DEEP LEARNING: MIRACLE OR SNAKE OIL?

Richard Harvey
IT Livery Company Professor of Information Technology, Gresham College

Professor of Computer Science, School of Computing Sciences,
University of East Anglia

#richardwharvey

A brief history of AI

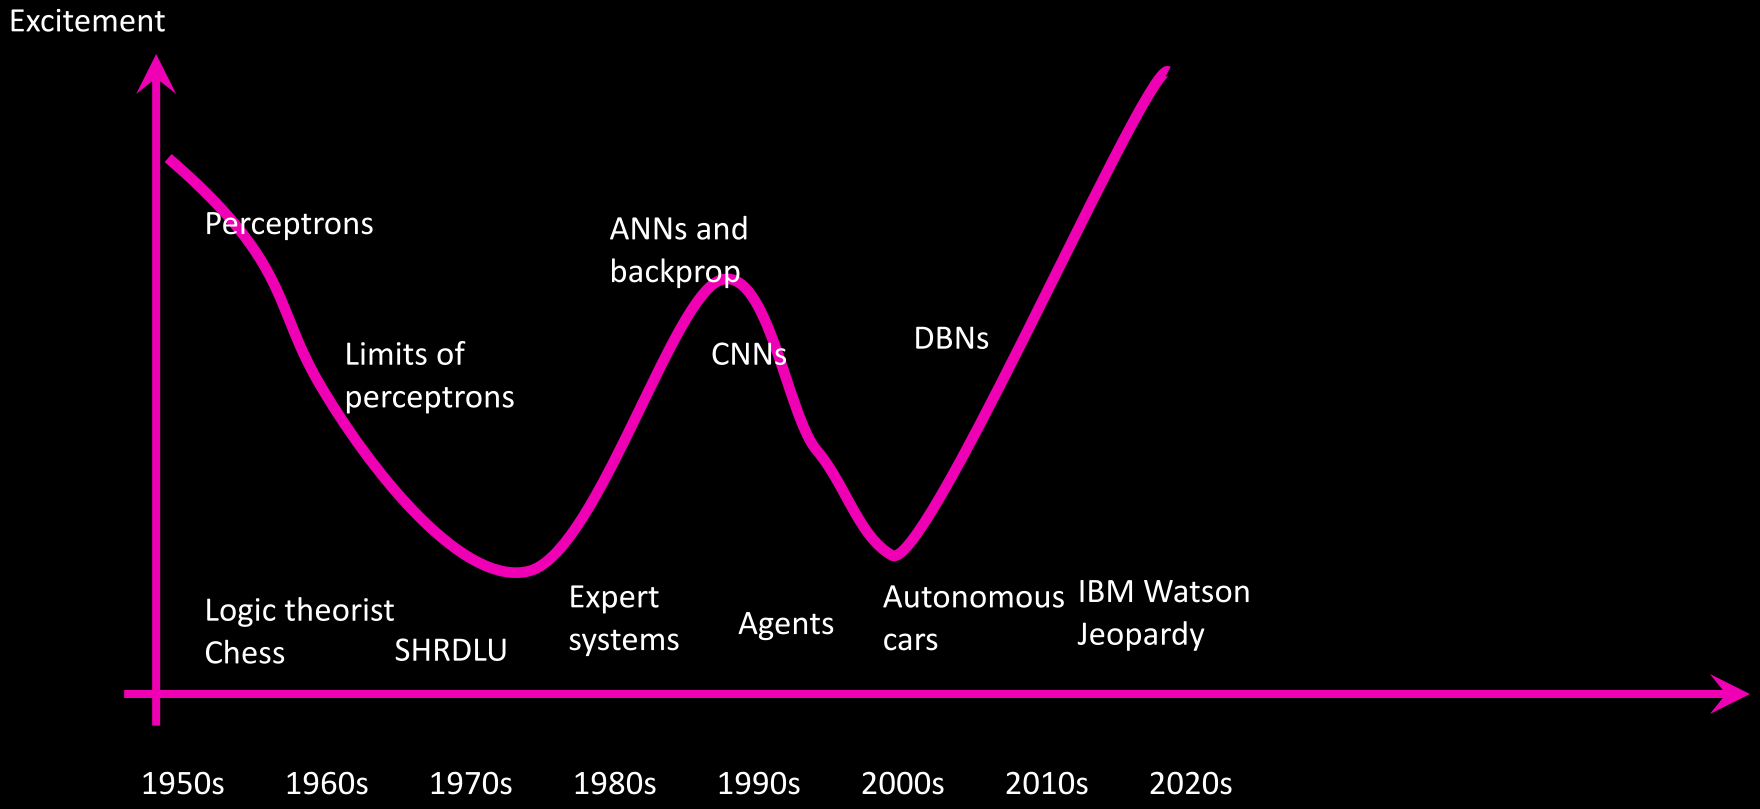

The Lighthill report

Scathing about “Category B systems”:

25 years of non-achievement against grandiose claims

“Artificial intelligence: a general survey,” James Lighthill, 1973, Artificial Intelligence

BBC "Controversy" series 1973 debate …
"The general purpose robot is a mirage"

[00:00:56] failures continually occurred also in computer recognition of human speech or handwritten letters and in automatic proving of theorems in higher mathematics. *[00:00:56][0.0]*

Automatic transcription provided by Trint.com

AND THAT WAS THE END OF AI

FOR A WHILE…

AI: Classification

Classification

A particular type of AI known as *machine learning*

Supervised classification aims to

Learn the relationship between some numbers (*feature*s) and *classe*

IMPORTANT CAVEAT

This lecture is about one class of classifiers, neural networks. There are others.

3 0.2 0.5

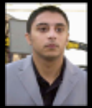

3 0.1 0.6

1 0.01 0.1

1 0.03 0.01

1 0.02 0.5

2 0.02 0.5

2 0.01 0.7

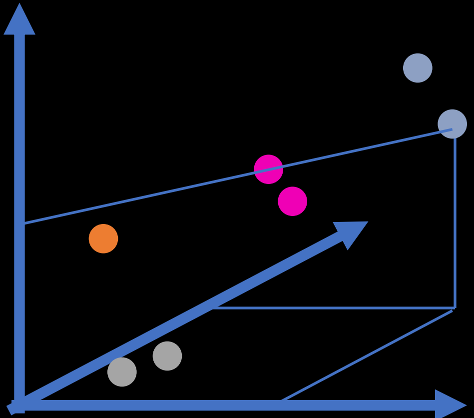

3 0.2 0.5

3 0.1 0.6

Feature vectors

Classes

1 0.01 0.1

Training data

1 0.03 0.01

1 0.02 0.5

2 0.02 0.5

Test data

0 0.21 0.3

?

2 0.01 0.7

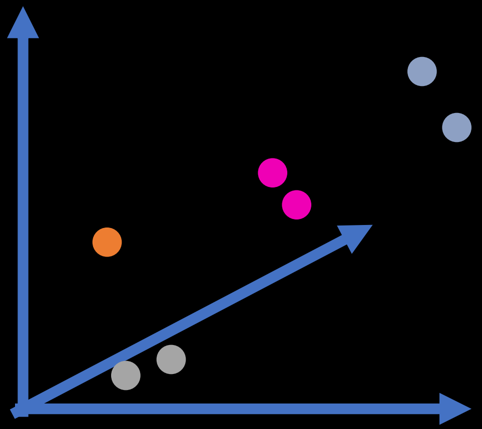

Why “neural” networks?

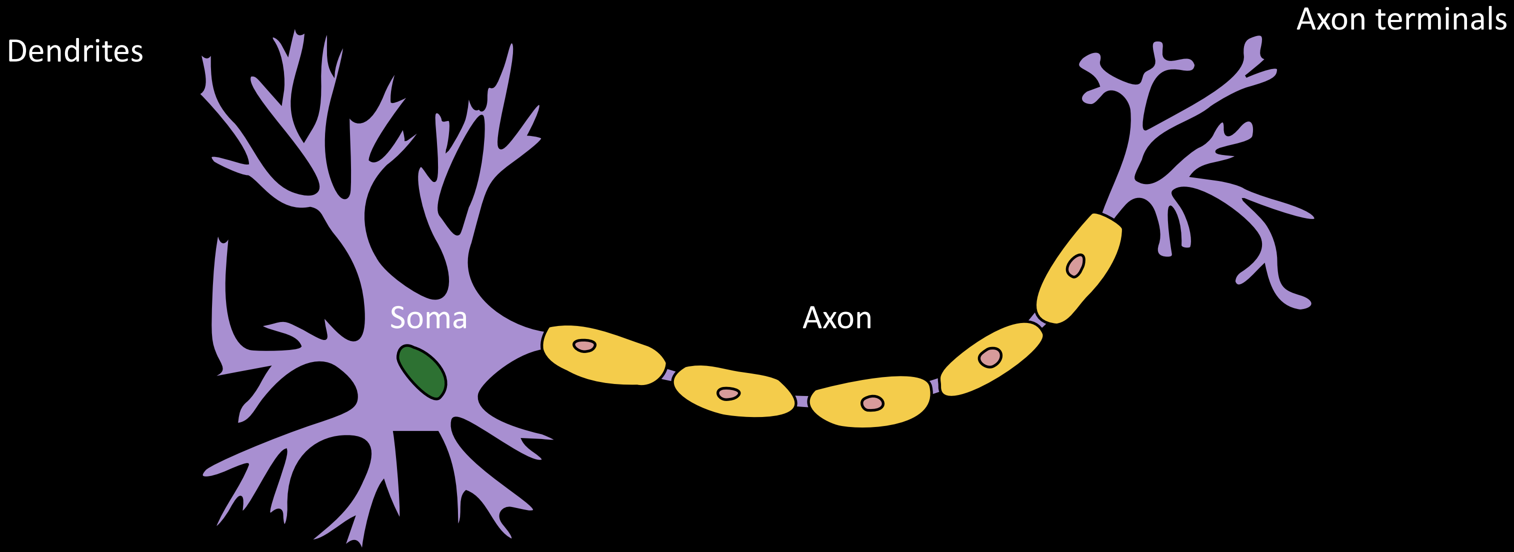

Why “neural” networks?

