u3a Computing Group

Alan Hopwood, 7 September 2023

Agenda

Welcome
Current News, Issues and Questions
Future Topics & Next Meeting
Topic: Artificial Intelligence
AOB and Follow up

Current News, Issues and Questions

Anything to discuss?

Future Topics

Topic	Votes
Streaming - the best way to go	4
Computer languages - functionality vs ease of use	4
Laptop vs Tablet	4
Microchip design	2
Being safe on the internet	2
History of computer development	2
Bluetooth	1
Digital communications / information encoding	1
Chromebook	1
Mac vs Windows vs Linux	1
Printing from an Android tablet	1

Next Month?

Volunteers or Cancel?

Presentation
Artificial Intelligence

Sources

Artificial Intelligence

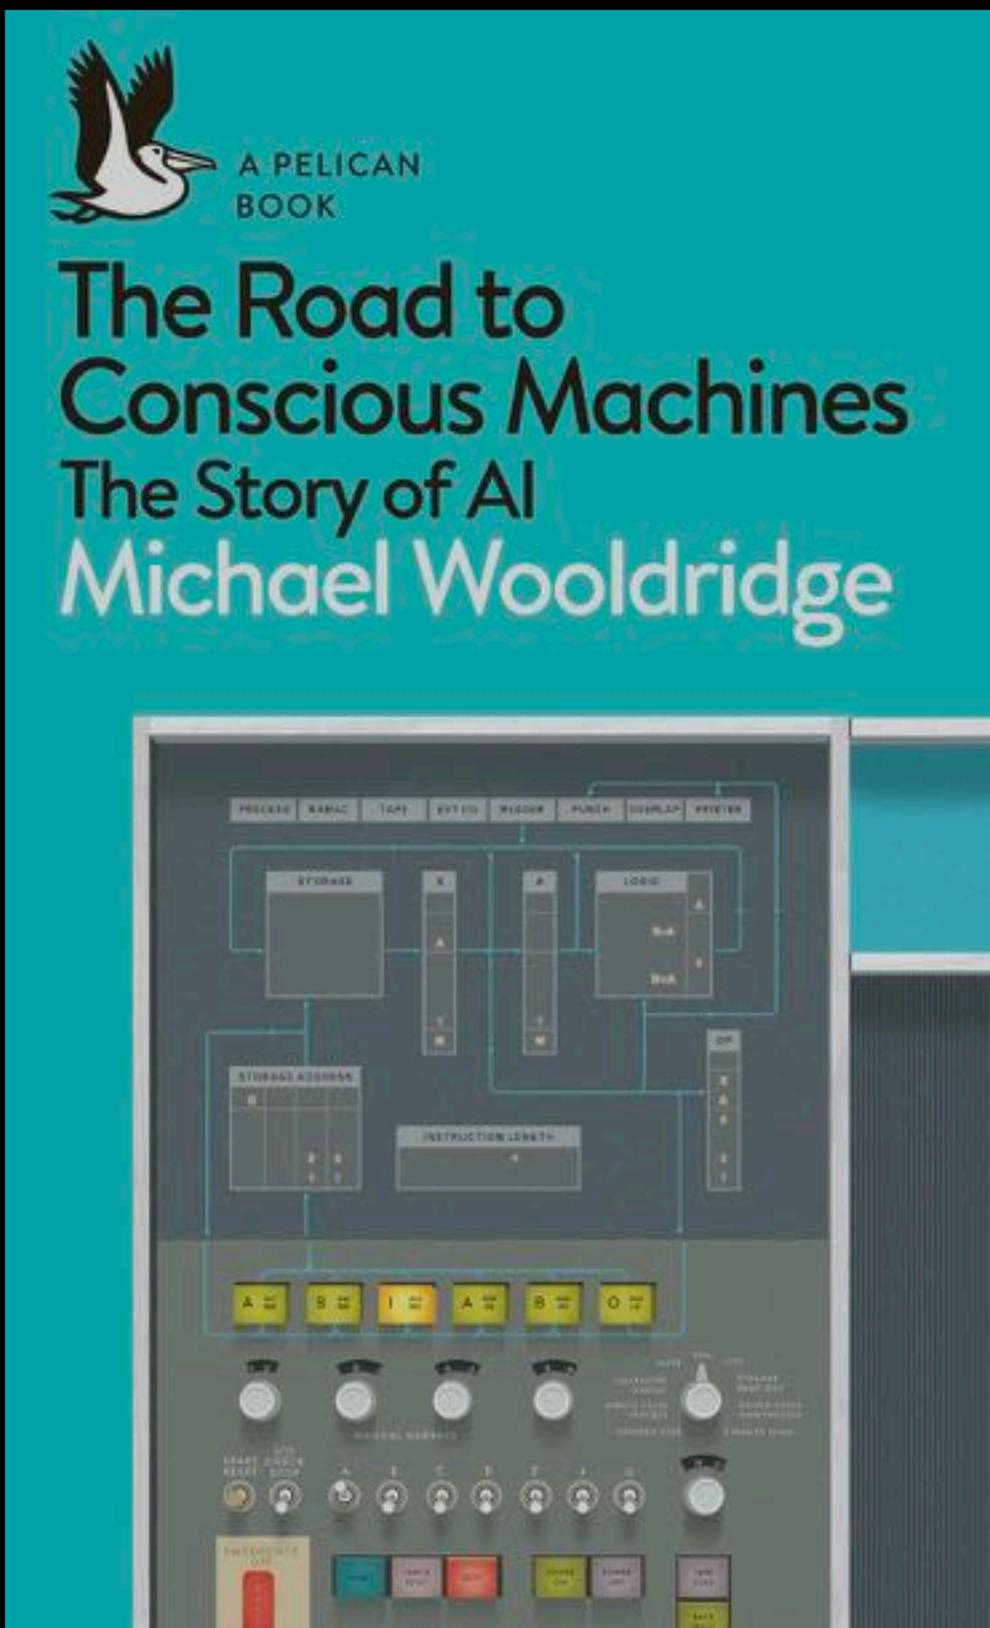

This presentation is based largely on:

- The book by Michael Wooldridge, professor of computer science at the University of Oxford (thanks John)
- Wikipedia
- Deep Learning ref:

```
@book{Goodfellow-et-al-2016,
    title={Deep Learning},
    author={Ian Goodfellow and Yoshua Bengio
and Aaron Courville},
    publisher={MIT Press},
    note={\url{http://
www.deeplearningbook.org}},
    year={2016}}
```

- NVIDIA website

If you want to read more about AI, I suggest you start with the Wooldbridge book.

Agenda

Artificial Intelligence

A Journey through AI

- What is Artificial Intelligence (what do we mean?)
- How does AI work - the problems and approaches
- How AI developed
  - Symbolic AI
  - Neural networks
- Machine Learning
- Large Language Models and Generative AI
- AI Limitations - Reality

Will start slowly/gently and ramp up.

