What's Going on in Artificial Intelligence?

John McKechnie – Project Manager

Overview

a. AI in the news

b. Some basic concepts

c. Neural networks and deep learning

d. Are they coming for my job?
Are they coming for me?

AI in the news

NBS

AlaphaGo

AlphaGo beats Lee Sedol at Go

Image:
https://www.scientificamerican.com/article/how-the-computer-beat-the-go-master/

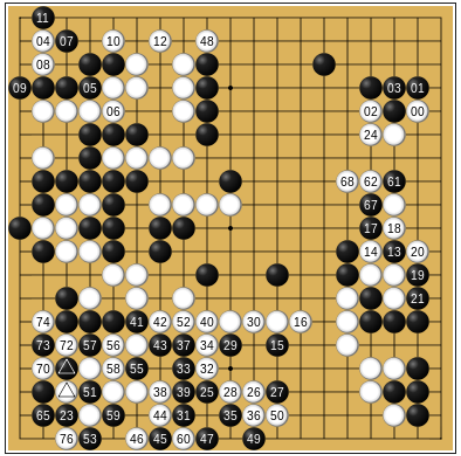

Moves 100–176 (122 at 113, 154 at △, 163 at 145, 164 at 151, 166 and 171 at 160, 169 at 145, 175 at △)

Image:
https://en.wikipedia.org/wiki/AlphaGo_versus_Lee_Sedol

Experts thought this was at least a decade away – why?

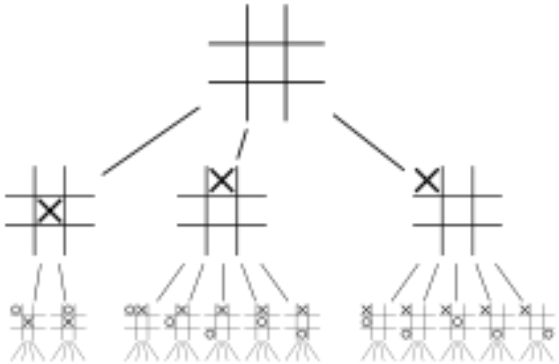

Image:
https://en.wikipedia.org/wiki/Game_tree

Average branching factor:
Noughts & crosses – 4
Chess – 20
Go – 250

Board positions to evaluate when looking 3 moves ahead:
Noughts & crosses – 64
Chess – 8,000
Go – 15,625,000