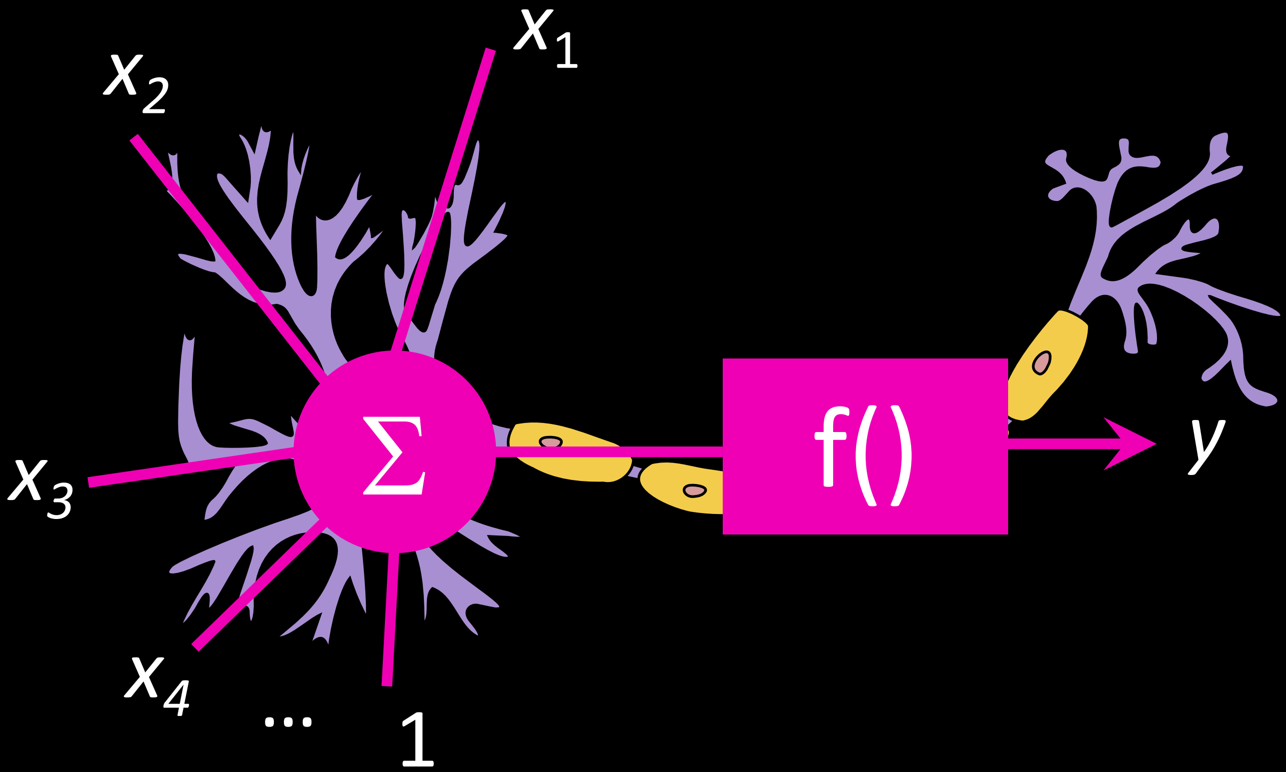

The perceptron

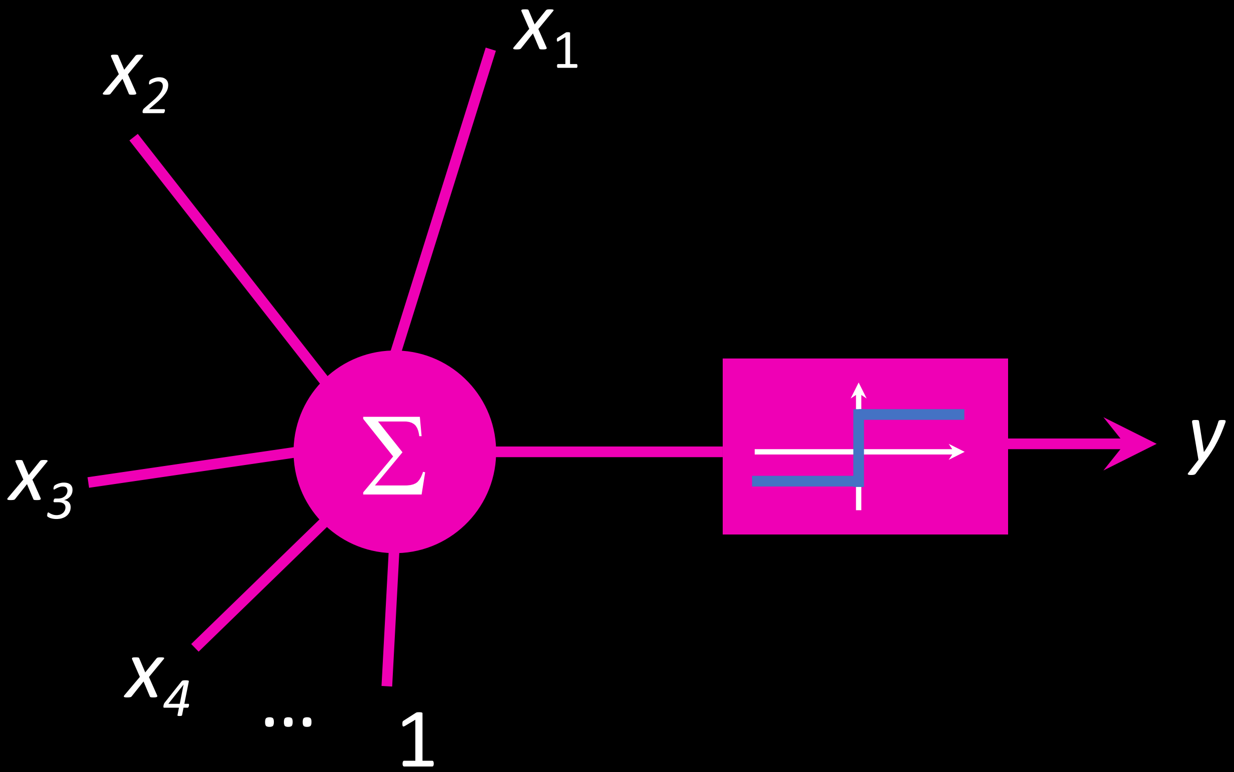

The perceptron

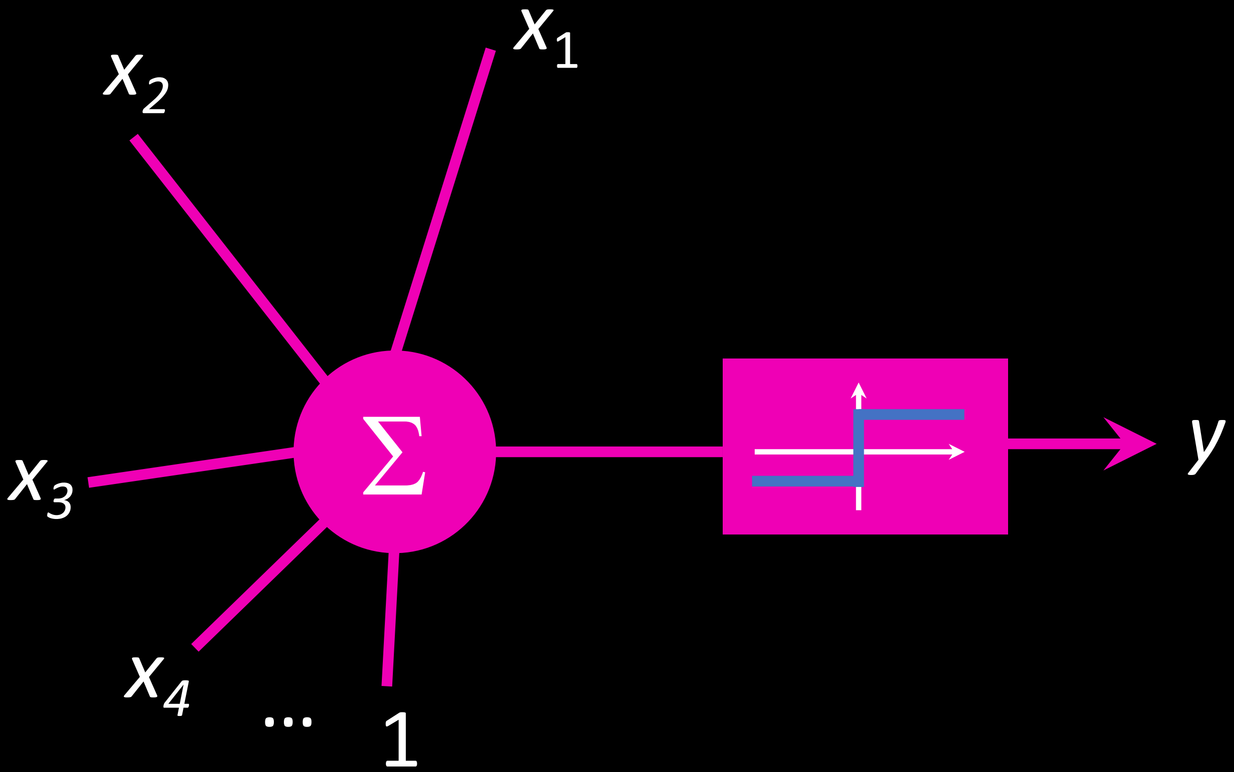

The perceptron

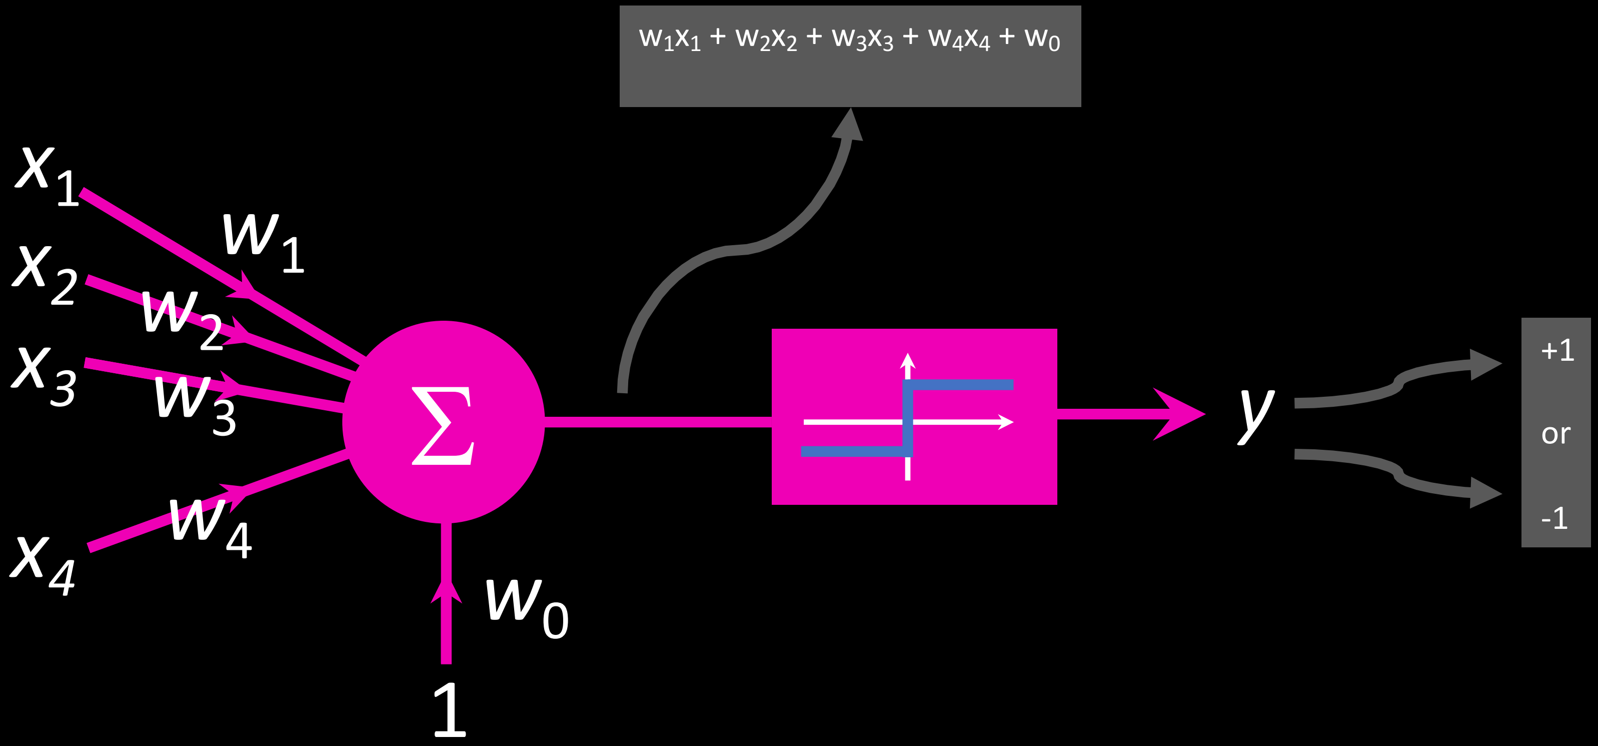

MARK I PERCEPTRON
PERCEPTRON

The perceptron rule

1. Find some training data
2. Set the weights to random values
3. Classify each pattern
4. For each pattern you got wrong

    Adjust the weights by adding a bit of pattern