May be too much for one meeting.

Check point at ~ 16:30

What is AI?

Artificial Intelligence

- What is artificial intelligence?

  The ability of machines to perform tasks that are typically associated with human intelligence, such as learning and problem-solving

- Basically, we measure AI against the capability of people

Intelligence tasks ranked by difficulty

Artificial Intelligence

Where we are

Easy

- Arithmetic (1945)
- Sorting lists of numbers (1959)

Solved after a lot of effort

- Playing simple board games (1959)
- Playing chess (1997)
- Recognising faces in pictures (2008)
- Usable automated translation (2010)
- Playing Go (2016)
- Usable real-time translation of spoken words (2016)

Real progress

- Driverless cars
- Automatically providing captions for pictures

Nowhere near solved

- Understanding a story & answer questions about it.
- Human -level automated translation
- Interpreting what is going on in a photograph
- Writing interesting stories
- Interpreting a work of art
- Human-level general intelligence

The Turing Test

Artificial Intelligence

- Based on a Victorian Parlour game - the Imitation Game Objective was to try to tell if a another person was a man or woman on the basis of answers they gave to questions posed to them

- Turing's version is: Is this human or computer with interaction via computer and screen.

- If the interrogators cannot tell if they are interacting with person or program, then the program must have some intelligence.

[Alan Turing played a huge part in the development of computing - worthy of a presentation in itself]

Consider the Distinction

Artificial Intelligence

In a dialogue like the Turing test:

- The program actually understands the dialogue, in much the same sense that a person does.

vs.

- The program does not have any understanding, but can simulate such understanding.

Capabilities Sought in AI development

Artificial Intelligence

- Perception
  - Building and using sensors to provide the analogue of our 5 senses: sight, hearing, touch, smell and taste.
- Machine Learning
  - Learning from and making predictions about data.
  - E.g. recognising faces in pictures
- Problem Solving and Planning
  - Achieve a goal using using a repertoire of actions
  - E.g. playing Chess or Go
- Natural Language Understanding

Compare What Computers do?

Artificial Intelligence

This is the basic building block

- Only follow instructions!
- Very basic instructions
  e.g. add A to B; if result is bigger than C, then do D
- But very fast
  PC: 100 billion instructions per second
- Compare manually at rate of one instruction for every 10 seconds would take 3,700 years (24 hours per day)to do what pc would in 1 second
- Without mistakes!

(This is quite different to Human characteristics)

Problem Solving

Artificial Intelligence

- Covers development of algorithms to solve problems or make logical deductions.
- Many mathematical problems are easy
  - Is 2+2=4? is 4x4=16?
- Some are more of a challenge but still "decidable".
  - is 7919 a prime number?
  - This requires checking every number that could possibly a divisor to see whether it divides 7919 exactly
- Other algorithms based on probability and economics programs to deal with uncertainty in problems

All "easily achievable, but not generally considered AI

Problem Solving, Planning

Artificial Intelligence

- AI progressed into game playing, e.g. chess & Go
- The fundamental approach taken is to **SEARCH** through every possible move to achieve the goal
- E.g. the Towers of Hanoi puzzle: (https://www.mathsisfun.com/games/towerofhanoi.html

A

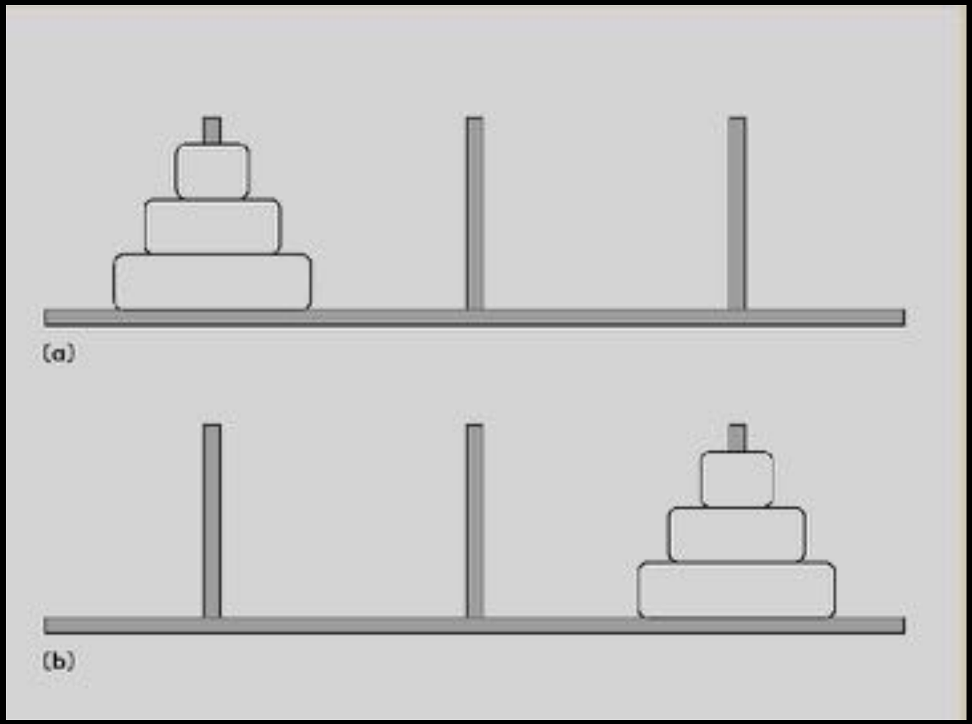

B

Challenge is to find the smallest number of moves to get the rings from A to B one by one without having a larger ring being on top of a smaller one