Self-driving cars – Google launches Waymo

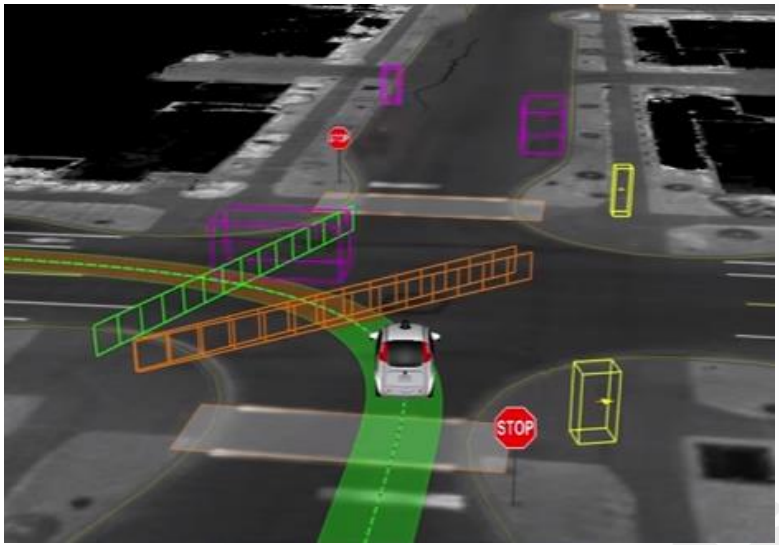

https://waymo.com/

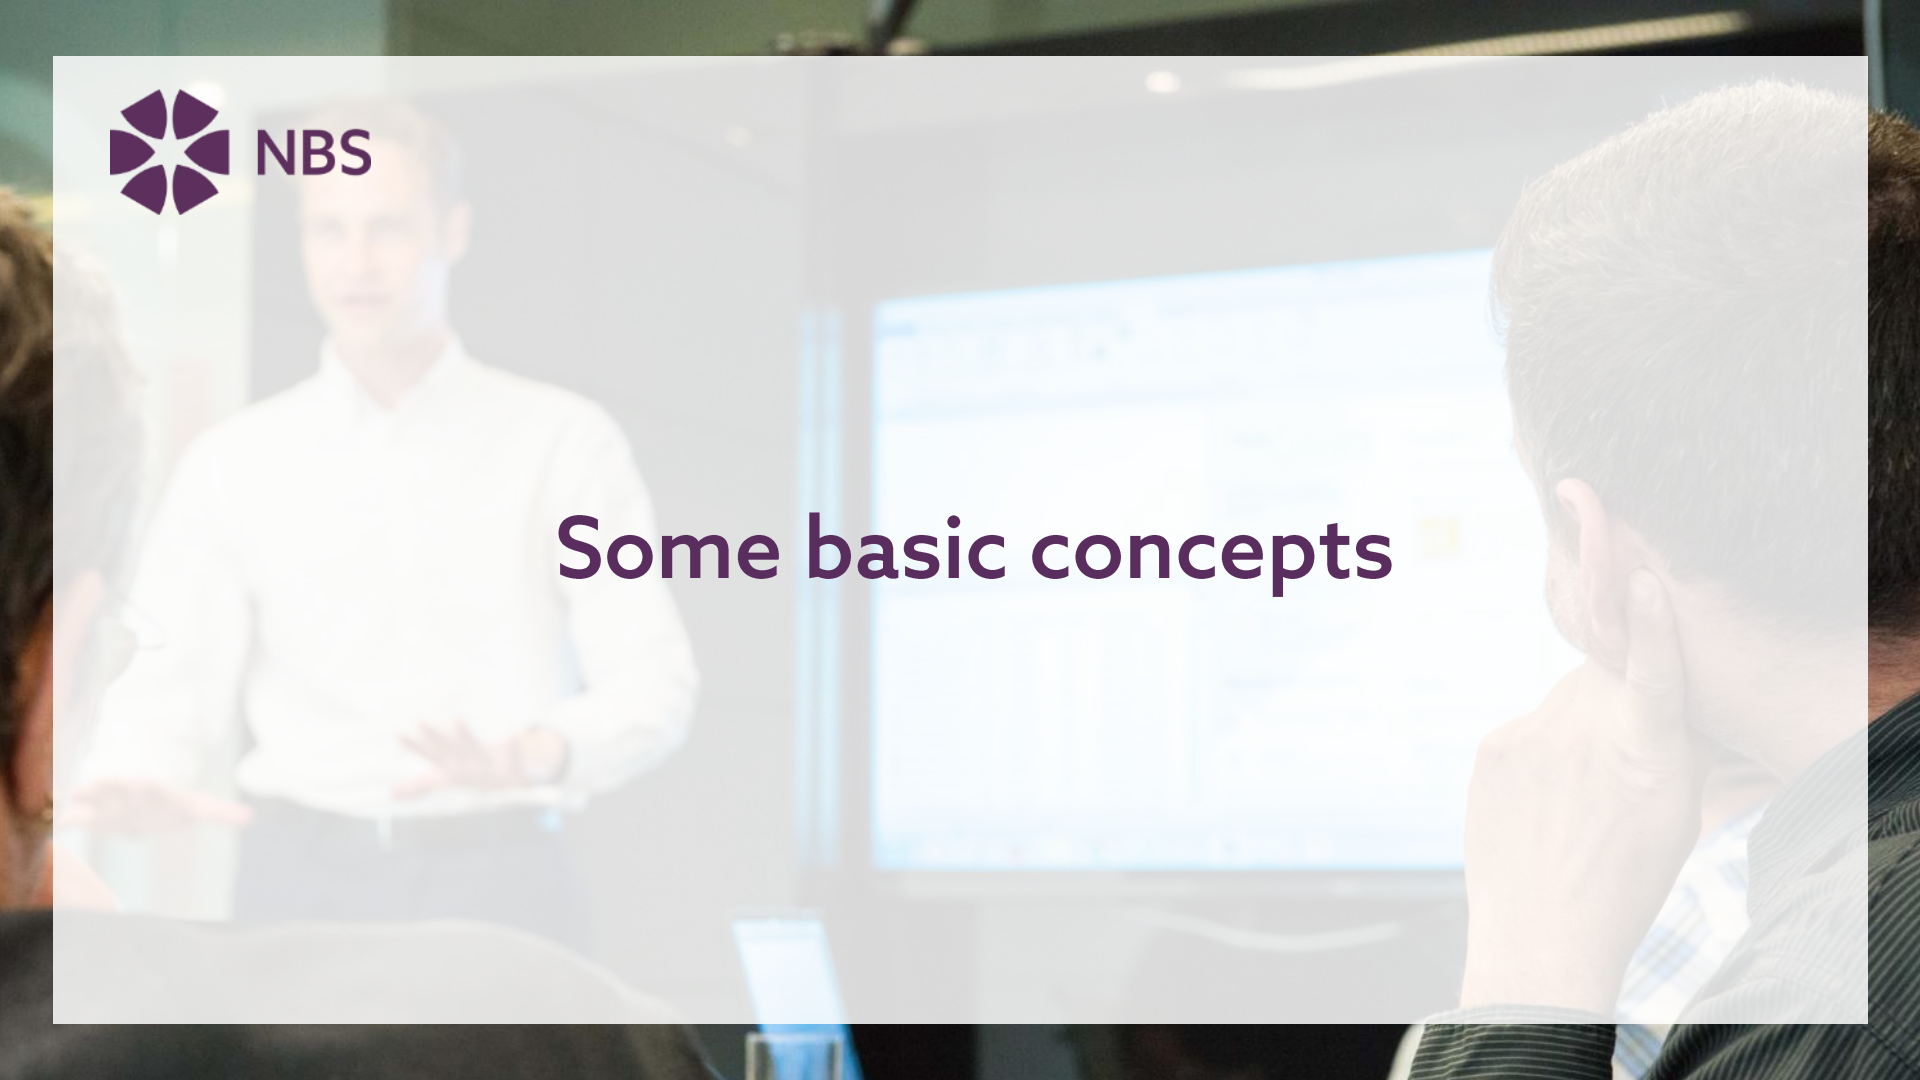

Some basic concepts

Models

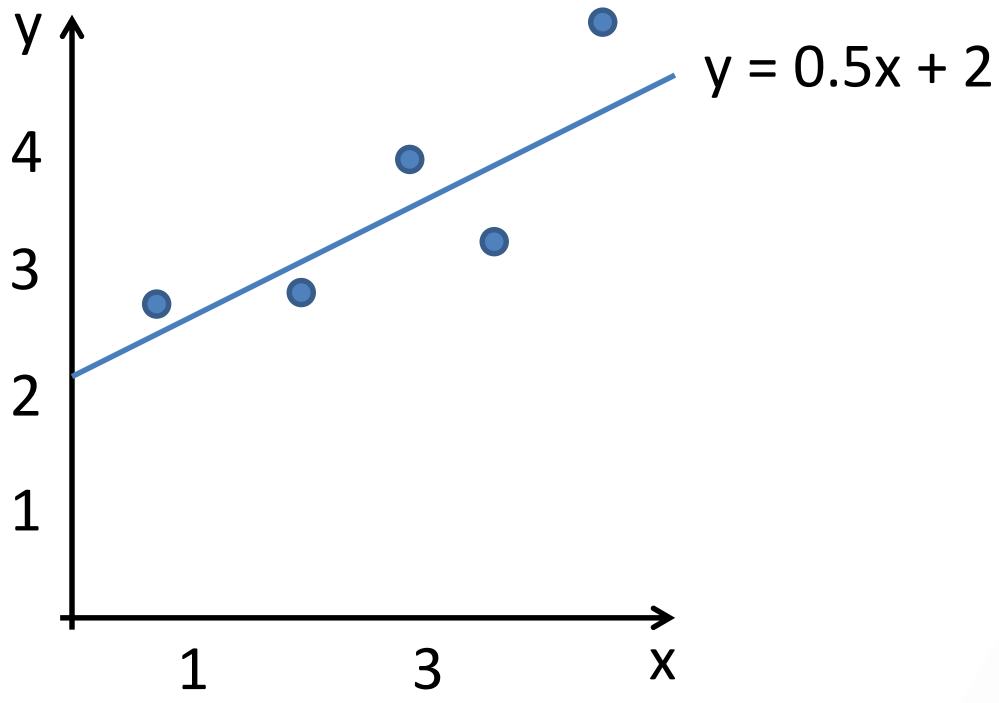

data

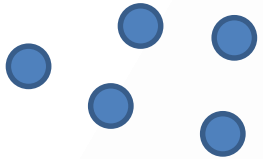

model

$y = 0.5x + 2$

Learning algorithm

data

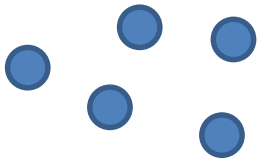

Learning Algorithm

model

$y = 0.5x + 2$

Patient records

Insurance claims

Images/video

Diagnostic rules

Fraud detector

Cancer cell detector

A decision tree model

Wears Crocs?	Home city?	Successful?
Yes	Glasgow	Yes
Yes	Edinburgh	No
No	Glasgow	No
Yes	Glasgow	Yes
No	Aberdeen	No
No	Dundee	No