5. Return to 3 until all correct.

Initial guess for the weight vector:

$$\mathbf{w} = [1.0, -0.8]^T$$

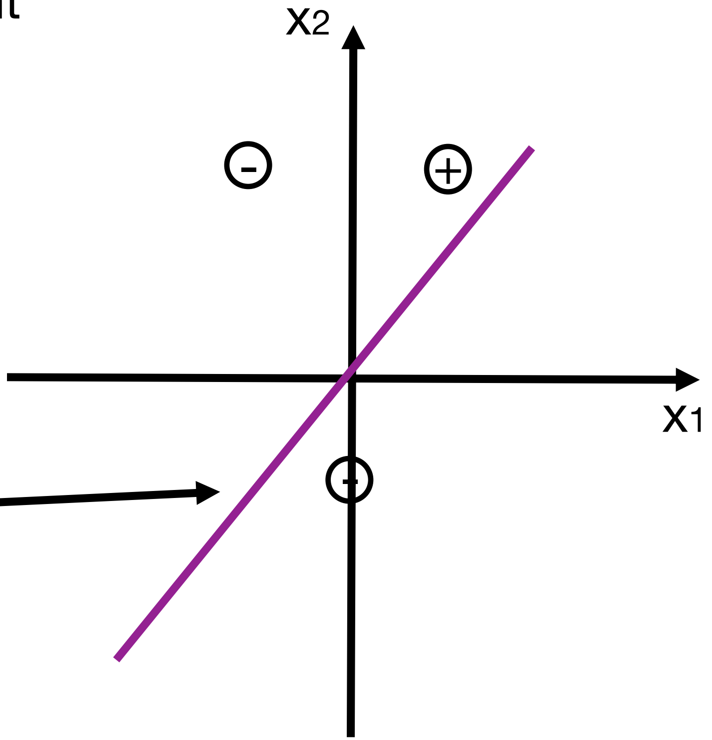

Decision boundary:

$$c(\mathbf{x}) = 0$$

$$w_1 x_1 + w_2 x_2 = 0$$

$$1.0 \times x_1 - 0.8 \times x_2 = 0$$

$$x_2 = \frac{1}{0.8} x_1$$

$$\mathbf{w} = \begin{bmatrix} 1.0 \\ -0.8 \end{bmatrix}$$

x_1	x_2	c	$w_0+\mathbf{w}^T\mathbf{x}$	c_{est}	
1	2	+1	-0.6	-1	✗
-1	2	-1	-2.6	-1	✓
0	-1	-1	0.8	+1	✗