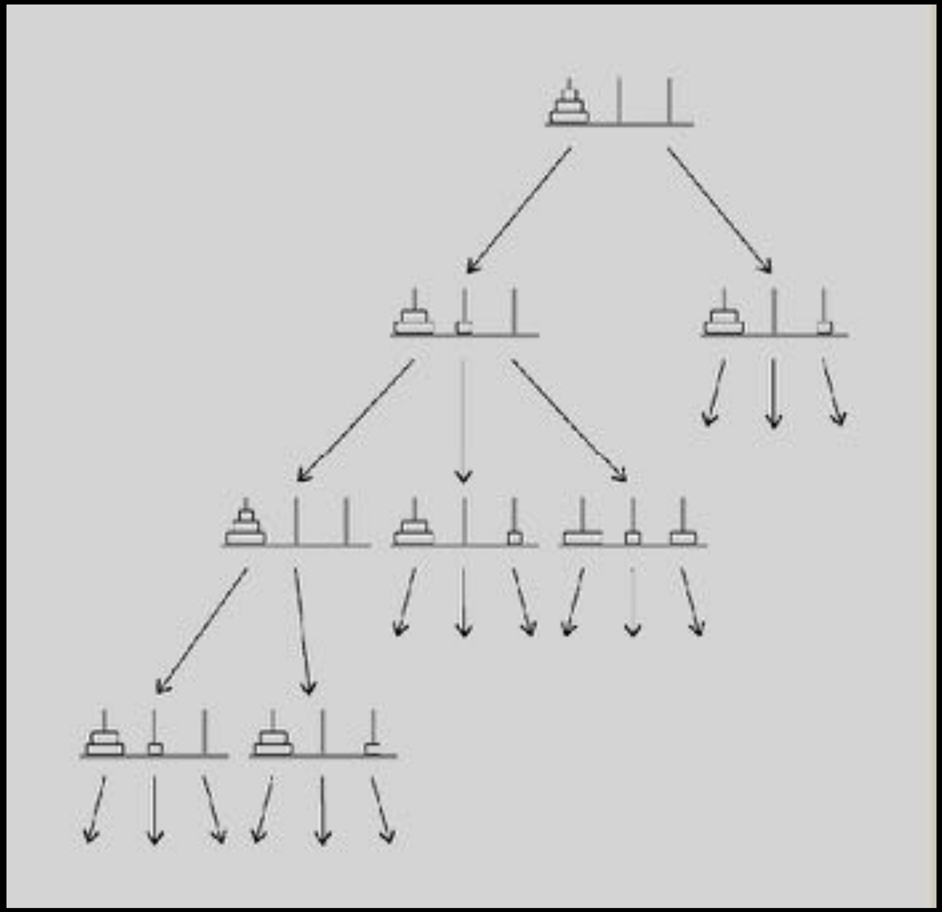

Search tree to solve the problem

AI Technique - Search

Artificial Intelligence

- **Search** is a fundamental problem solving technique which involves systematically considering all possible courses of action.
- Any program that plays a game, E.g. Chess is based on Search, as is sat nav.
- Basically the Search approach is:
    - From initial state, consider every available action on the initial state
    - Each action transforms problem to a new state
    - If a new state is the goal state, then solution is found
    - Otherwise repeat the process for every new state just generated.
- The challenge of Search is Combinatorial Explosion

Combinatorial Explosion

Artificial Intelligence

The size of Search tree depends on branching factor and number of levels (moves)

- The Branching Factor is the number of possible moves available each step.
- For the Tower of Hanoi, the Branching Factor is 3
- For Go, the Branching Factor is 250
- The Go search tree has 250 states at the 1st level, 62,500 at the 2nd level, 15.6 million at the 3rd, 3.9 billion for the 4th
- A typical game of Go lasts for about 20 moves

The Search Tree is rather large and the basic approach is to have the full model in memory

Beating Combinatorial Explosion

Artificial Intelligence

- **Depth first search**: expand a branch until a solution is reached or there is a dead end:
  - Don’t have to store the whole of the search tree
  - But may get to end of branch without finding a solution.
- **Heuristic search**: use “rules of thumb” (heuristics) to indicate where to focus search.
  - The heuristics are specific to the problem
  - May, for e.g., involve having an overall measure of quality of a board position (chess, drafts, Go) built up on a set of factors.

Of course there is also an opponent to take account of.

- **Minimax** search also tests against all possible next moves of the opponent assuming they would make the worst move for you.

Knowledge Representation

Artificial Intelligence

- Two models of AI treatment of knowledge:
  - Approach 1:
    Symbolic, where symbols represent concepts
  - Approach 2:
    Neural Nets, captures knowledge as a network of perceptrons (a simplified model of a biological neuron)
- Also thought of in terms of modelling the mind vs. modelling the brain

Approache1: Symbolic AI

Artificial Intelligence

- Called symbolic AI, because it makes use of symbols that stand for things that the system is reasoning about.
- For example:
  - symbol ‘room451’ within a robot’s control system might be the name that the robot uses for your bedroom
  - symbol ‘cleanRoom’ might be used as the name for the activity of cleaning a room.
- Key advantage (compared to Neural Nets) is transparency: when the robot concludes that it will ‘cleanRoom(room451)’, then we can immediately understand what it has decided to do.
- from mid-1950s until the late 1980s, symbolic AI was the most popular approach.

Rules based knowledge

Symbolic AI

- Example of rules:
  - IF animal gives milk THEN animal is mammal
  - IF animal has feathers THEN animal is bird….
- MYCIN (1970s) was an expert system providing doctors with advice about blood diseases used this foundation:
- Typical MYCIN rule:
  - *If:*
    1. *The organism does not strain using Gram method AND*
    2. *The morphology of the organism is rod AND*
    3. *The organism is anaerobic*

    Then:

    There is evidence that the organism is bactericides.
- MYCIN knowledge base had hundreds of rules developed over 5 years by experts in AI and experts from Stanford’s medical school

Variations and developments of “Rules”

Symbolic AI

- Semantic nets: Captures the relationships between concepts and entities using a graphical notation