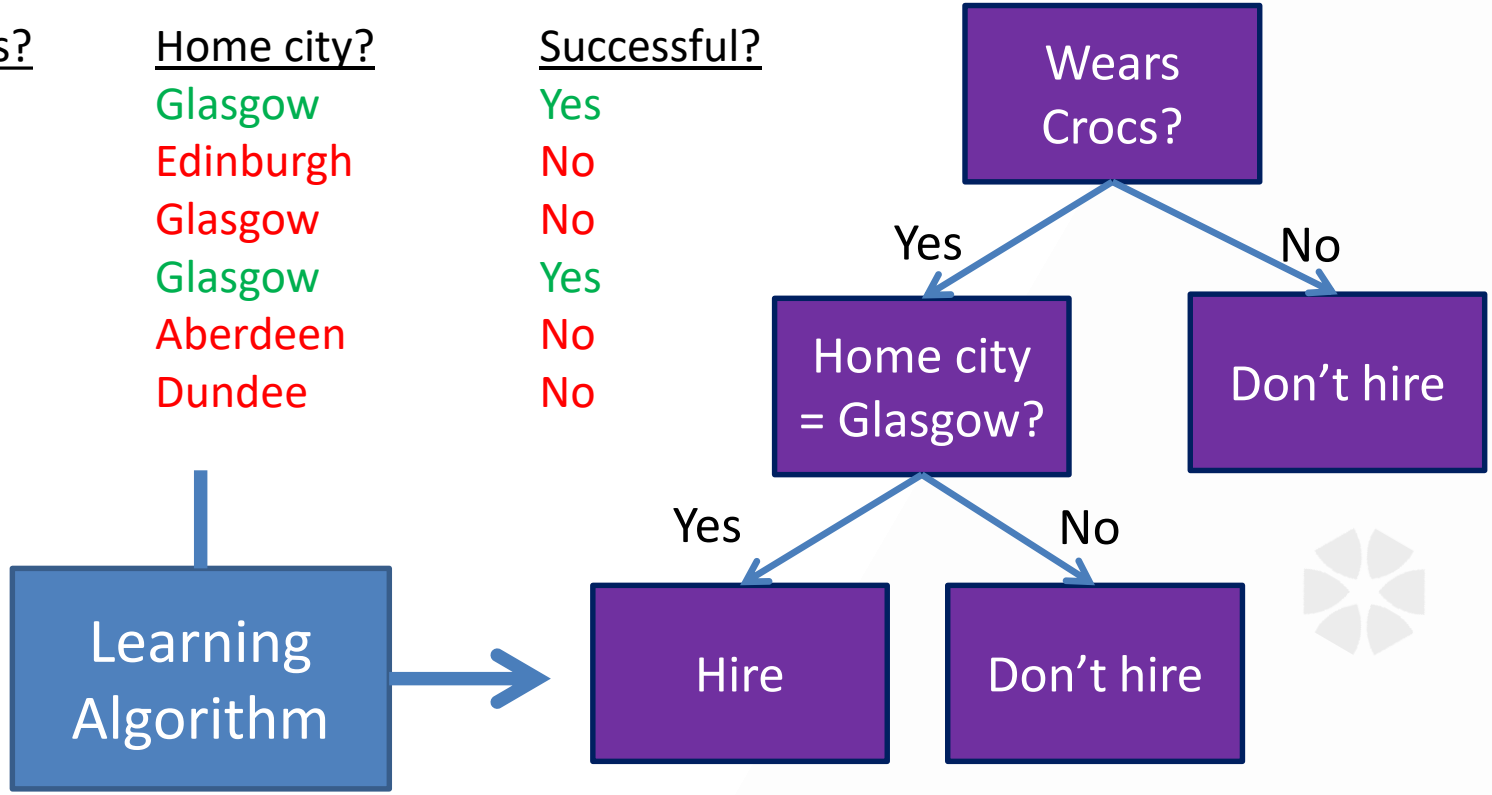

Supervised and unsupervised learning

Wears Crocs?	Home city?	Successful?
Yes	Glasgow	Yes
Yes	Edinburgh	No
No	Glasgow	No
Yes	Glasgow	Yes
No	Aberdeen	No
No	Dundee	No

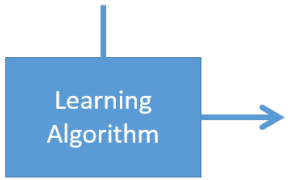

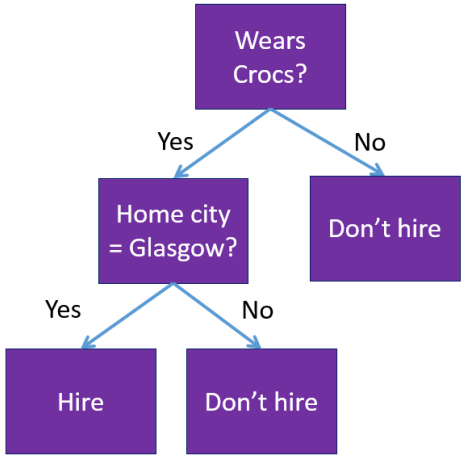

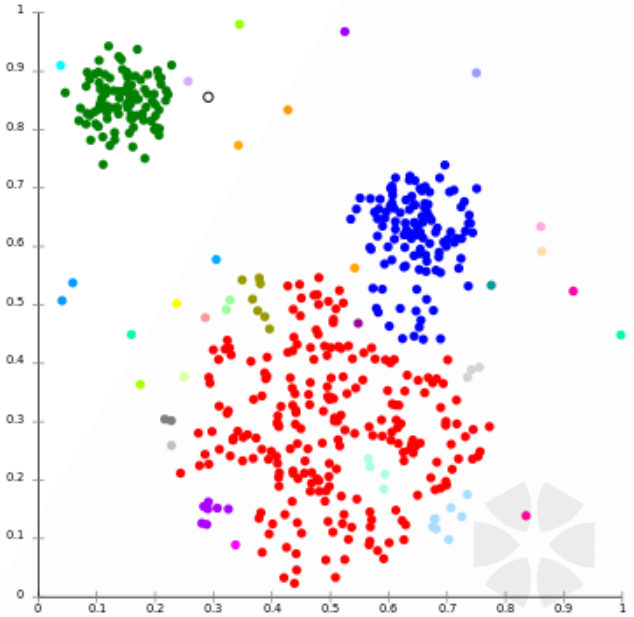

Image: https://en.wikipedia.org/wiki/Cluster_analysis#/media/File:SLINK-Gaussian-data.svg

