyday computer systems good at…and not so good at?

Good at	*Not so good at*
Rule-based systems: doing what the programmer wants them to do	Dealing with noisy data
	Dealing with unknown environment data
	Reacting quickly
	Fault tolerance
	Adapting to circumstances

What are neural networks?

- Models of the brain and nervous system
  - Highly simplified
- Highly parallel
  - Process information much more like the brain than a serial computer
- Learning
- Very simple principles: very complex behaviours

Where can neural networks help?

- When we can not formulate an algorithmic solution
- When we can get lots of examples of the behavior we require
- Learning from experience

Inspiration from neurobiology

- A neuron: many-inputs and 1-output unit
- Output can be excited or not excited
- Incoming signals from other neurons determine if the neuron shall excite (fire)
- Output subject to attenuation in the synapses, which are junction parts of the neuron

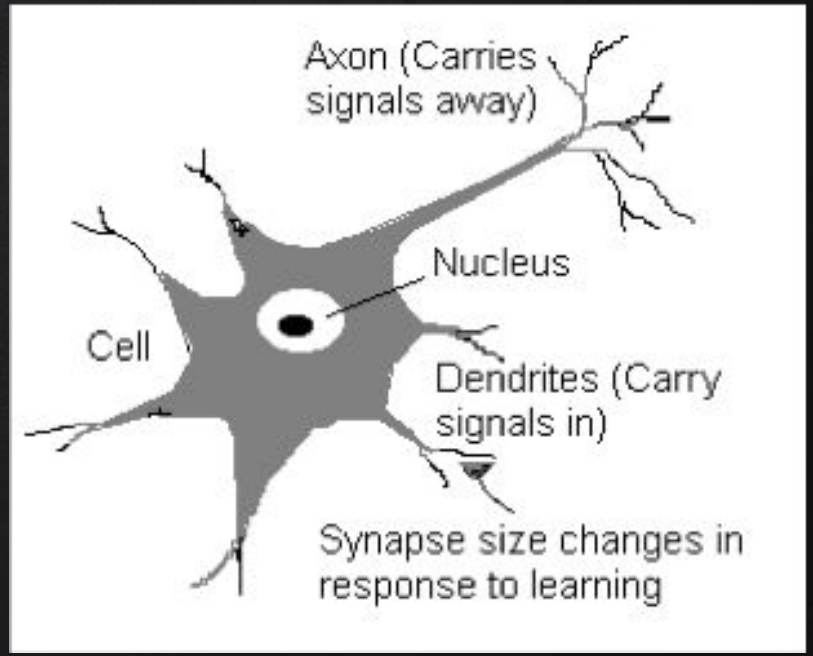

Biological Structure

- Neurons connect in a vast network
  - Billions of neurons
  - Thousands of connections for each

Numbers of Neurons

- Fruit fly: 100, 000
- Cockroach: 1 million
- Even small networks show complex behavior