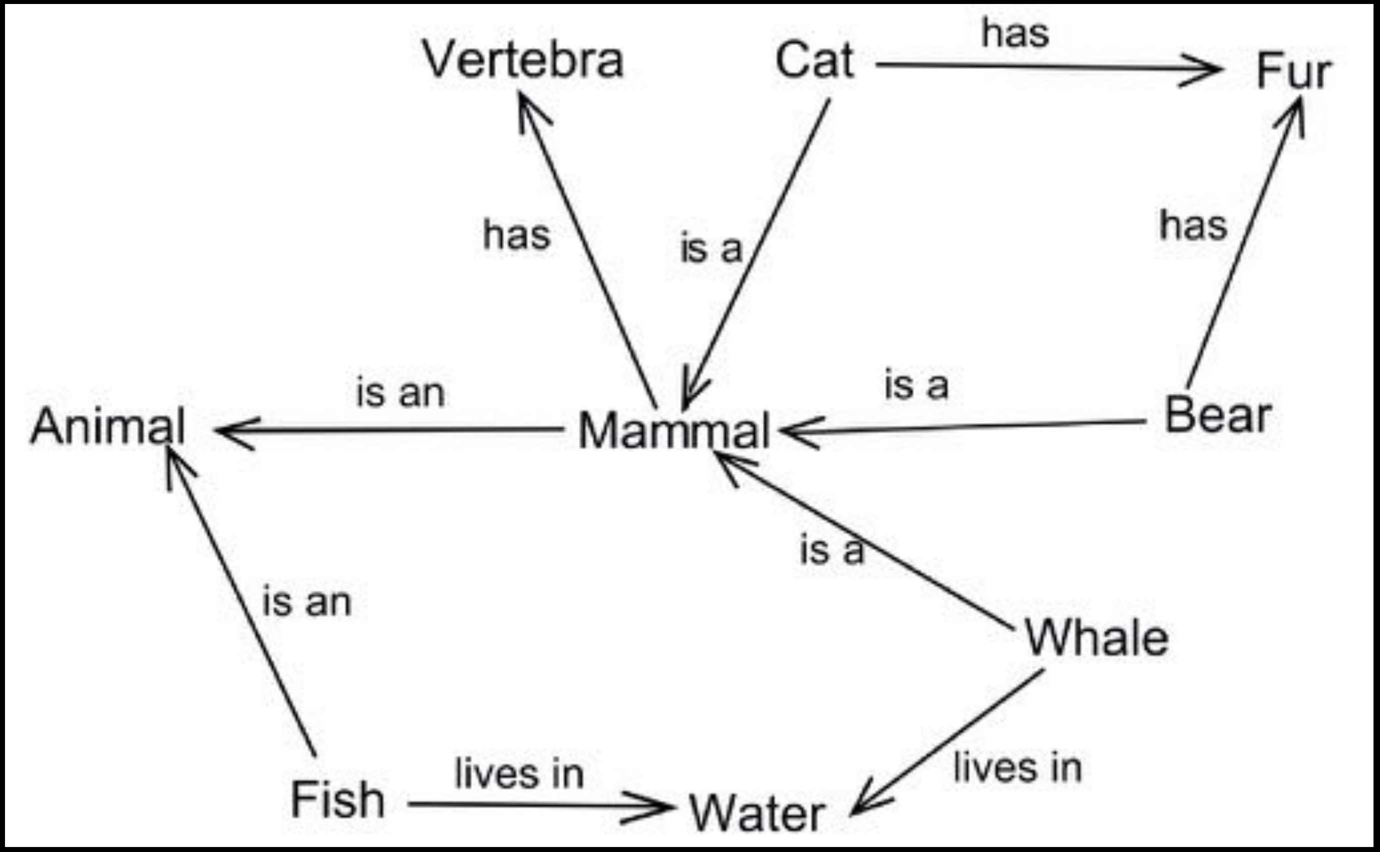

- Logic: added further depth of reasoning to rules, allowing conclusions to be drawn from premises. Eg:
  - All humans are mortal
  - Emma is human
  - Therefore, Emma is mortal
- Various ways of capturing concepts have been developed
- PROLOG is a language developed specifically for logic programming

Cyc: The ultimate Expert System

Symbolic AI

- The **Cyc project** set out to create a complete description of “consensus reality” - the world as we understand it - EVERYTHING
- Original estimated 200 person years to input knowledge, E.g.
    - A plane that runs out of fuel will crash
    - red taps usually produce hot water
    - blue taps usually produce cold water
- 10 years later, the system had 0.5 million rules.
- The system could for e.g. retrieve images according to requests in ordinary English
    - “someone relaxing” retrieved picture of 3 men on beach holding surfboards.
- But huge blank spots - Cyc “knew” various activities caused death but knew nothing about starvation.

Google's Knowledge Graph

Symbolic AI

- Released 30 years after Cyc, **The Google Knowledge Graph** is a semantic network used to support Google Search.
- Much is confidential, but Google says:
  - The database holds billions of facts about people, places, and things.
  - The Knowledge Graph allows us to answer factual questions such as "How tall is the Eiffel Tower?" or "Where were the 2016 Summer Olympics held."
  - Our [Google's] goal with the Knowledge Graph is for our systems to discover and surface publicly known, factual information when it's determined to be useful.

Robotics

Artificial Intelligence

- If you would like to see current robotic capability view:
- https://www.youtube.com/watch?v=htc6OdomjtY

Approach 2: Neural Nets

Artificial Intelligence

- “Models the brain”
- Take inspiration from some structures that occur in the brain, and model these as components in intelligent systems (Perceptrons).
- The brain has 100 billion interconnected components

Single Layer Perceptrons

Neural Nets Generation 1

- Neural networks model the brain's neurons = "perceptron"
  - Each input is given a weight
  - Each neuron operates independently
  - The neuron will "fire" if the sum of the inputs reaches a set threshold
- Single layer NNs very limited - cannot "learn" many simple relationships between inputs and outputs