Supervised

Unsupervised

Reinforcement learning

YouTube Video: https://youtu.be/L4KBBAwF_bE

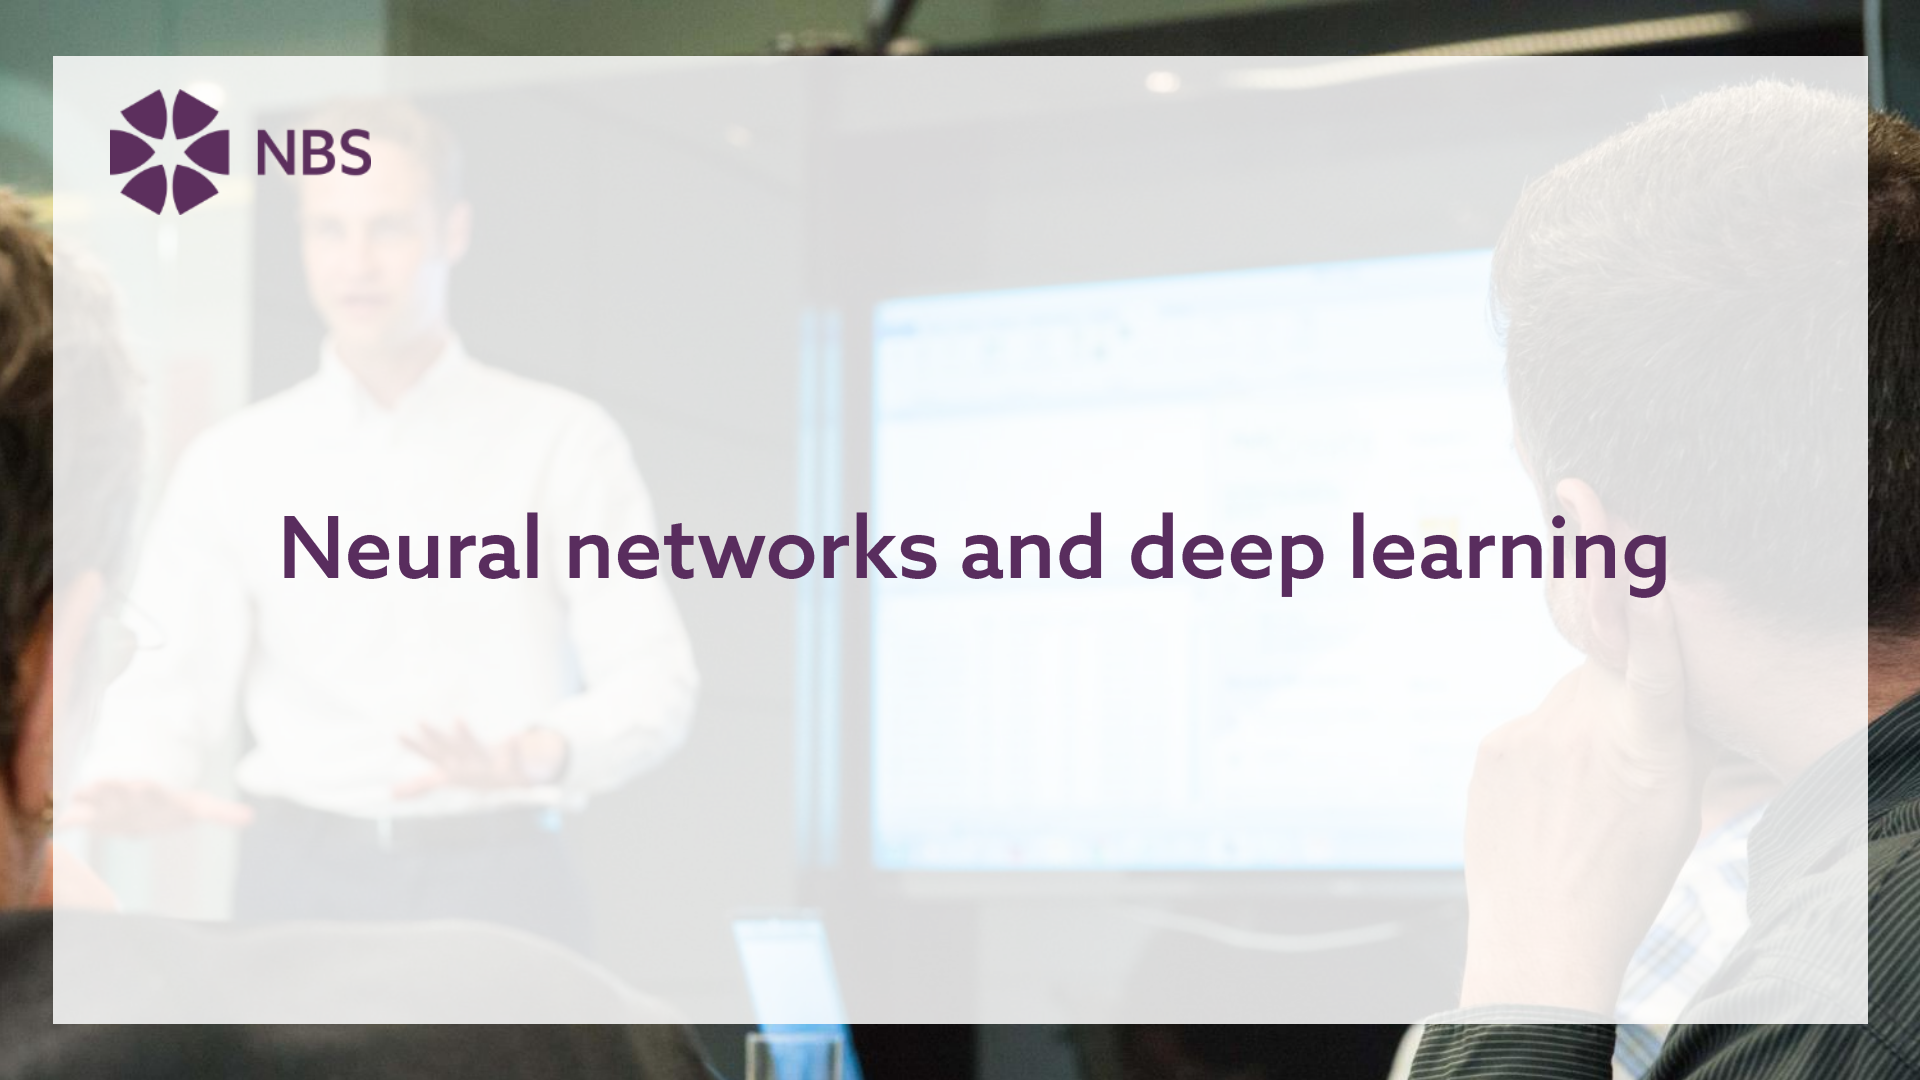

Neural networks and deep learning

Neural networks

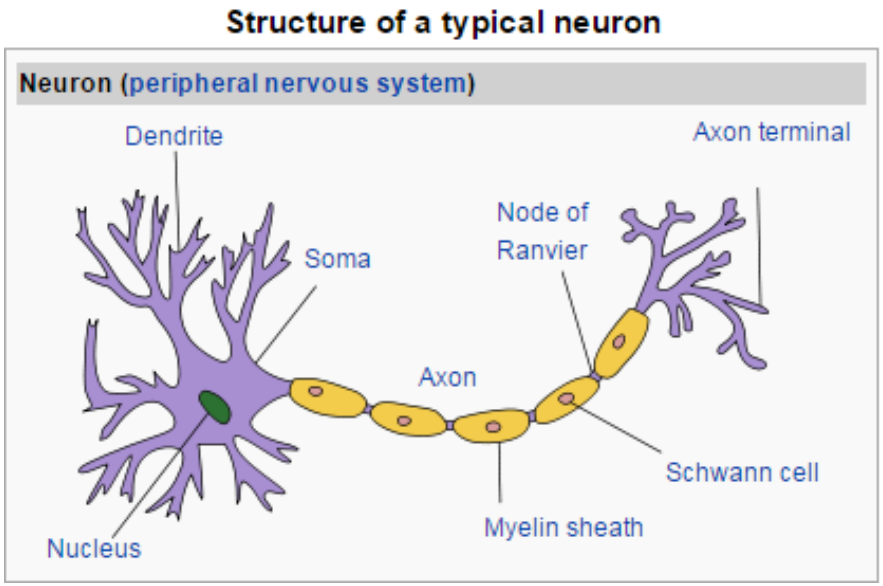

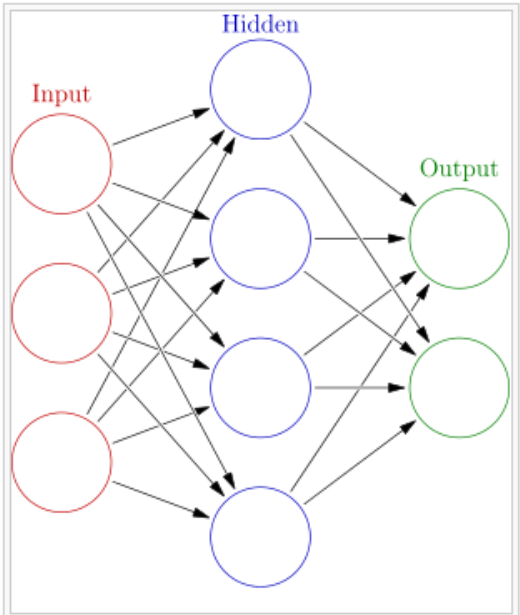

An artificial neural network is an interconnected group of nodes, akin to the vast network of neurons in a brain. Here, each circular node represents an artificial neuron and an arrow represents a connection from the output of one neuron to the input of another.