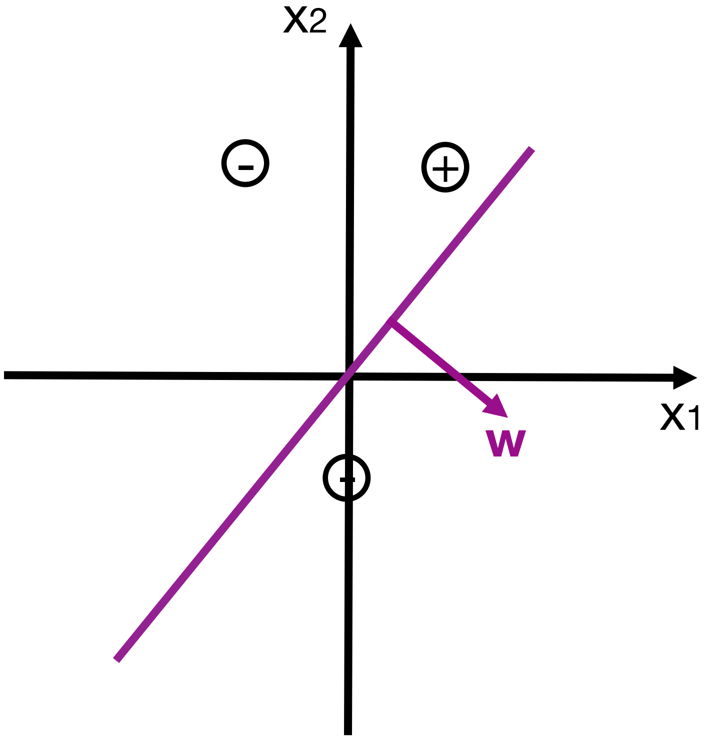

Take the first wrong point...

x_1	x_2	c	$w_0+\mathbf{w}^T\mathbf{x}$	c_{est}	
1	2	+1	-0.6	-1	✖

And use it to adjust **w**:

$\mathbf{w} = \mathbf{w} + \eta\, c(\mathbf{x}_1)\, \mathbf{x}_1$

$$\mathbf{w} = \begin{bmatrix} 1.0 \\ -0.8 \end{bmatrix} + \begin{bmatrix} 1 \\ 2 \end{bmatrix} \qquad \mathbf{w} = \begin{bmatrix} 2 \\ 1.2 \end{bmatrix}$$

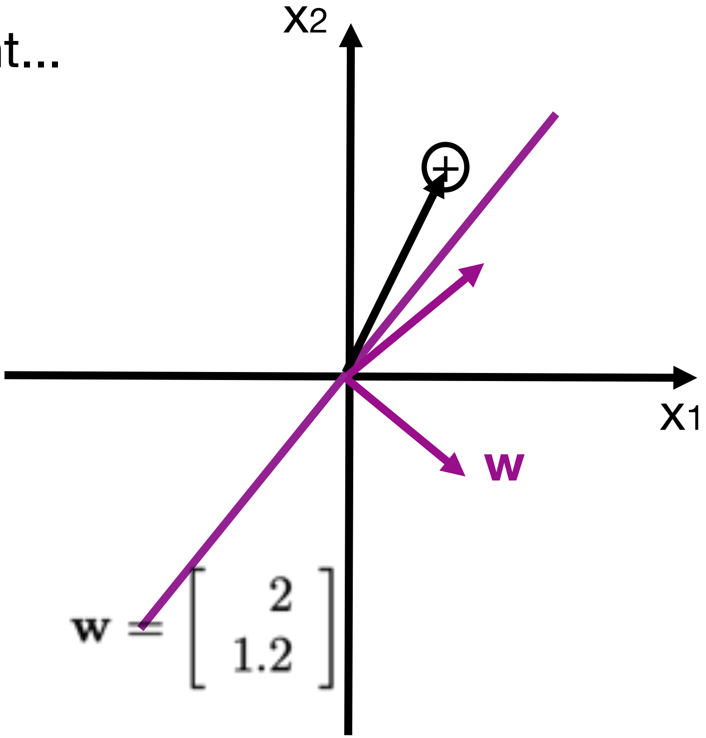

Here η=1 and ($c(x_1)$=+1.

$$\mathbf{w} = \begin{bmatrix} 2.0 \\ 1.2 \end{bmatrix}$$

x_1	x_2	c	$w_0+\mathbf{w}^T\mathbf{x}$	c_{est}	
1	2	+1	4.4	+1	✓
-1	2	-1	0.4	+1	✗
0	-1	-1	-1.2	-1	✓

$$\mathbf{w} = \begin{bmatrix} 2.0 \\ 1.2 \end{bmatrix} - \begin{bmatrix} -1 \\ 2 \end{bmatrix} = \begin{bmatrix} 3.0 \\ -0.8 \end{bmatrix}$$

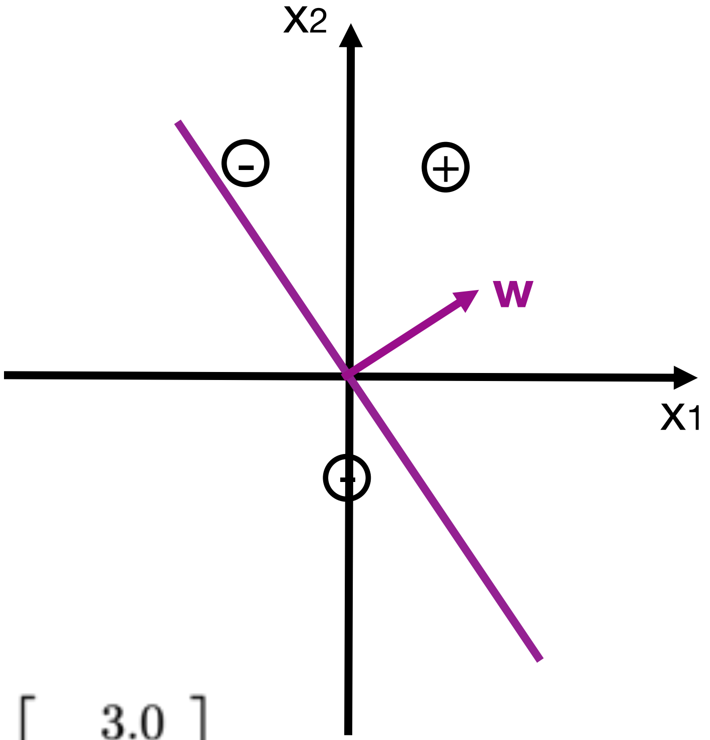

$$\mathbf{w} = \begin{bmatrix} 3.0 \\ -0.8 \end{bmatrix}$$

x_1	x_2	c	$w_0+\mathbf{w}^T\mathbf{x}$	c_{est}	
1	2	+1	1.4	+1	✓
-1	2	-1	-4.6	-1	✓
0	-1	-1	0.8	+1	✗

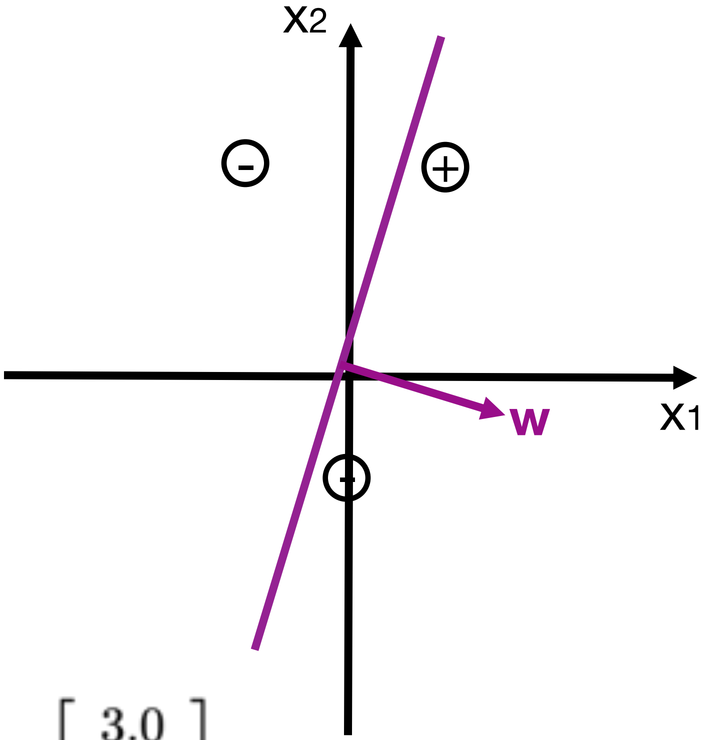

$$\mathbf{w} = \begin{bmatrix} 3.0 \\ -0.8 \end{bmatrix} - \begin{bmatrix} 0 \\ -1 \end{bmatrix} = \begin{bmatrix} 3.0 \\ 0.2 \end{bmatrix}$$

$$\mathbf{w} = \begin{bmatrix} 3.0 \\ 0.2 \end{bmatrix}$$

X_1	X_2	C	$W_0+\mathbf{w}^T\mathbf{x}$	C_{est}	
1	2	+1	3.4	+1	✓
-1	2	-1	-2.6	-1	✓
0	-1	-1	-0.2	-1	✓

No more points. Stop!