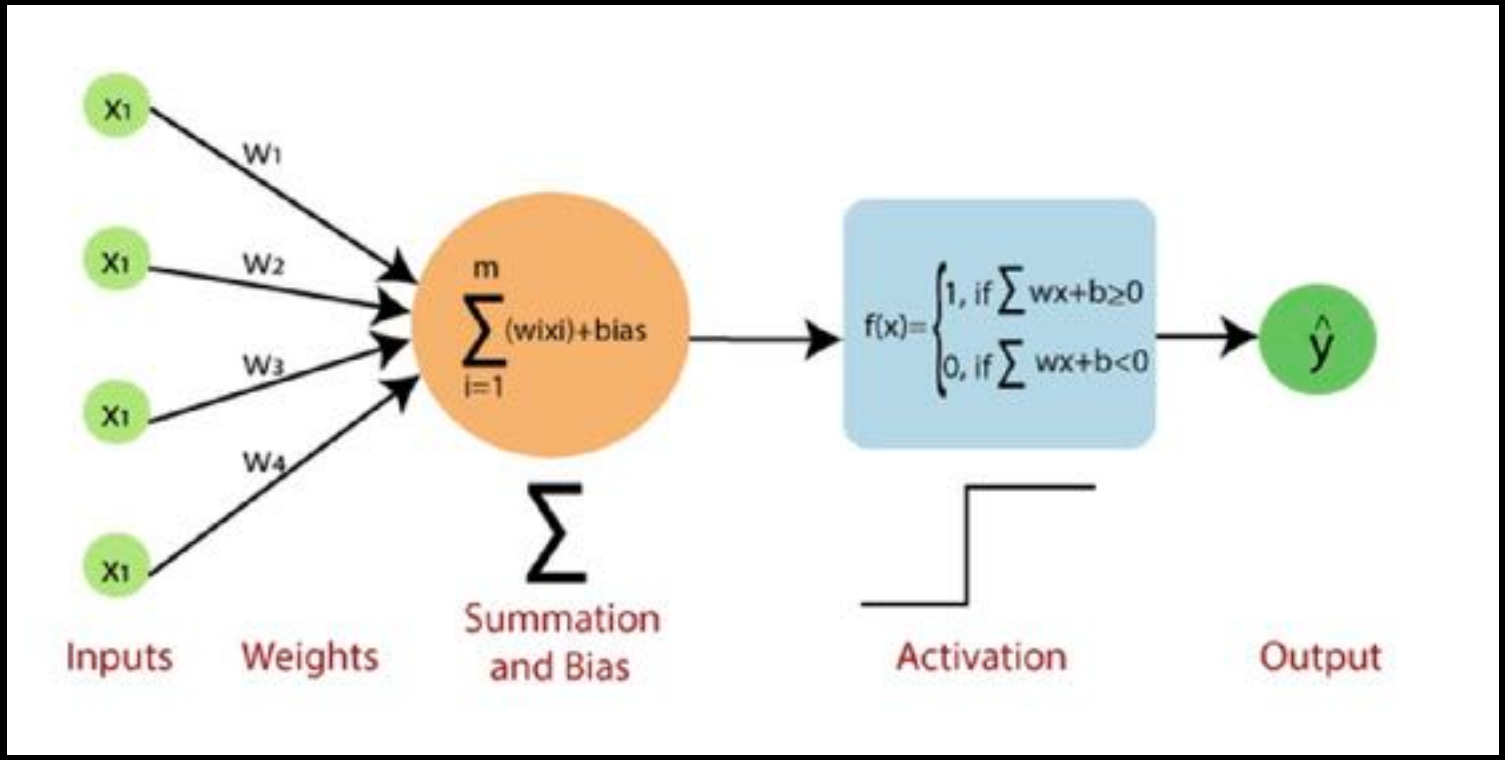

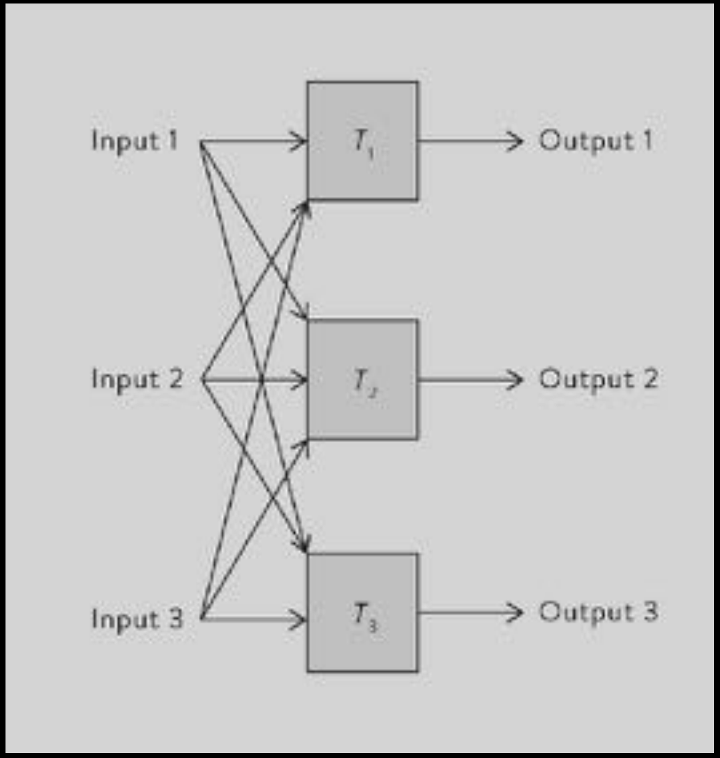

Multi-Layer Perceptrons

Neural Nets Generation 1

- Multi-layer perceptron:
  - every neuron receives all outputs from preceding layer
  - weights and trigger thresholds operate for every neuron.
- Multi-layer perceptrons are able to model complex relationships

but

- Complexity very quickly increases with inputs and layers.
- No one knew how to train neural networks with more than one layer and stopped progress late 1960s

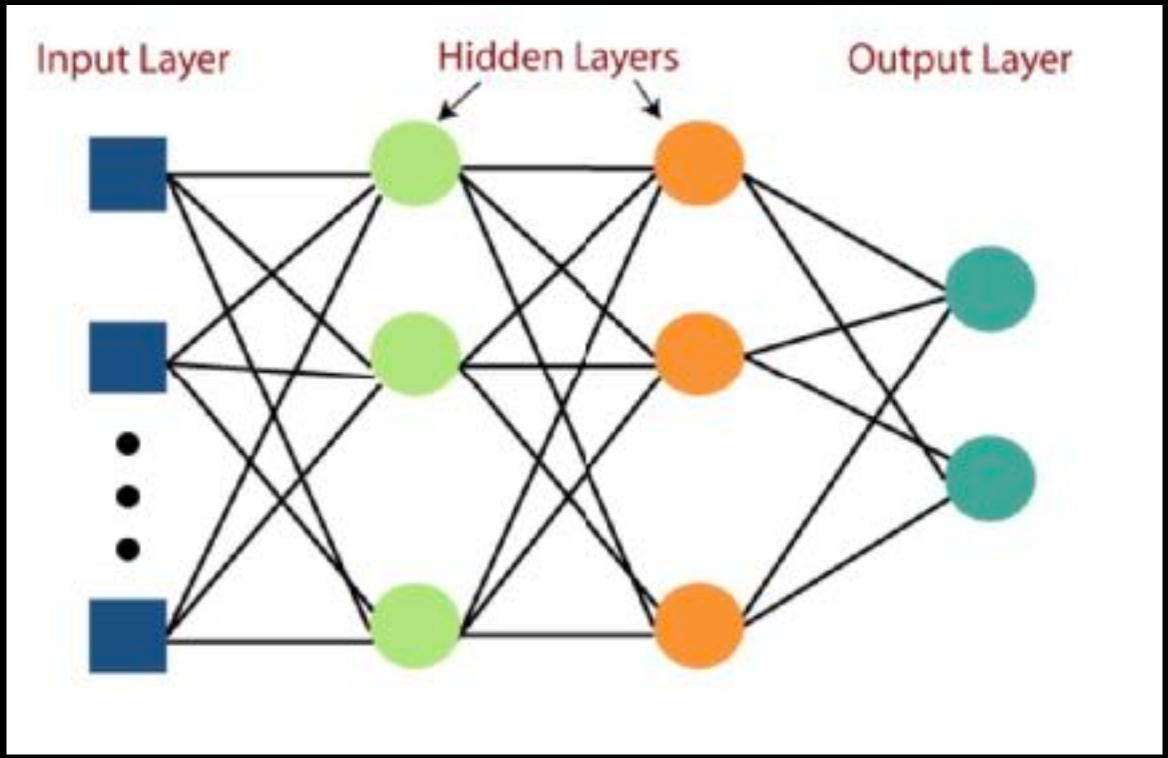

Backprop Training

Neural Nets Generation 2

Next breakthrough came in 1980s with a an algorithm to correct errors in classification by multi-layer perceptrons.