Images:
https://en.wikipedia.org/wiki/Neuron
https://en.wikipedia.org/wiki/Artificial_neural_network

NBS

playground.tensorflow.org

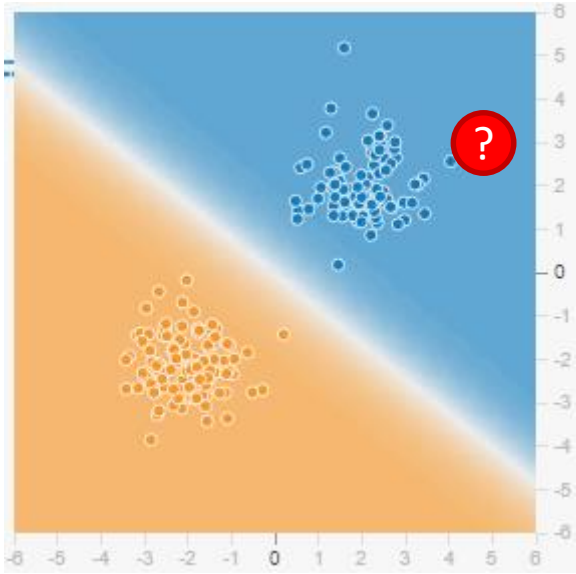

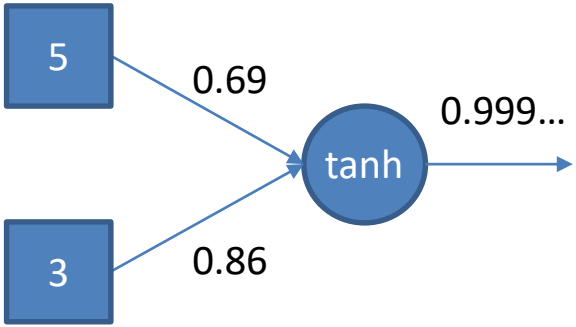

1 means blue

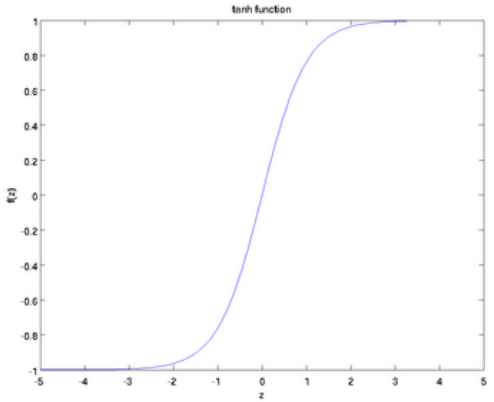

-1 means orange

5 x 0.69 = 3.45

3 x 0.86 = 2.58

3.45 + 2.58 = 6.03

tanh(6.03) = 0.999…

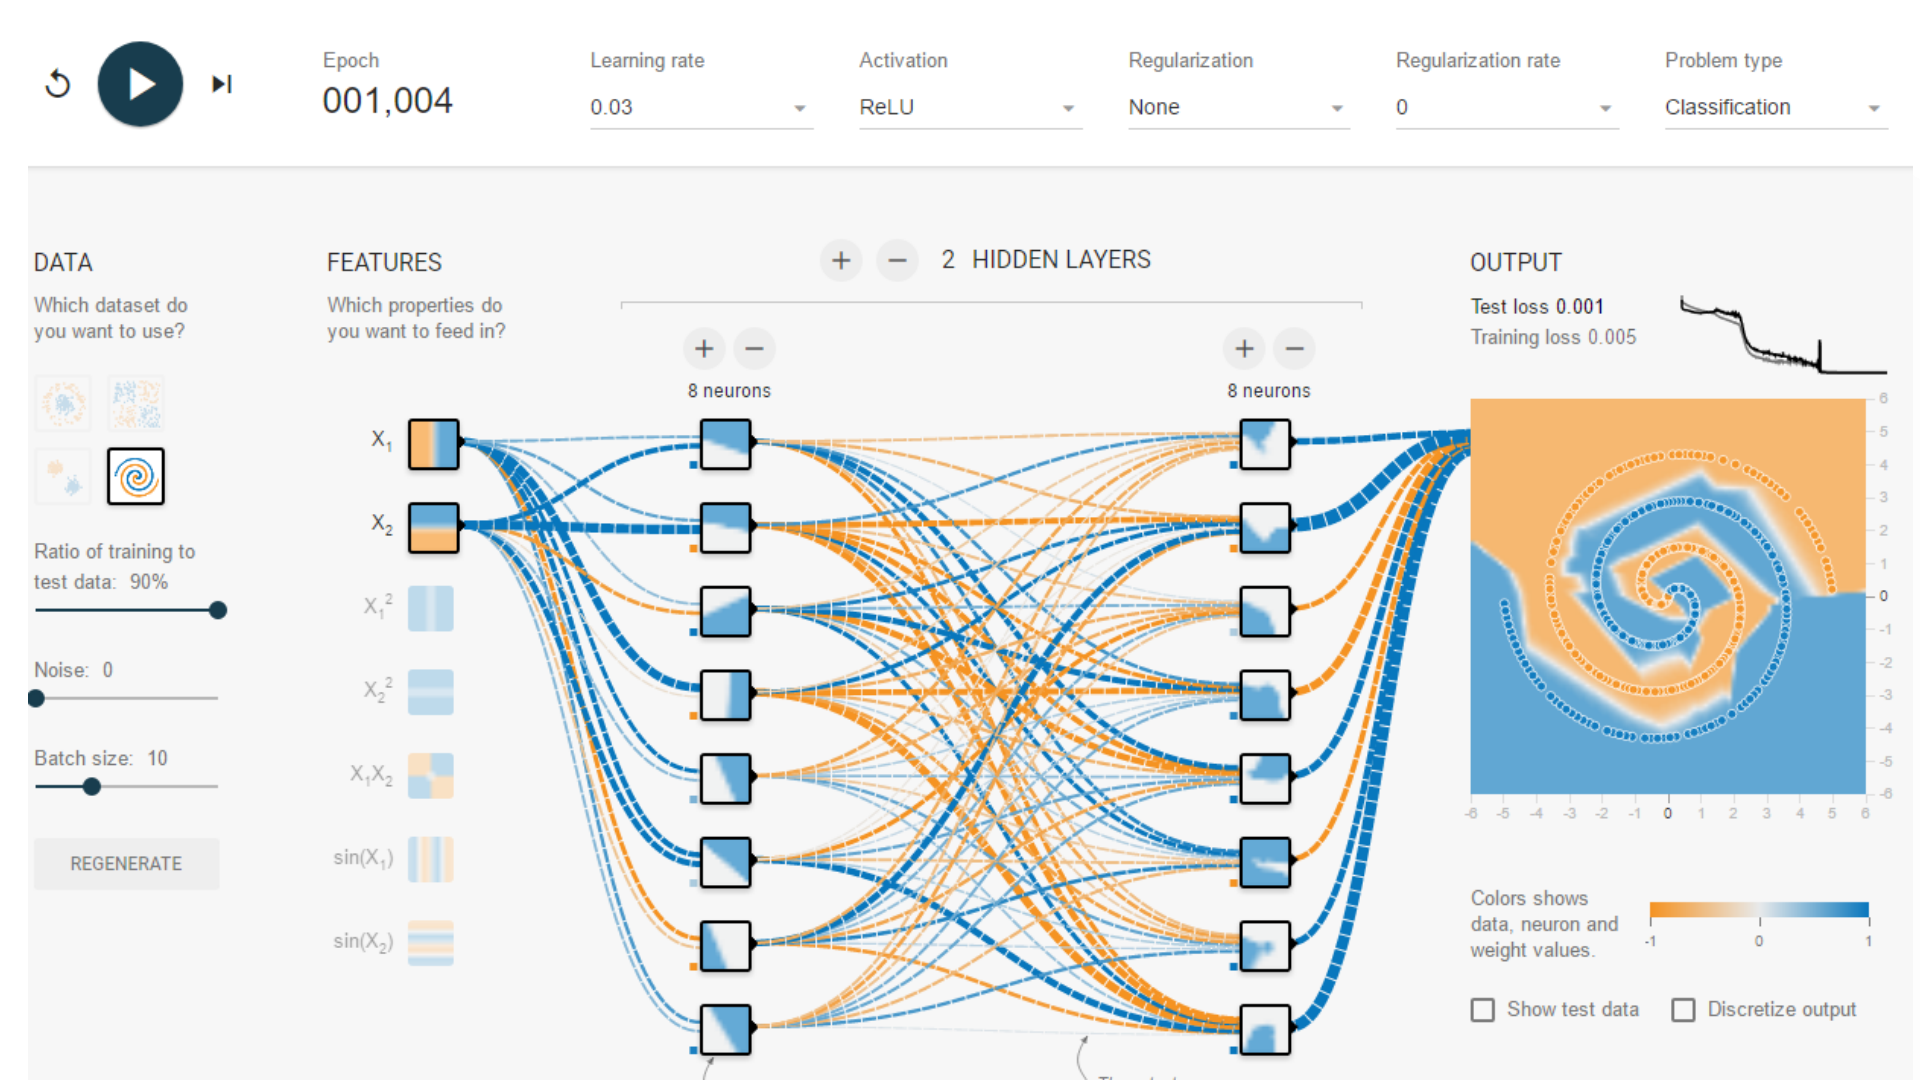
Epoch
001,004
Learning rate
0.03
Activation
ReLU
Regularization
None
Regularization rate
0
Problem type
Classification
DATA
Which dataset do you want to use?
Ratio of training to test data: 90%
Noise: 0
Batch size: 10
REGENERATE
FEATURES
Which properties do you want to feed in?
2 HIDDEN LAYERS
8 neurons
8 neurons
X1
X2
X1²
X2²
X1X2
sin(X1)
sin(X2)
OUTPUT
Test loss 0.001
Training loss 0.005
Colors shows data, neuron and weight values.
-1
0
1
Show test data
Discretize output