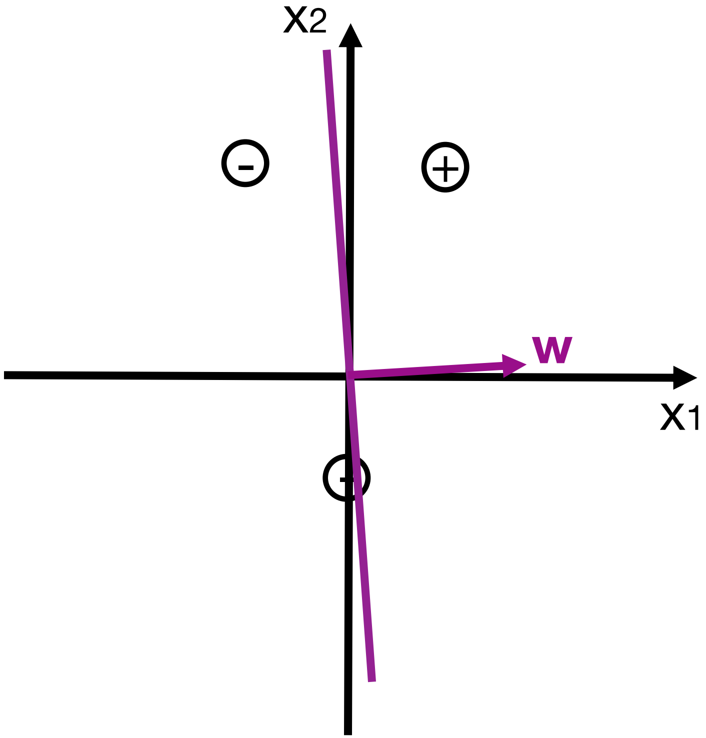

Perceptrons

- Can model logical AND and OR functions
- So provide a connection with logical thinking!

Perceptrons

- Can model logical AND and OR functions
- So provide a connection with logical thinking!

THIS WAS VERY EXCITING!

BUT…

Perceptrons

- Cannot do some simple things
- Percepetron rule does not allow max margin

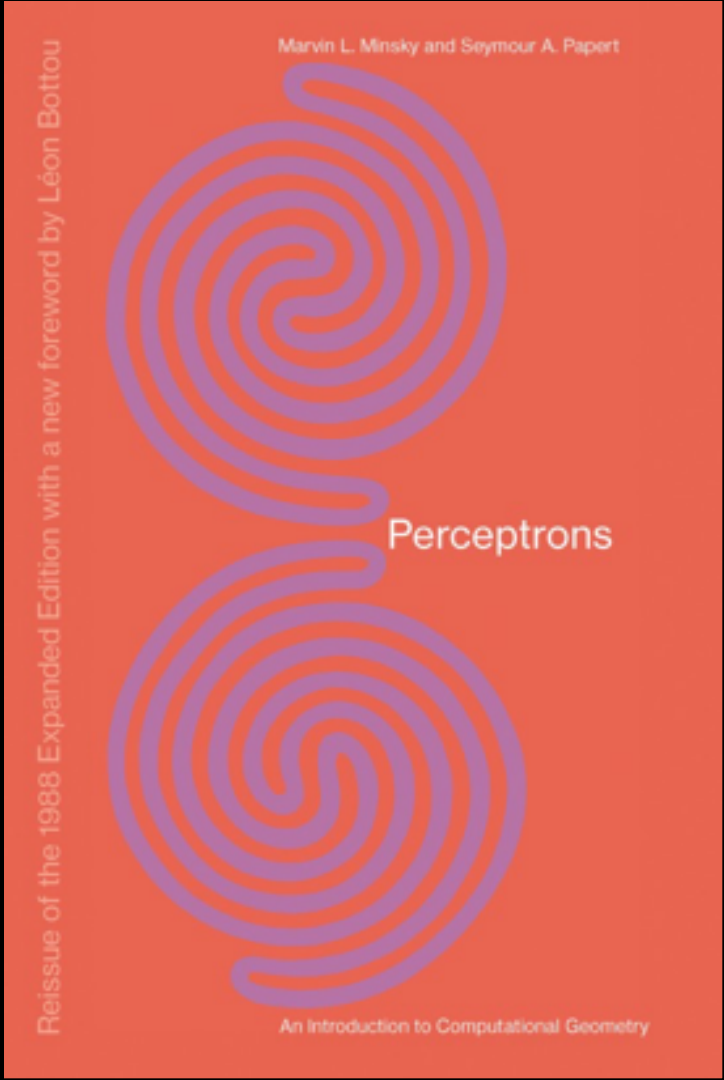

AND THAT WAS THE END OF AI

FOR A WHILE…

The perceptron (version 2)

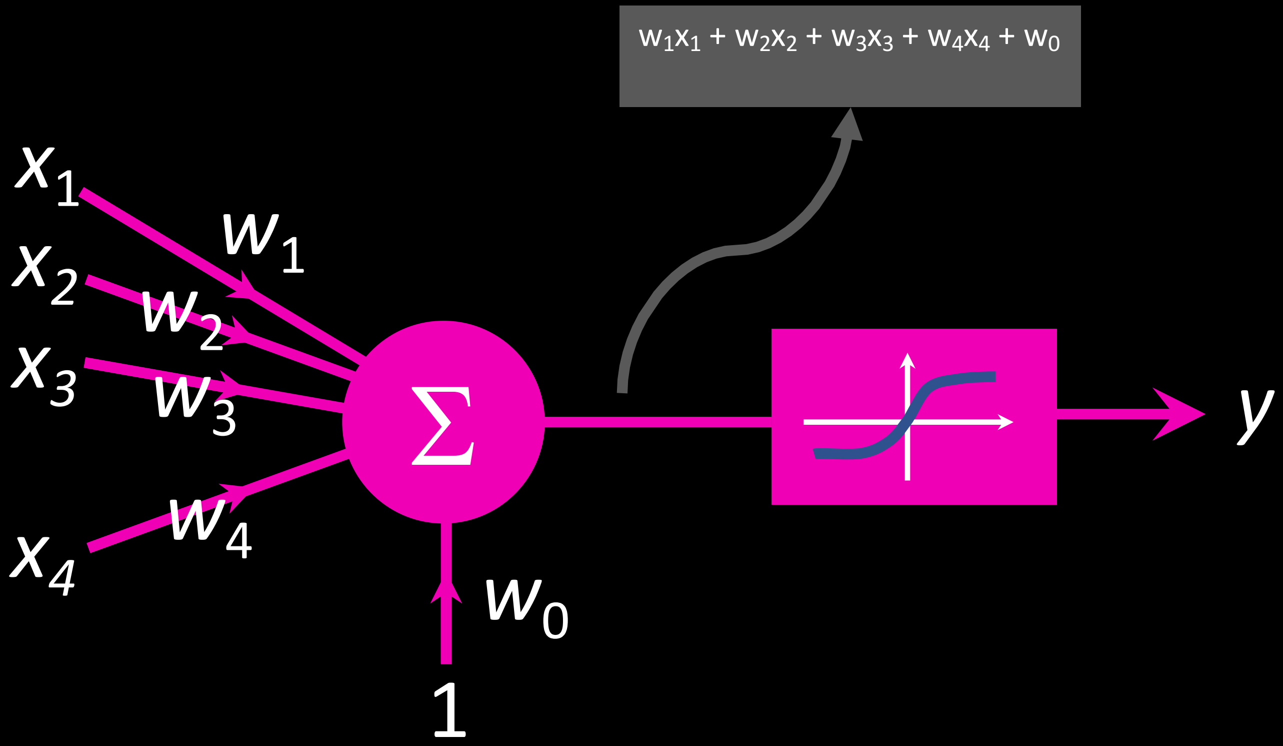

The multi-layer perceptron (MLP)

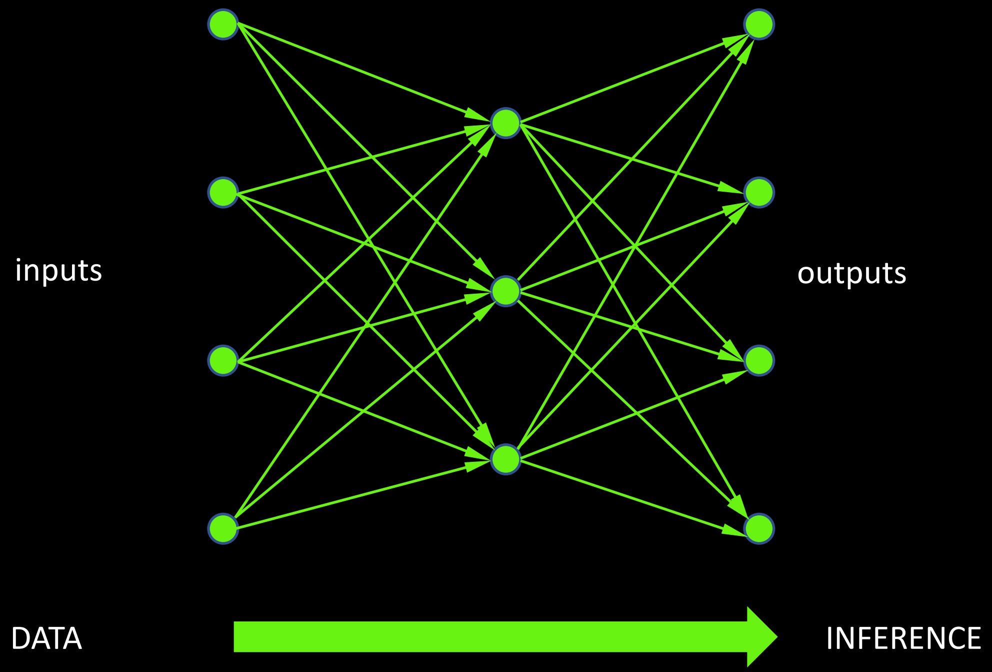

MLPs

- Can model logical functions
- Curved decision boundaries
- Trainable via backprop

THIS WAS VERY EXCITING!