Process is:

- Propagate training data through the network
- Find pattern of errors - difference between prediction and actual output
- Feed errors back layer by layer correcting weights according to the "slope" of the errors
- repeat until network converges on "correct" prediction

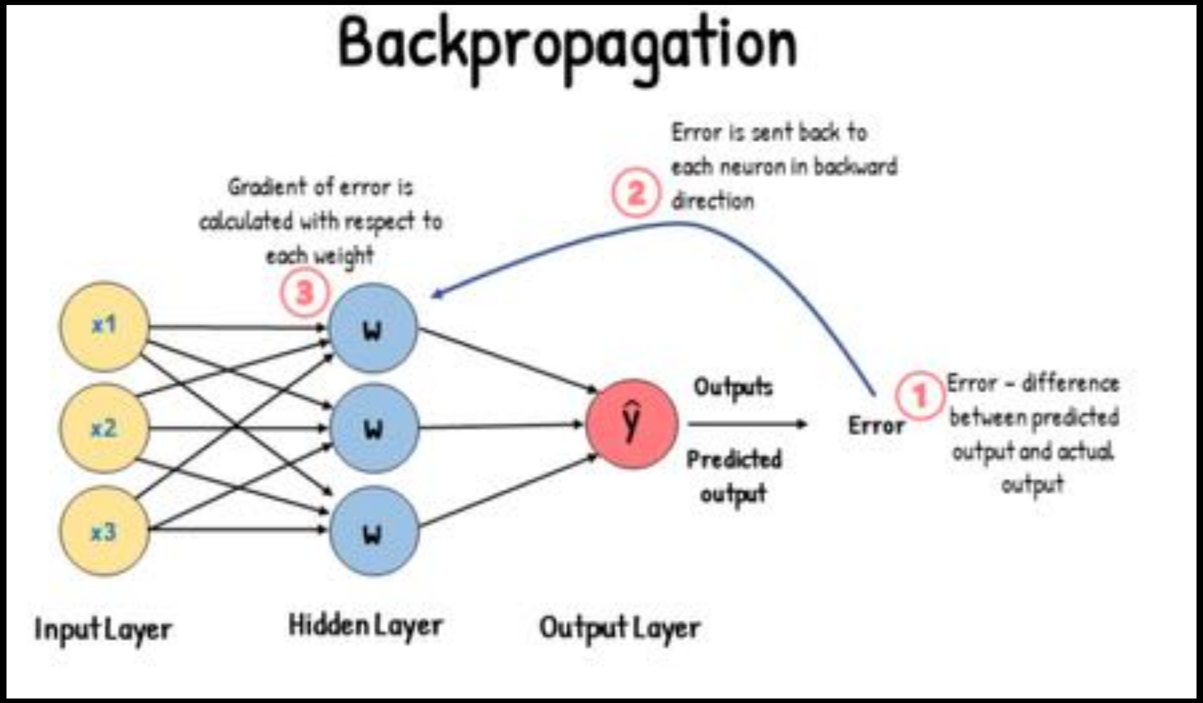

Wikipedia provides a good description of backpropogation

Backprop plus other innovations made application of neural networks viable - *except that computers were still not sufficiently powerful.*

DeepBlue

Neural Nets Generation 2

- IBM, demonstrated DeepBlue in 1997
- Was able to beat Russian grand master Garry Kasparov in the game of chess.
- DeepBlue took its first match from Kasparov in February 1996 in a six-game tournament, but Kasparov won overall, four games to two.
- An upgraded version of the system played Kasparov just over a year later, and this time, DeepBlue defeated Kasparov overall.

Deep Learning

Neural Nets Generation 3

- Deep Learning means:
  - More layers
  - Larger number of neurons
  - Neurons having more connections

- Deep Learning required
  - Improved new techniques for training the nets
  - Huge amounts of data
  - Huge Computing power

Deep Learning

Neural Nets Generation 3

Deep Learning means:

- More layers
  - each layer can process a problem at a different level of abstraction…
- Larger number of neurons
  - 2016 - 1 million - bee level
- Each Neuron having more connections
  - Current - ~10,000 - cat level