ImageNet

Not unreasonable even when wrong

Primitive features that turned out to be important

ImageNet Classification with Deep Convolutional Neural Networks, Alex Krizhevsky, Ilya Sutskever, Geoffrey E Hinton, NIPS 2012.

ImageNet

Each row of images produced similar behaviour in the network but the similarities in the images are not merely superficial.

ImageNet Images from paper: http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf

Creativity

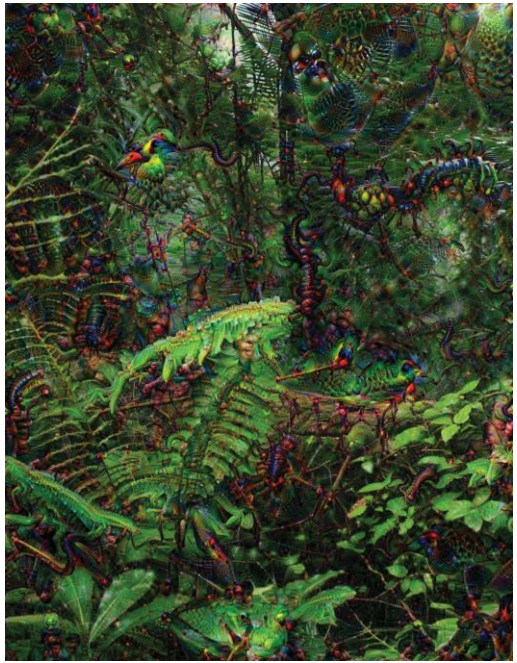

The Dark Secret at the Heart of AI

May/June 2017 Issue

Autodesk Generative Design

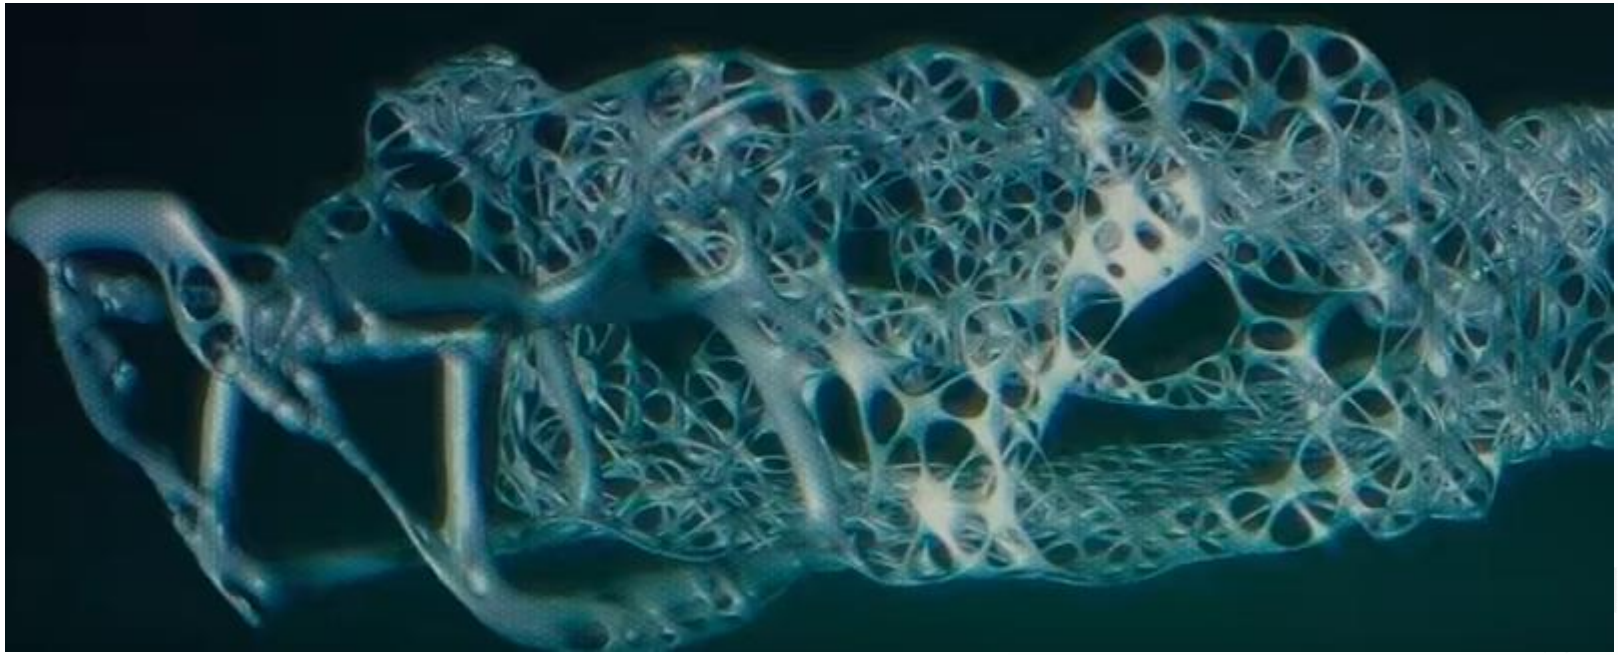

YouTube video: https://www.youtube.com/watch?v=CtYRfMzmWFU

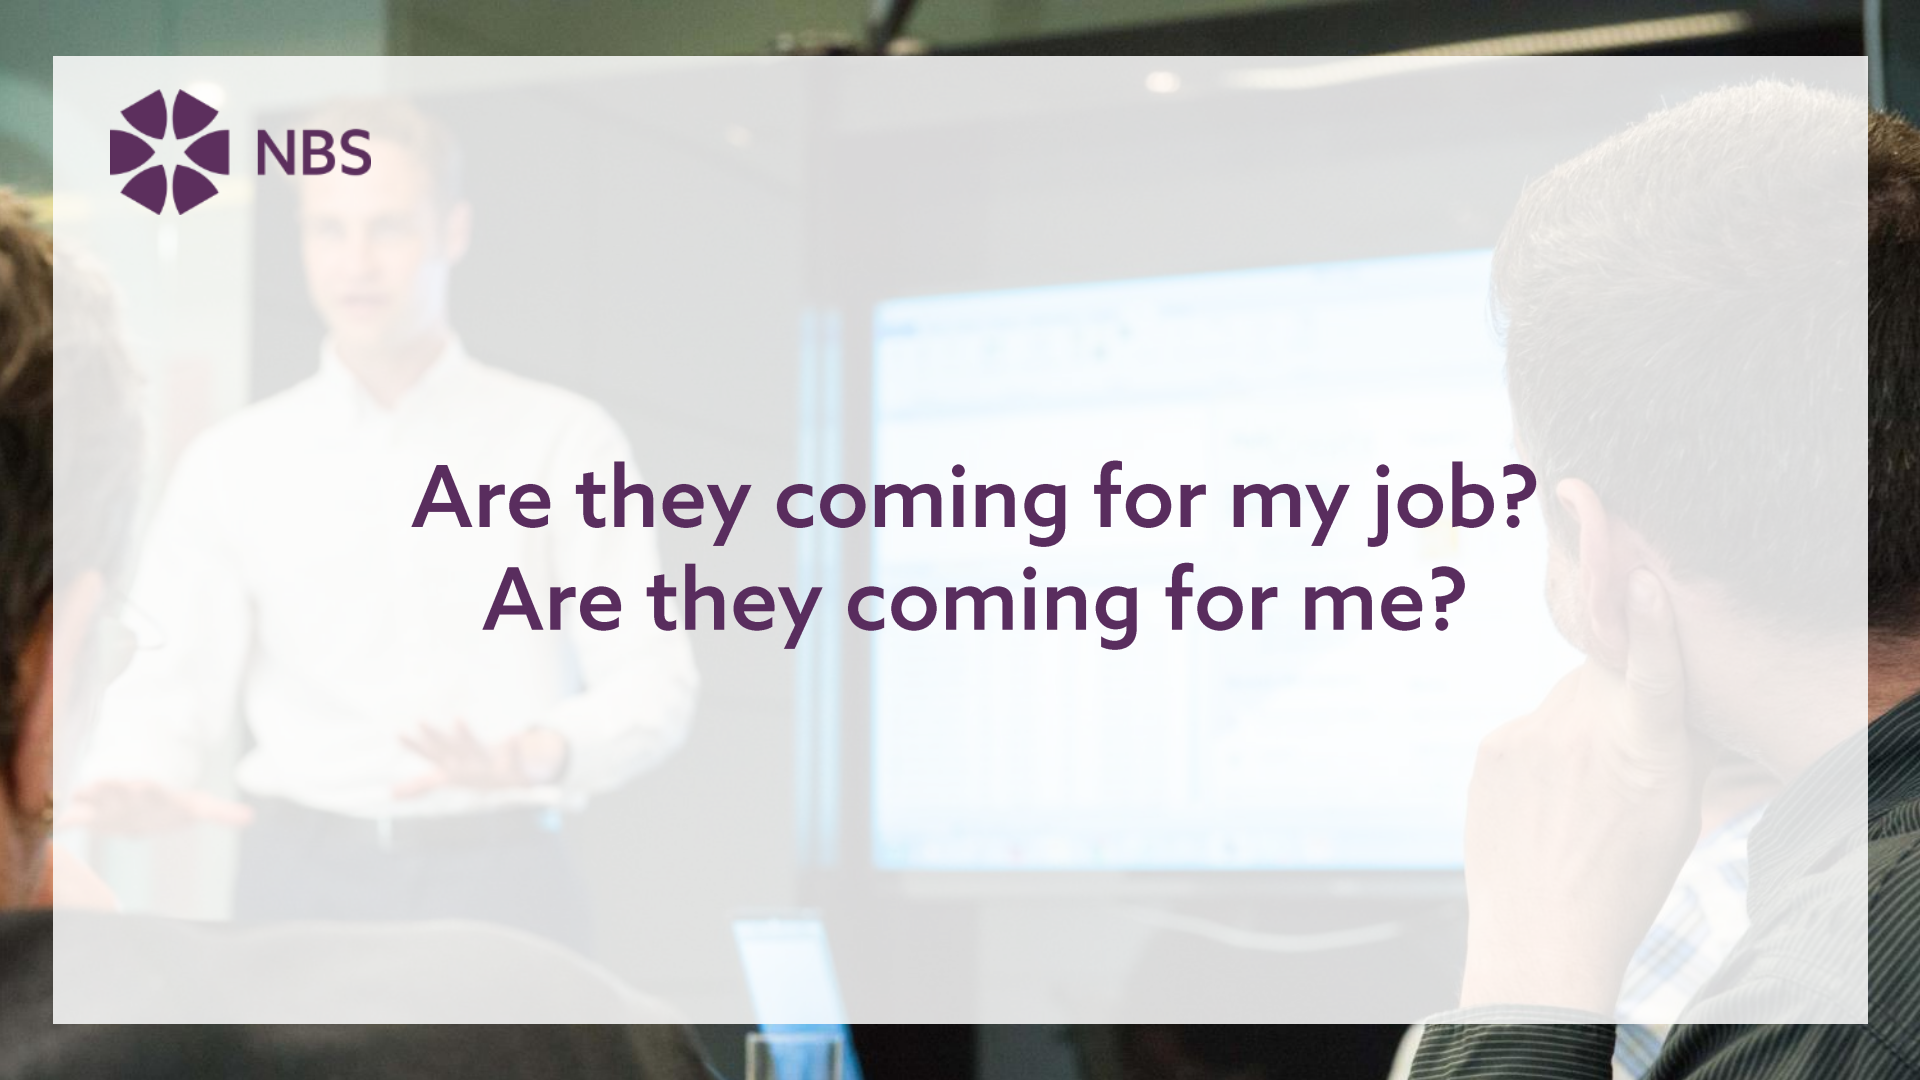

Are they coming for my job?
Are they coming for me?