BUT…

Multilayer perceptrons

- Were rather unprincipled
- Not as effective as other classifiers which had probabilistic interpretations

- Seemed to be difficult to train.

AND THAT WAS THE END OF THEM

FOR A WHILE…

The deep network

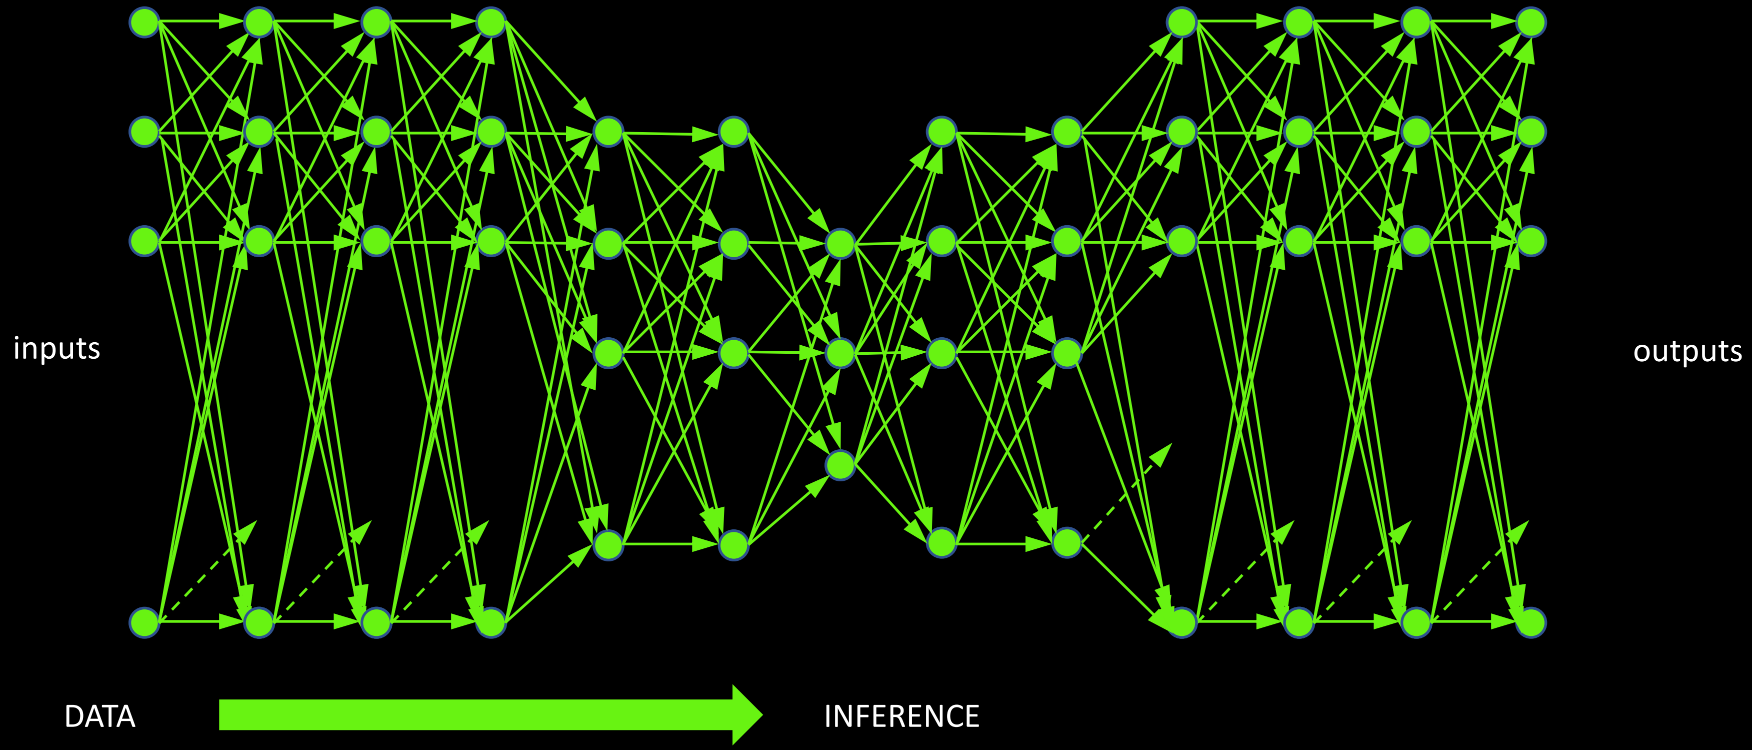

The reality

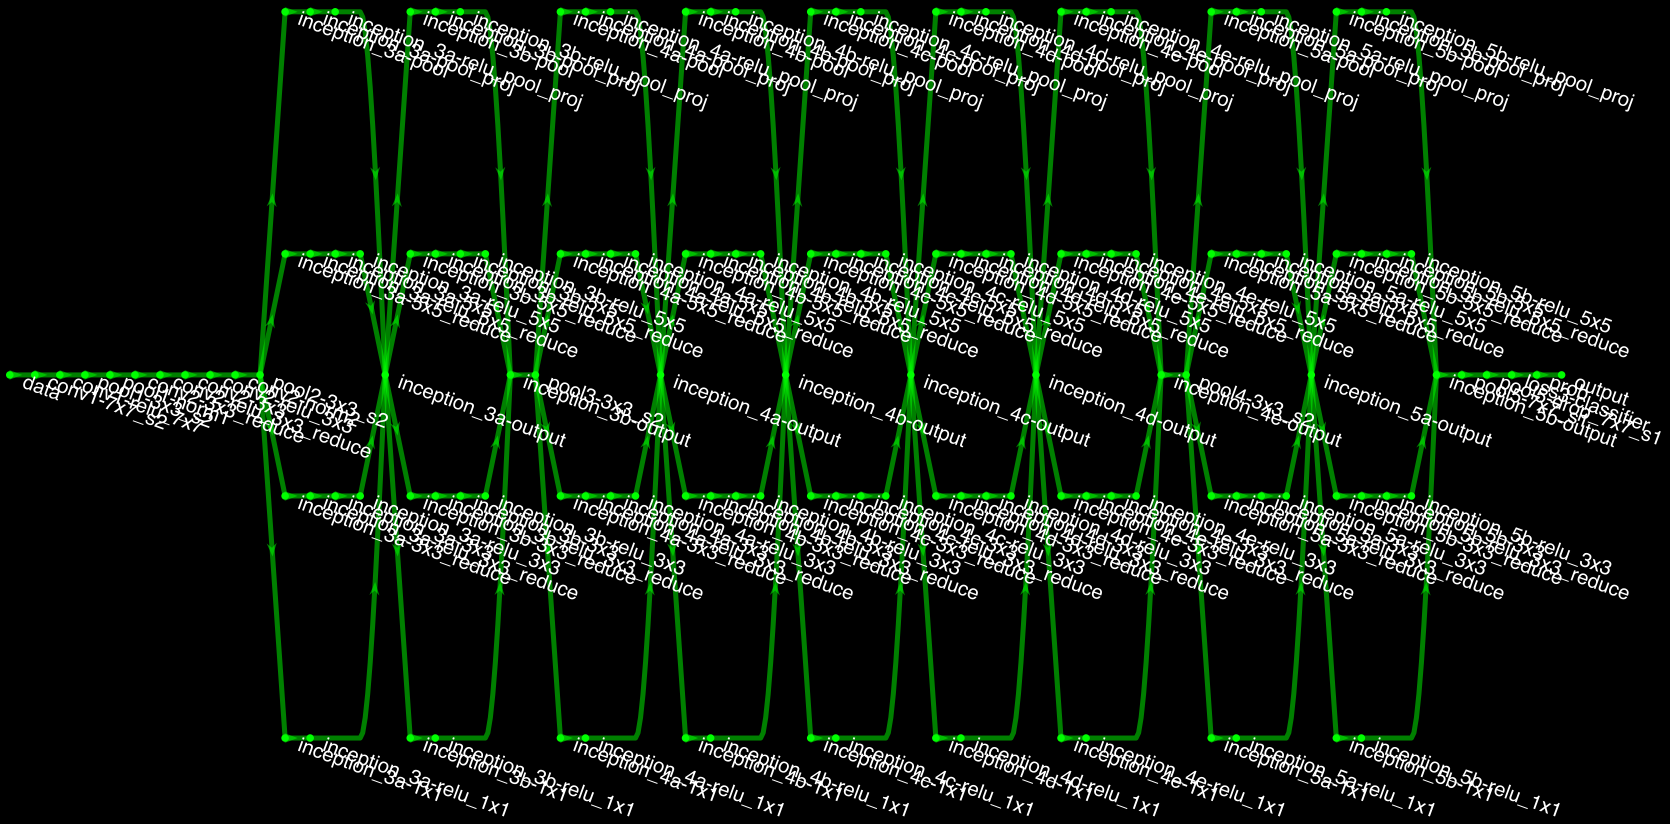
data
inception_3a-output
inception_4a-output
inception_4b-output
inception_4c-output
inception_4d-output
inception_5a-output
inception_5b-relu_pool_proj
inception_3a-relu_1x1
inception_3b-relu_1x1
inception_4a-relu_1x1
inception_4b-relu_1x1
inception_4c-relu_1x1
inception_4d-relu_1x1
inception_4e-relu_1x1
inception_5a-relu_1x1
inception_5b-relu_1x1
output