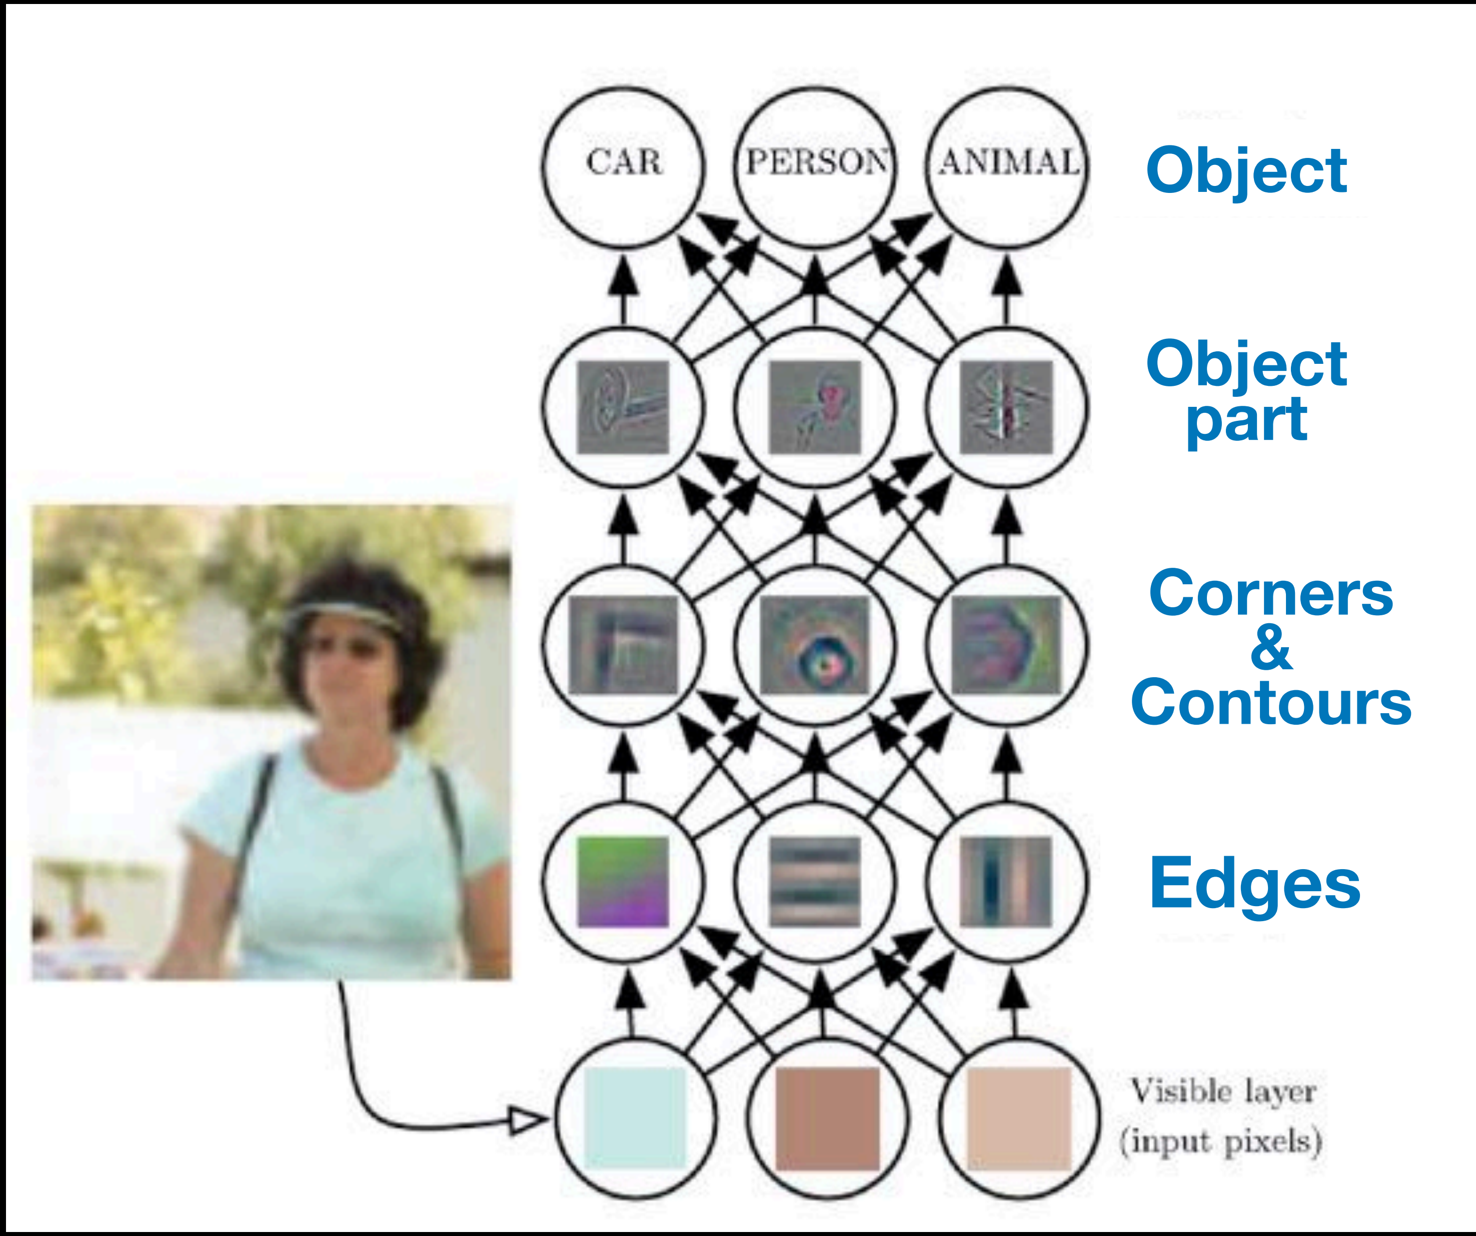

Version 3 - Deep Learning

Deep Learning means:

- More layers
  - each layer can process a problem at a different level of abstraction…
- Larger number of neurons
  - 2016 - 1 million - bee level
- Each Neuron having more connections
  - Current - ~10,000 - cat level

Number of neurons

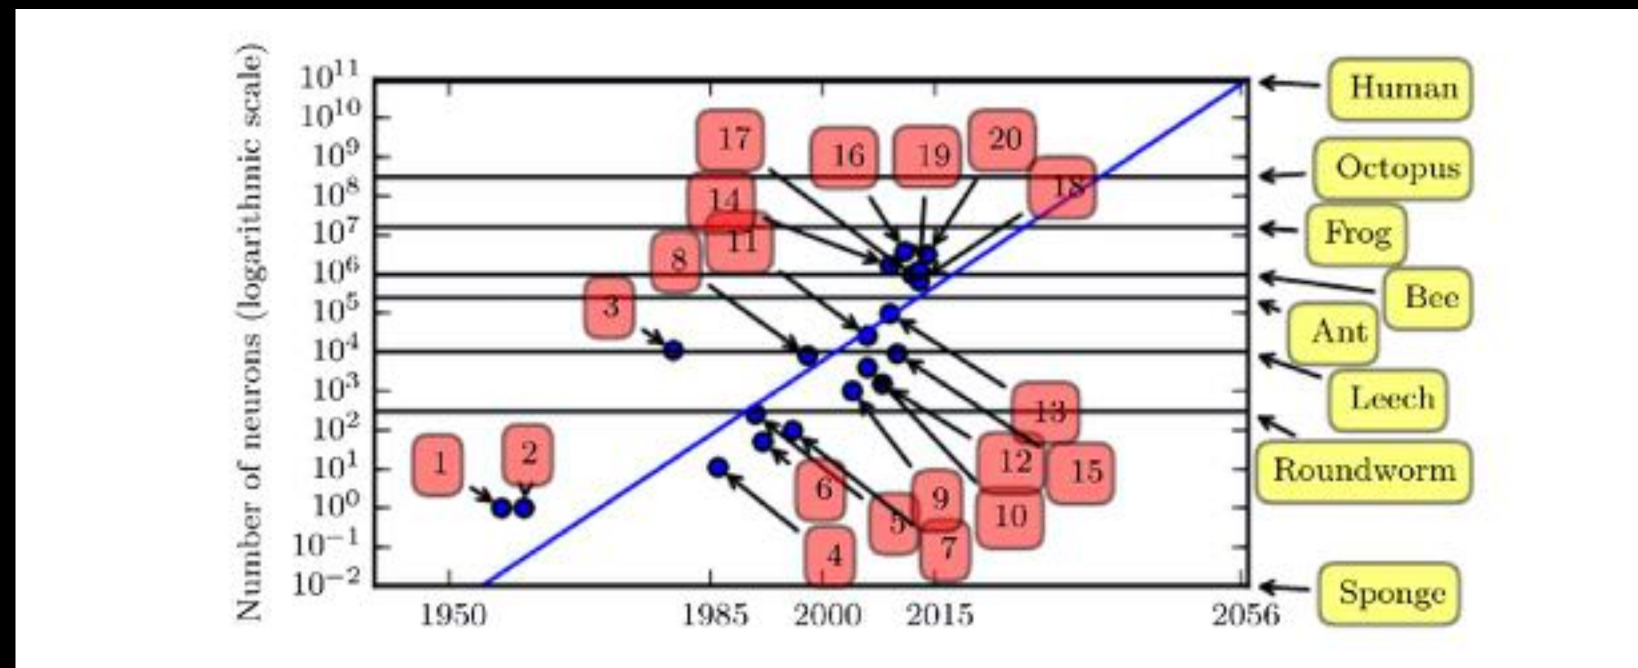

1. Perceptron (Rosenblatt, 1958, 1962)
2. Adaptive linear element (Widrow and Hoff, 1960)
3. Neocognitron (Fukushima, 1980)
4. Early back-propagation network (Rumelhart et al., 1986b)
5. Recurrent neural network for speech recognition (Robinson and Fallside, 1991)
6. Multilayer perceptron for speech recognition (Bengio et al., 1991)
7. Mean field sigmoid belief network (Saul et al., 1996)
8. LeNet-5 (LeCun et al., 1998b)
9. Echo state network (Jaeger and Haas, 2004)
10. Deep belief network (Hinton et al., 2006)
11. GPU-accelerated convolutional network (Chellapilla et al., 2006)
12. Deep Boltzmann machine (Salakhutdinov and Hinton, 2009a)
13. GPU-accelerated deep belief network (Raina et al., 2009)
14. Unsupervised convolutional network (Jarrett et al., 2009)