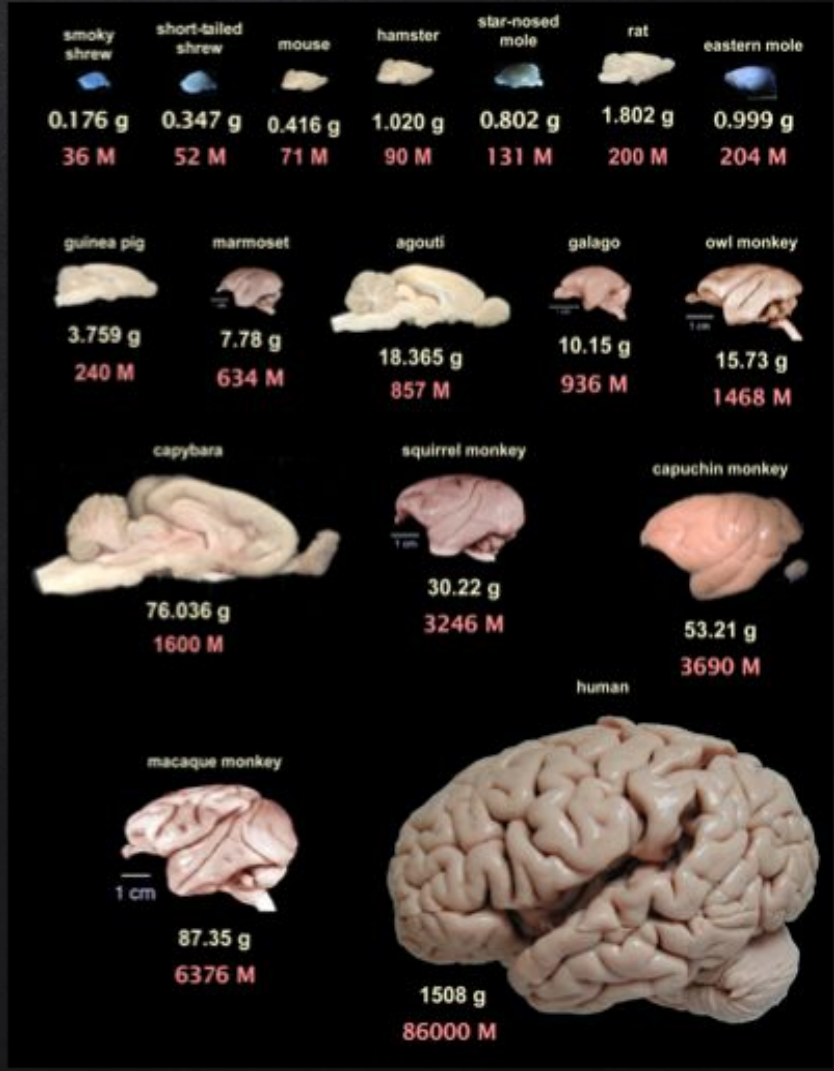

Biological Neural Nets

- Pigeons as art experts (Watanabe et al. 1995)
- Experiment:
  a. Pigeon in Skinner box
  b. Present paintings of 2 different artists (Chagall v.s. Van Gogh)
  c. Reward for pecking when presented a particular artist (e.g. Van Gogh)