GoogLeNet

- 1.2M training images
- 1000 classes

tench goldfish great white shark tiger shark hammerhead electric ray stingray cock hen ostrich brambling goldfinch house finch junco indigo bunting robin bulbul jay magpie chickadee water ouzel kite bald eagle vulture great grey owl European fire salamander common newt eft spotted salamander axolotl bullfrog tree frog tailed frog loggerhead leatherback turtle mud turtle terrapin box turtle banded gecko common iguana American chameleon whiptail agama frilled lizard alligator lizard Gila monster green lizard African chameleon Komodo dragon African crocodile American alligator triceratops thunder snake ringneck snake hognose snake green snake king snake garter snake water snake vine snake night snake boa constrictor rock python Indian cobra green mamba sea snake horned viper diamondback sidewinder trilobite harvestman scorpion black and gold garden spider barn spider garden spider black widow tarantula wolf spider tick centipede black grouse ptarmigan ruffed grouse prairie chicken peacock quail partridge African grey macaw sulphur-crested cockatoo lorikeet coucal bee eater hornbill hummingbird jacamar toucan drake red-breasted merganser goose black swan tusker echidna platypus wallaby koala wombat jellyfish sea anemone brain coral flatworm nematode conch snail slug sea slug chiton chambered nautilus Dungeness crab rock crab fiddler crab king crab American lobster spiny lobster crayfish hermit crab isopod white stork black stork spoonbill flamingo little blue heron American egret bittern crane limpkin European gallinule American coot bustard ruddy turnstone red-backed sandpiper redshank dowitcher oystercatcher pelican king penguin albatross grey whale killer whale dugong sea lion Chihuahua Japanese spaniel Maltese dog Pekinese Shih-Tzu Blenheim spaniel papillon toy terrier Rhodesian ridgeback Afghan hound basset beagle bloodhound bluetick black-and-tan coonhound Walker hound English foxhound redbone borzoi Irish wolfhound Italian greyhound whippet Ibizan hound Norwegian elkhound otterhound Saluki Scottish deerhound Weimaraner Staffordshire bullterrier American Staffordshire terrier Bedlington terrier Border terrier Kerry blue terrier Irish terrier Norfolk terrier Norwich terrier Yorkshire terrier wire-haired fox terrier Lakeland terrier Sealyham terrier Airedale cairn Australian terrier Dandie Dinmont Boston bull miniature schnauzer giant schnauzer standard schnauzer Scotch terrier Tibetan terrier silky terrier soft-coated wheaten terrier West Highland white terrier Lhasa flat-coated retriever curly-coated retriever golden retriever Labrador retriever Chesapeake Bay retriever German short-haired pointer vizsla English setter Irish setter Gordon setter Brittany spaniel clumber English springer Welsh springer spaniel cocker spaniel Sussex spaniel Irish water spaniel kuvasz schipperke groenendael malinois briard kelpie komondor Old English sheepdog Shetland sheepdog collie Border collie Bouvier des Flandres Rottweiler German shepherd Doberman miniature pinscher Greater Swiss Mountain dog Bernese mountain dog Appenzeller EntleBucher boxer bull mastiff Tibetan mastiff French bulldog Great Dane Saint Bernard Eskimo dog malamute Siberian husky dalmatian affenpinscher basenji pug Leonberg Newfoundland Great Pyrenees Samoyed Pomeranian chow keeshond Brabancon griffon Pembroke Cardigan toy poodle miniature poodle standard poodle Mexican hairless timber wolf white wolf red wolf coyote dingo dhole African hunting dog hyena red fox kit fox Arctic fox grey fox tabby tiger cat Persian cat Siamese cat Egyptian cat cougar lynx leopard snow leopard jaguar lion tiger cheetah brown bear American black bear ice bear sloth bear mongoose meerkat tiger beetle ladybug ground beetle long-horned beetle leaf beetle dung beetle rhinoceros beetle weevil fly bee ant grasshopper cricket walking stick cockroach mantis cicada leafhopper lacewing dragonfly damselfly admiral ringlet monarch cabbage butterfly sulphur butterfly lycaenid starfish sea urchin sea cucumber wood rabbit hare Angora hamster porcupine fox squirrel marmot beaver guinea pig sorrel zebra hog wild boar warthog hippopotamus ox water buffalo bison ram bighorn ibex hartebeest impala gazelle Arabian camel llama weasel mink polecat black-footed ferret otter skunk badger armadillo three-toed sloth orangutan gorilla chimpanzee gibbon siamang guenon patas baboon macaque langur colobus proboscis monkey marmoset capuchin howler monkey titi spider monkey squirrel monkey Madagascar cat indri Indian elephant African elephant lesser panda giant panda barracouta eel coho rock beauty anemone fish sturgeon gar lionfish puffer abacus abaya academic gown accordion acoustic guitar aircraft carrier airliner airship altar ambulance amphibian analog clock apiary apron ashcan assault rifle backpack bakery balance beam balloon ballpoint Band Aid banjo bannister barbell barber chair barbershop barn barometer barrel barrow baseball basketball bassinet bassoon bathing cap bath towel bathtub beach wagon beacon beaker bearskin beer bottle beer glass bell cote bib bicycle-built-for-two bikini binder binoculars birdhouse boathouse bobsled bolo tie bonnet bookcase bookshop bottlecap bow bow tie brass brassiere breakwater breastplate broom bucket buckle bulletproof vest bullet train butcher shop cab caldron candle cannon canoe can opener cardigan car mirror carousel carpenter's kit carton car wheel cash machine cassette cassette player castle catamaran CD player cello cellular telephone chain chainlink fence chain mail chain saw chest chiffonier chime china cabinet Christmas stocking church cinema cleaver cliff dwelling cloak clog cocktail shaker coffee mug coffeepot coil combination lock computer keyboard confectionery container ship convertible corkscrew cornet cowboy boot cowboy hat cradle crane (machine) crash helmet crate crib Crock Pot croquet ball crutch cuirass dam desk desktop computer dial telephone diaper digital clock digital watch dining table dishrag dishwasher disk brake dock dogsled dome doormat drilling platform drum drumstick dumbbell Dutch oven electric fan electric guitar electric locomotive entertainment center envelope espresso maker face powder feather boa file fireboat fire engine fire screen flagpole flute folding chair football helmet forklift fountain fountain pen four-poster freight car French horn frying pan fur coat garbage truck gasmask gas pump goblet go-kart golf ball golfcart gondola gong gown grand piano greenhouse grille grocery store guillotine hair slide hair spray half track hammer hamper hand blower hand-held computer handkerchief hard disc harmonica harp harvester hatchet holster home theater honeycomb hook hoopskirt horizontal bar horse cart hourglass iPod iron jack-o'-lantern jean jeep jersey jigsaw puzzle jinrikisha joystick kimono knee pad knot lab coat ladle lampshade laptop lawn mower lens cap letter opener library lifeboat lighter limousine liner lipstick Loafer lotion loudspeaker loupe lumbermill magnetic compass mailbag mailbox maillot maillot, tank suit manhole cover maraca marimba mask matchstick maypole maze measuring cup medicine chest megalith microphone microwave military uniform milk can minibus miniskirt minivan missile mitten mixing bowl mobile home Model T modem monastery monitor moped mortar mortarboard mosque mosquito net motor scooter mountain bike mountain tent mouse mousetrap moving van muzzle nail neck brace necklace nipple notebook obelisk oboe ocarina odometer oil filter organ oscilloscope overskirt oxcart oxygen mask packet paddle paddlewheel padlock paintbrush pajama palace panpipe paper towel parachute parallel bars park bench parking meter passenger car patio pay-phone pedestal pencil box pencil sharpener perfume Petri dish photocopier pick pickelhaube picket fence pickup pier piggy bank pill bottle pillow ping-pong ball pinwheel pirate pitcher plane planetarium plastic bag plate rack plow plunger Polaroid camera pole police van poncho pool table pop bottle pot potter's wheel power drill prayer rug printer prison projectile projector puck punching bag purse quill quilt racer racket radiator radio radio telescope rain barrel recreational vehicle reel reflex camera refrigerator remote control restaurant revolver rifle rocking chair rotisserie rubber eraser rugby ball rule running shoe safe safety pin saltshaker sandal sarong sax scabbard scale school bus schooner scoreboard screen screw screwdriver seat belt sewing machine shield shoe shop shoji shopping basket shopping cart shovel shower cap shower curtain ski ski mask sleeping bag slide rule sliding door slot snorkel snowmobile snowplow soap dispenser soccer ball sock solar dish sombrero soup bowl space bar space heater space shuttle spatula speedboat spider web spindle sports car spotlight stage steam locomotive steel arch bridge steel drum stethoscope stole stone wall stopwatch stove strainer streetcar stretcher studio couch stupa submarine suit sundial sunglass sunglasses sunscreen suspension bridge swab sweatshirt swimming trunks swing switch syringe table lamp tank tape player teapot teddy television tennis ball thatch theater curtain thimble thresher throne tile roof toaster tobacco shop toilet seat torch totem pole tow truck toyshop tractor trailer truck tray trench coat tricycle trimaran tripod triumphal arch trolleybus trombone tub turnstile typewriter keyboard umbrella unicycle upright vacuum vase vault velvet vending machine vestment viaduct violin volleyball waffle iron wall clock wallet wardrobe warplane washbasin washer water bottle water jug water tower whiskey jug whistle wig window screen window shade Windsor tie wine bottle wing wok wooden spoon wool worm fence wreck yawl yurt web site comic book crossword puzzle street sign traffic light book jacket menu plate guacamole consomme hot pot trifle ice