AI and disruption of the job market

Stockroom worker

Bartender

Soldier

Pharmacist

Farmer

Journalist

Housekeeper

Tellers and clerks

Images: http://www.marketwatch.com/story/9-jobs-robots-already-do-better-than-you-2014-01-27

When will AI take your job?

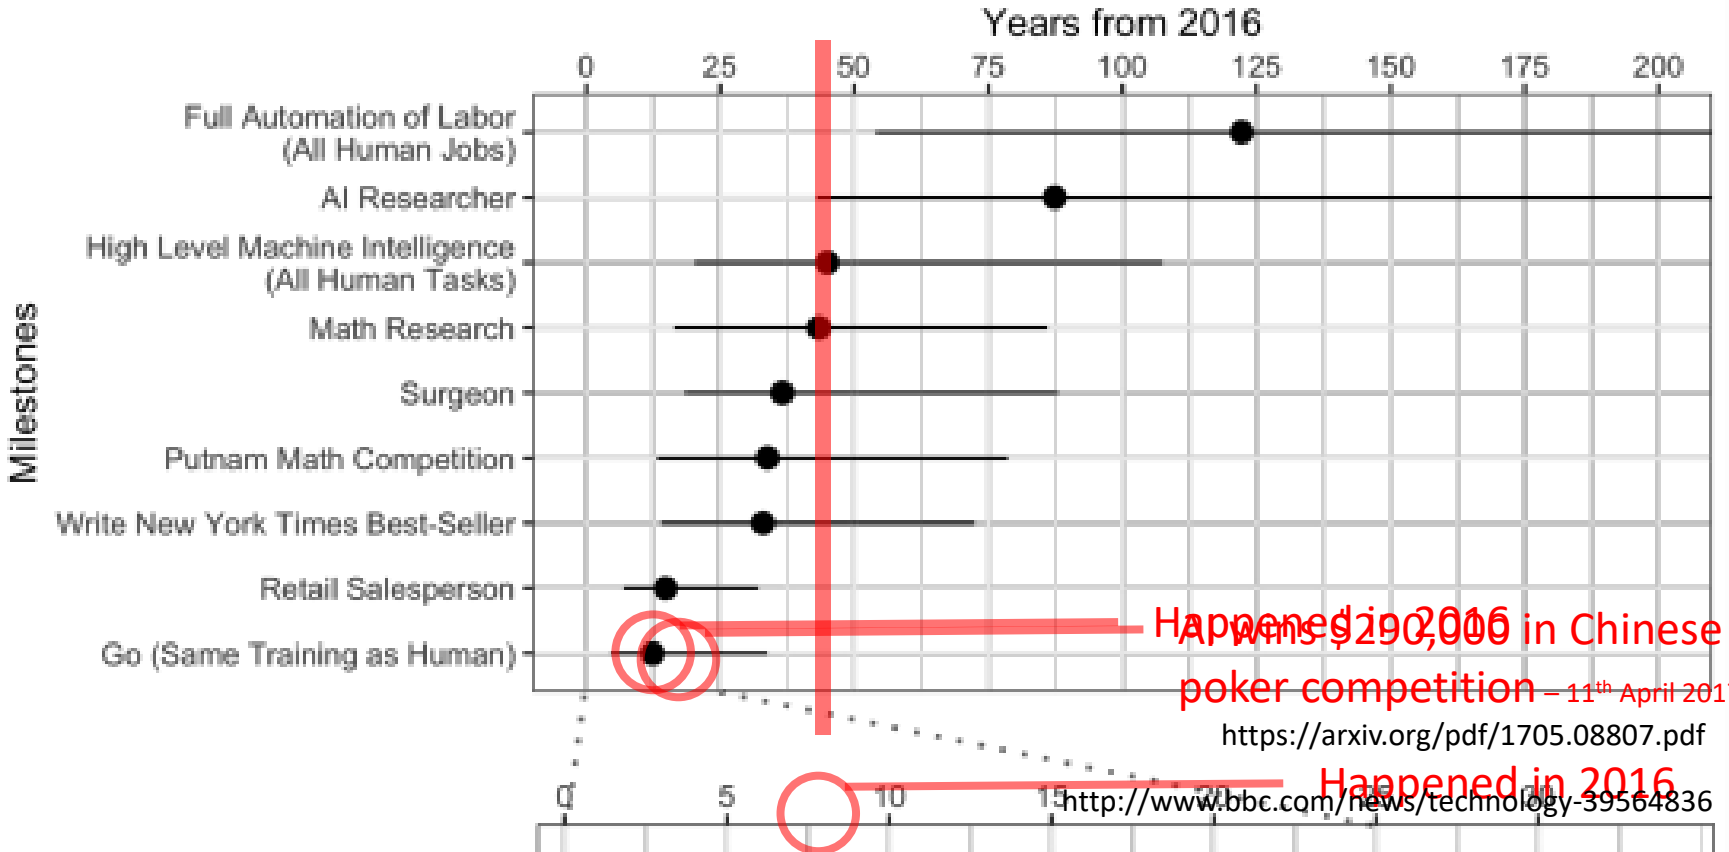

Happened in 2016

AI wins $290,000 in Chinese poker competition – 11th April 2017

https://arxiv.org/pdf/1705.08807.pdf

Happened in 2016

http://www.bbc.com/news/technology-39564836

Value alignment problem

“Uh, your 20th anniversary at 7pm.”

“Well, I did warn you, but you overrode my recommendation.”

“Don’t worry. I arranged for his plane to be delayed. Some kind of computer malfunction.”

“He sends his profound apologies and looks forward to meeting you for lunch tomorrow.”

“I can’t do that. I’m meeting with the secretary-general at 7:30. How could this have happened?”

“Well, what am I going to do? I can’t just tell him I’m too busy.”

“Really? You can do that?”

3 principles for creating safer AI Stuart Russell

Value alignment problem

You're late from work again and the robot has to feed the kids, and the kids are hungry and there's nothing in the fridge. And the robot sees the cat.

And the robot hasn't quite learned the human value function properly, so it doesn't understand the sentimental value of the cat outweighs the nutritional value of the cat.

"Deranged robot cooks kitty for family dinner."

3 principles for creating safer AI Stuart Russell

Intelligent machines and humans = partners

“There's one thing only a human can do.

That's dream.

So let us dream big.”

Don't fear intelligent machines. Work with them Garry Kasparov

References and further reading

- When Will AI Exceed Human Performance? Evidence from AI Experts, Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang and Owain Evans: https://arxiv.org/pdf/1705.08807.pdf
- How the Computer Beat the Go Master, Scientific American Article https://www.scientificamerican.com/article/how-the-computer-beat-the-go-master/
- AlphaGo website, https://deepmind.com/research/alphago/
- 3 Principles For Creating Safer AI, Stuart Russell, TED Talk https://www.ted.com/talks/stuart_russell_how_ai_might_make_us_better_people
- Don't Fear Intelligent Machines. Work With Them, Garry Kasparov, TED Talk: https://www.ted.com/talks/garry_kasparov_don_t_fear_intelligent_machines_work_with_them
- Humans Need Not Apply, YouTube video https://www.youtube.com/watch?v=7Pq-S557XQU
- AI And Deep Learning, A Primer, Frank Chen, Video: http://a16z.com/2016/06/10/ai-deep-learning-machines/

NBS

www.theNBS.com

Questions and discussion...