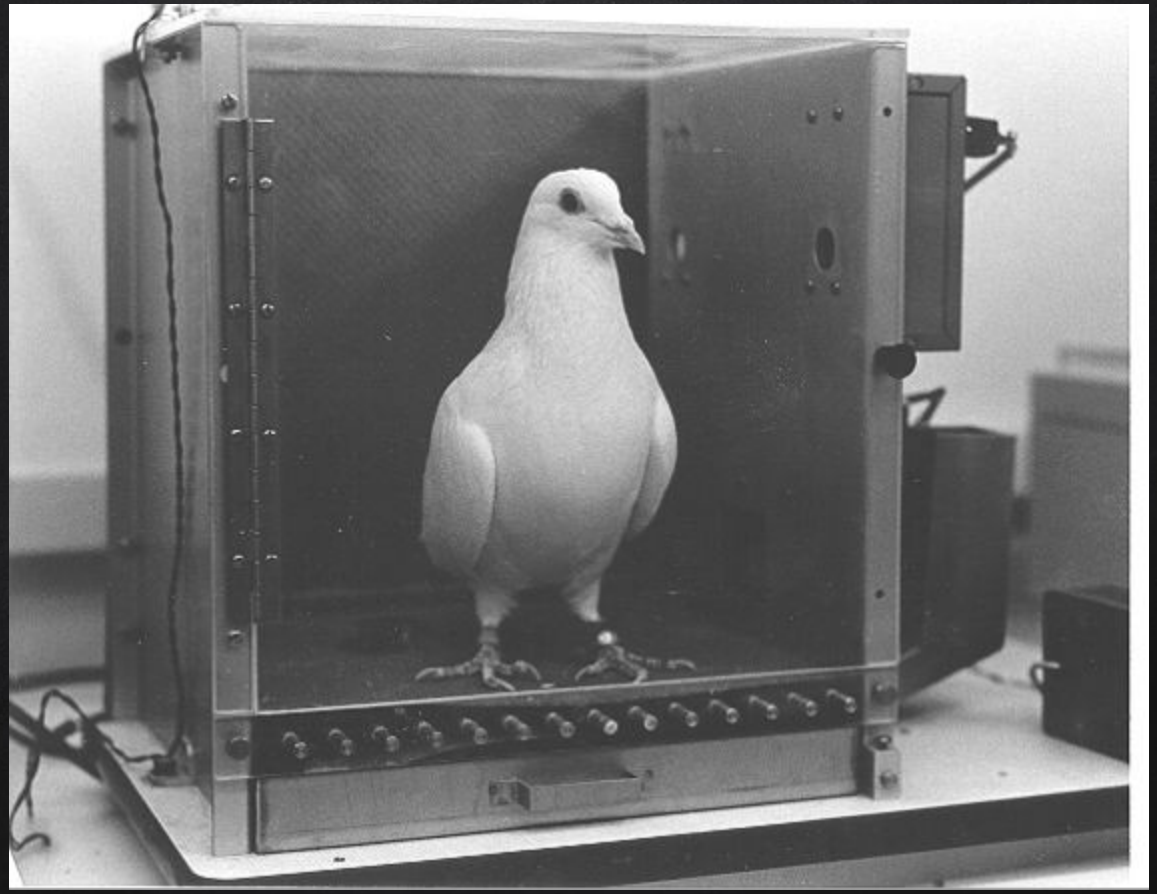

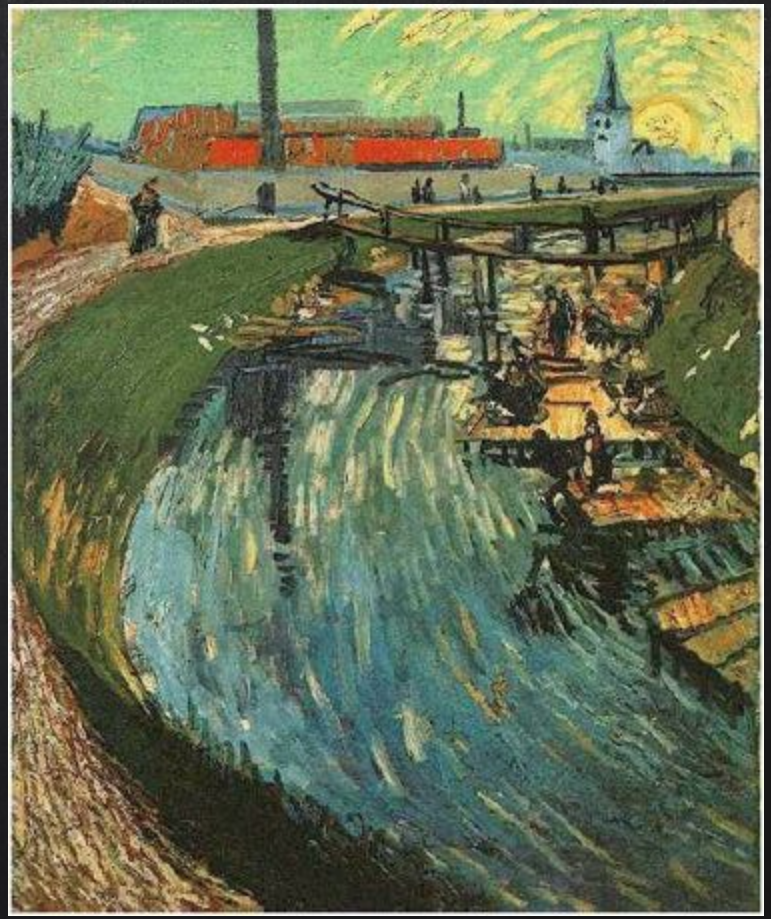

- Pigeons were able to discriminate between Van Gogh and Chagall with 95% accuracy
  - When presented with pictures they had been trained on
- Discrimination still 85% successful for previously unseen paintings of the artists

- Pigeons do not simply memorise the pictures
- They can extract and recognise patterns (the style)
- They generalise from the already seen to make predictions

This is what neural networks (biological and artificial) are good at,

Unlike conventional computer...