GoogLeNet

- 1.2M training images
- 1000 classes
- 22 layers deep
- Training time – not declared (several weeks probably)
- Run-time 0.25s on my laptop

The reality

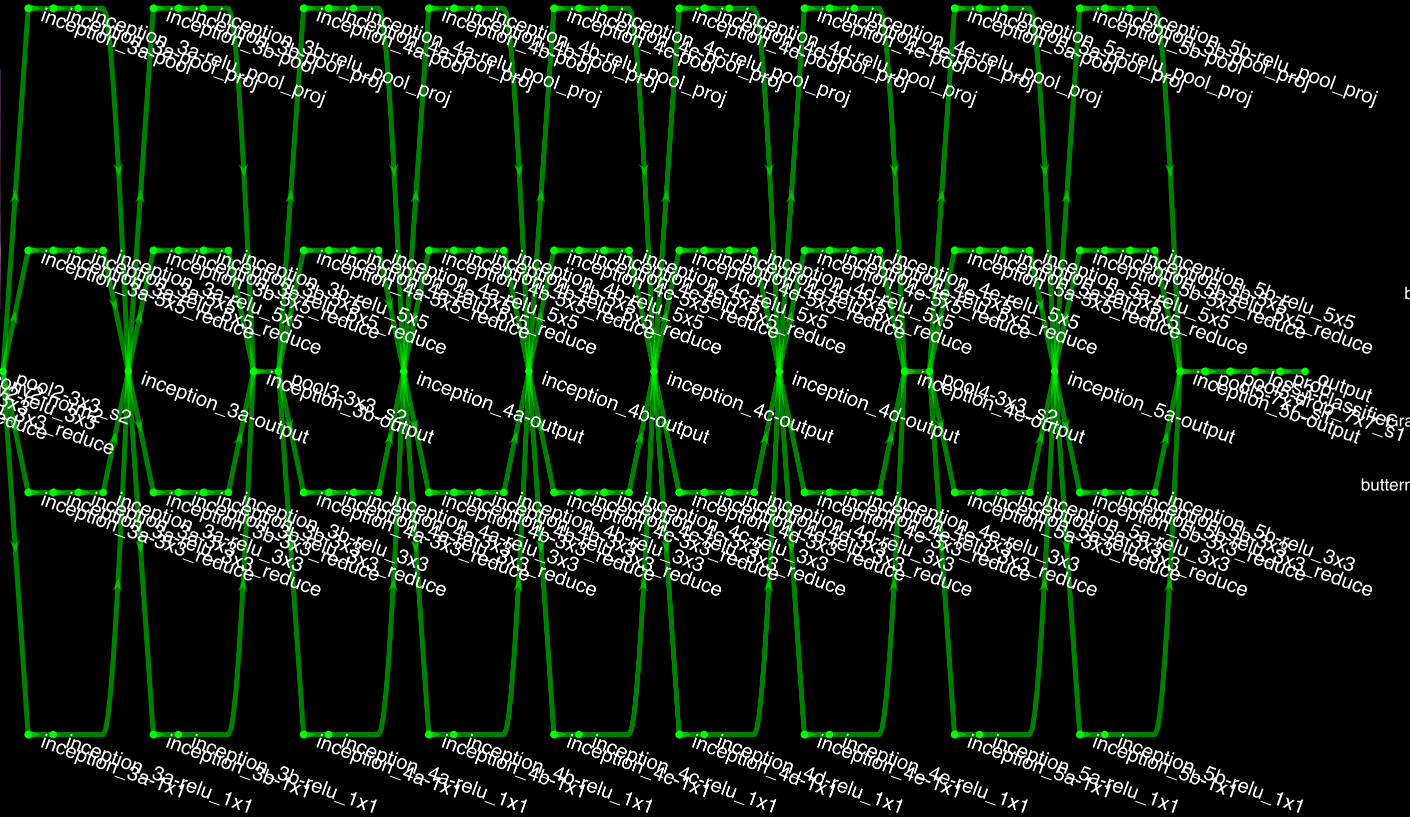

inception_3a-output
inception_3b-output
inception_4a-output
inception_4b-output
inception_4c-output
inception_4d-output
inception_4e-output
inception_5a-output
inception_5b-output
inception_5b-relu_pool_proj
inception_5b-relu_5x5
inception_5b-relu_3x3
inception_5b-relu_1x1

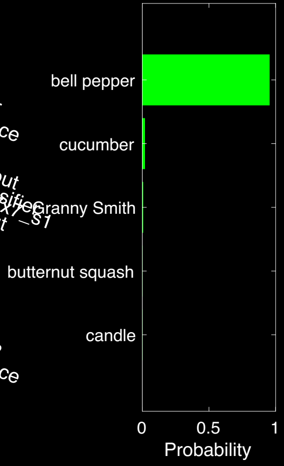

bell pepper
cucumber
Granny Smith
butternut squash
candle
0
0.5
1
Probability

The reality

BLACK BOX

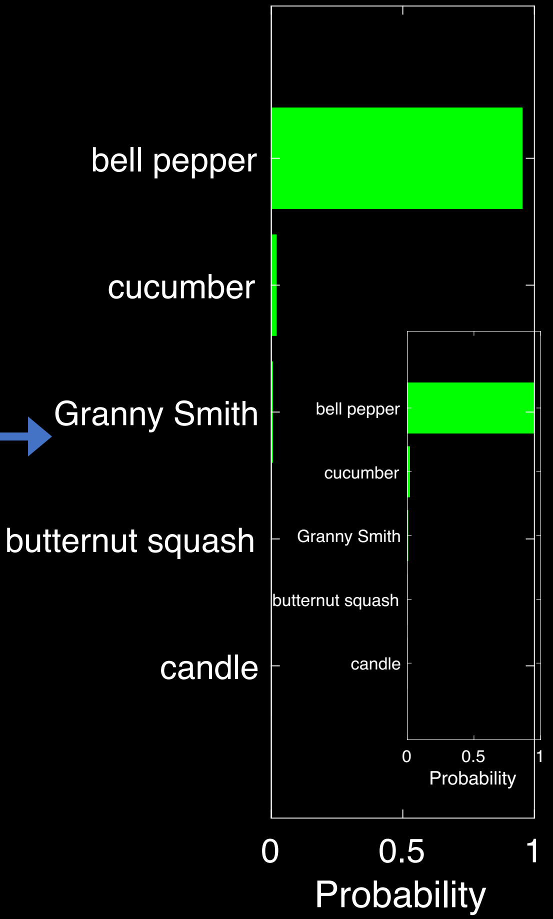

Deep neural networks

- Can model very complex functions
- Have a rich set of architectures (many not mentioned here)
- Can be “transferred” from one domain to another
- Are fantastically successful
- Can be implemented by a high school student.

Google AlphaGo vs Li Sedol

THIS IS VERY EXCITING!

BUT…

Fairness in machine learning

- Joy Boulamwini, MIT
- IBM Fairness 360 toolkit

Deep learning

- Has progressed so quickly that various people worry about ethical issues
- Is “black box”
- Is often applied without adequate testing

Debate:

There is no such thing as AI Ethics: just ethics
Richard Harvey versus Chris Rees
Worshipful Company of IT, 18:00, 2nd April 2019

Next lectures:

text (16th April)
creativity (28th May)