Artificial Neural Nets (ANNs)

The basics: ANNs incorporate the 2 fundamental components of biological neural nets:

1. Neurons (nodes)
2. Synapses (connection weights)

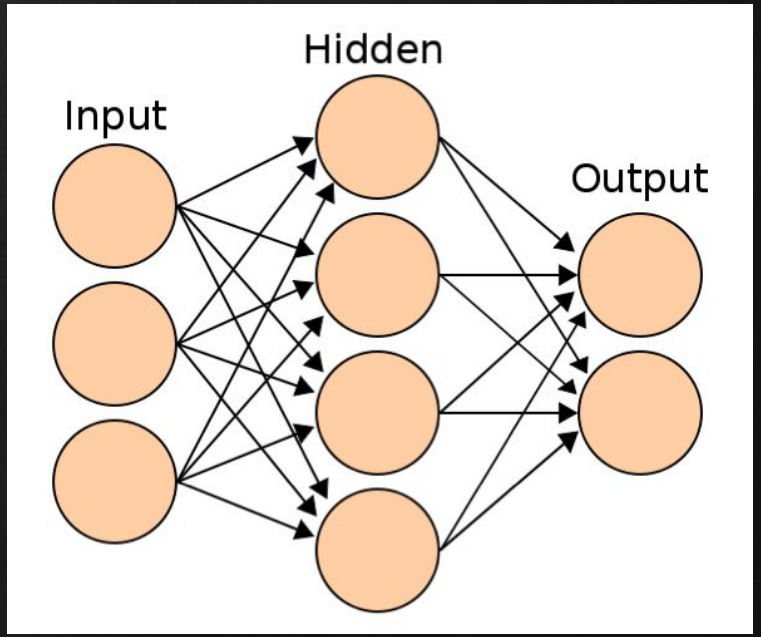

Structure:

- Input layer
- Middle layer(s)
- OUtput layer

Training Neural Nets

- Present training data to input layer
- Adjust connection weights until output shows correct result
- Many different mathematical techniques
- Example
  - Recognize digits
  - http://www.sund.de/netze/applets/bpn/bpn2/ochre.html

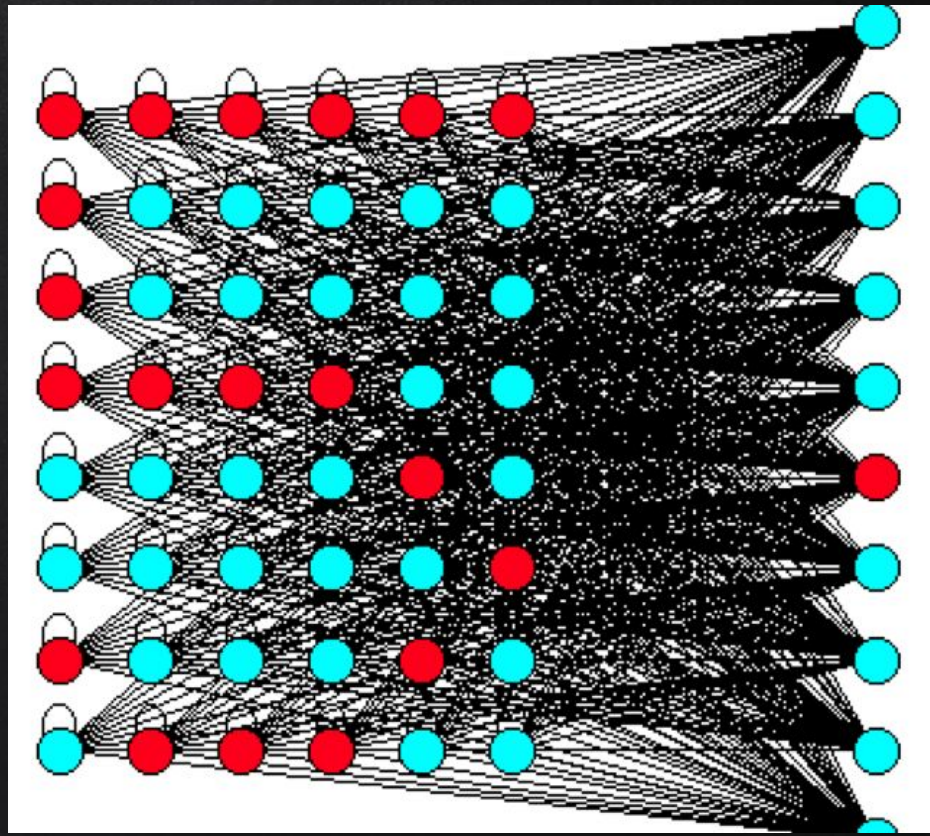

Example Application

Improving portfolio returns

A major Japanese securities company decided to user neural computing in order to develop better prediction models. A neural network was trained on 33 months' worth of historical data. This data contained a variety of economic indicators such as turnover, previous share values, interest rates and exchange rates. The network was able to learn the complex relations between the indicators and how they contribute to the overall prediction. Once trained it was then in a position to make predictions based on "live" economic indicators.

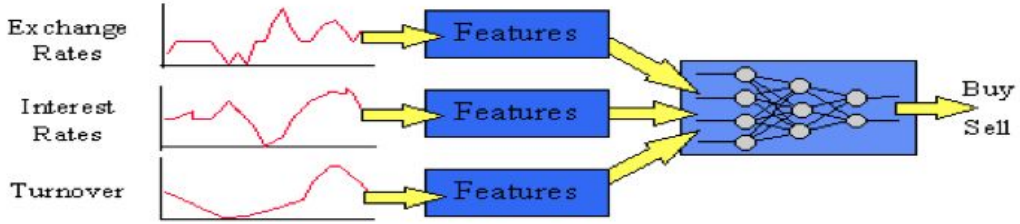

The neural network-based system is able to make faster and more accurate predictions than before. It is also more flexible since it can be retrained at any time in order to accommodate changes in stock market trading conditions. Overall the system outperforms statistical methods by a factor of 19%, which in the case of a £1 million portfolio means a gain of £190,000. The system can therefore make a considerable difference on returns.

Making predictions based on key indicators

- Predicting gas and electricity supply and demand
- Predicting sales and customer trends
- Predicting the route of a projectile
- Predicting crop yields

Artificial Life

Artificial Life

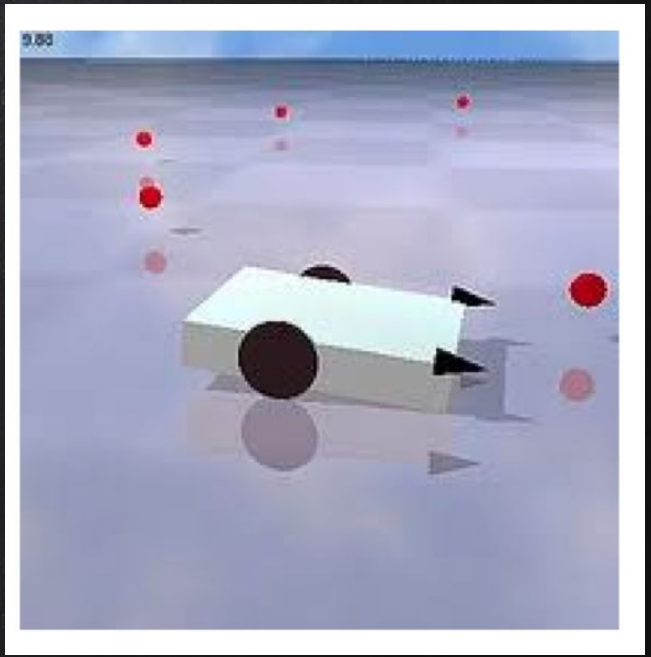

- Simulate creatures in 3D environment
  - Virtual sensors (eyes) connect to ANN
  - ANN controls virtual muscles
- Let creatures:
  - Chase / eat each other
  - Evolve movement
  - Encode design using artificial genetics
- Can we evolve mechanical designs? City layouts? Consciousness?

THANKS!

Any questions?

You can find me at
beiwang@sci.utah.edu

http://www.sci.utah.edu/~beiwang/teaching/cs1060.html

Credits

Special thanks to all the people who made and released these awesome resources for free:

- Presentation template by SlidesCarnival
- Photographs by Unsplash