18. Multi-GPU convolutional network (Krizhevsky et al., 2012)
19. COTS HPC unsupervised convolutional network (Coates et al., 2013)
20. GoogLeNet (Szegedy et al., 2014a)

Version 3 - Deep Learning

Deep Learning means:

- More layers
  - each layer can process a problem at a different level of abstraction…
- Larger number of neurons
  - 2016 - 1 million - bee level
- Each Neuron having more connections
  - Current - ~10,000 - cat level

Connections per neuron

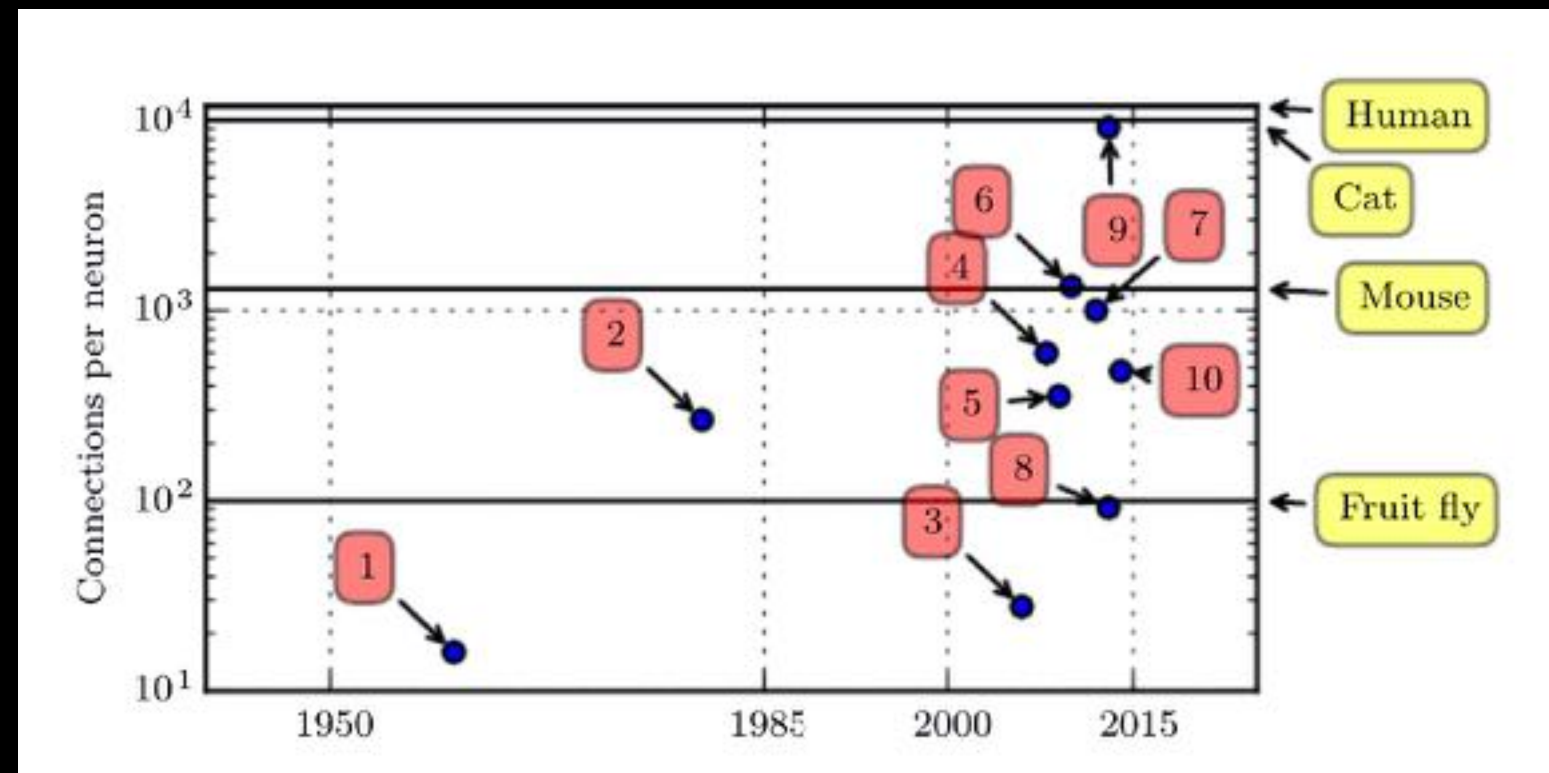

1. Adaptive linear element (Widrow and Hoff, 1960)
2. Neocognitron (Fukushima, 1980)
3. GPU-accelerated convolutional network (Chellapilla *et al.*, 2006)
4. Deep Boltzmann machine (Salakhutdinov and Hinton, 2009a)
5. Unsupervised convolutional network (Jarrett *et al.*, 2009)

8. Multi-GPU convolutional network (Krizhevsky *et al.*, 2012)
9. COTS HPC unsupervised convolutional network (Coates et
10. GoogLeNet (Szegedy *et al.*, 2014a)

Data for Training

Neural Nets version 3

ImageNet / AlexNet was an early example of successful image training

- ImageNet was:
  - An online archive of ~ 14 million images
  - Classified into 22,000 categories using a thesaurus call WordNet
  - Example ~ 1,032 pictures in “Volcanic Crater” category
  - 122 pictures in “frisbee” category - all different pictures, except that all featured a frisbee.
- AlexNet system used ImageNet for training and (2012) demonstrated highly accurate image recognition capability

Processing Power

Neural Nets version 3

- Training a deep neural net requires a huge amount of computer-processing time.
- The work that needs to be done in training is not particularly complex – but there is an enormous amount of it
- Graphical Processing Units (GPUs):
  - Designed to accelerate computer graphics and image processing for games consoles, PCs, mobile phones.
  - Have a very parallel structure and simple instruction set.
  - “Discovered” to be ideal for training of neural nets
  - GPUs are 100 times faster than CPUs for some highly parallel problems

(Nvidia holds 80% of market in specialist AI chips, market value ~ $400 Billion.)

Deep Learning

Neural Networks

- The next slides show some developments in use of Deep Learning.

Deep Learning progress example

Deepmind Technologies (Google) developed a series of neural network programs

2016: AlphaGo program beat a world champion Go player Lee Sedol in a 5-game match

Program was “given” game examples, domain knowledge and rules

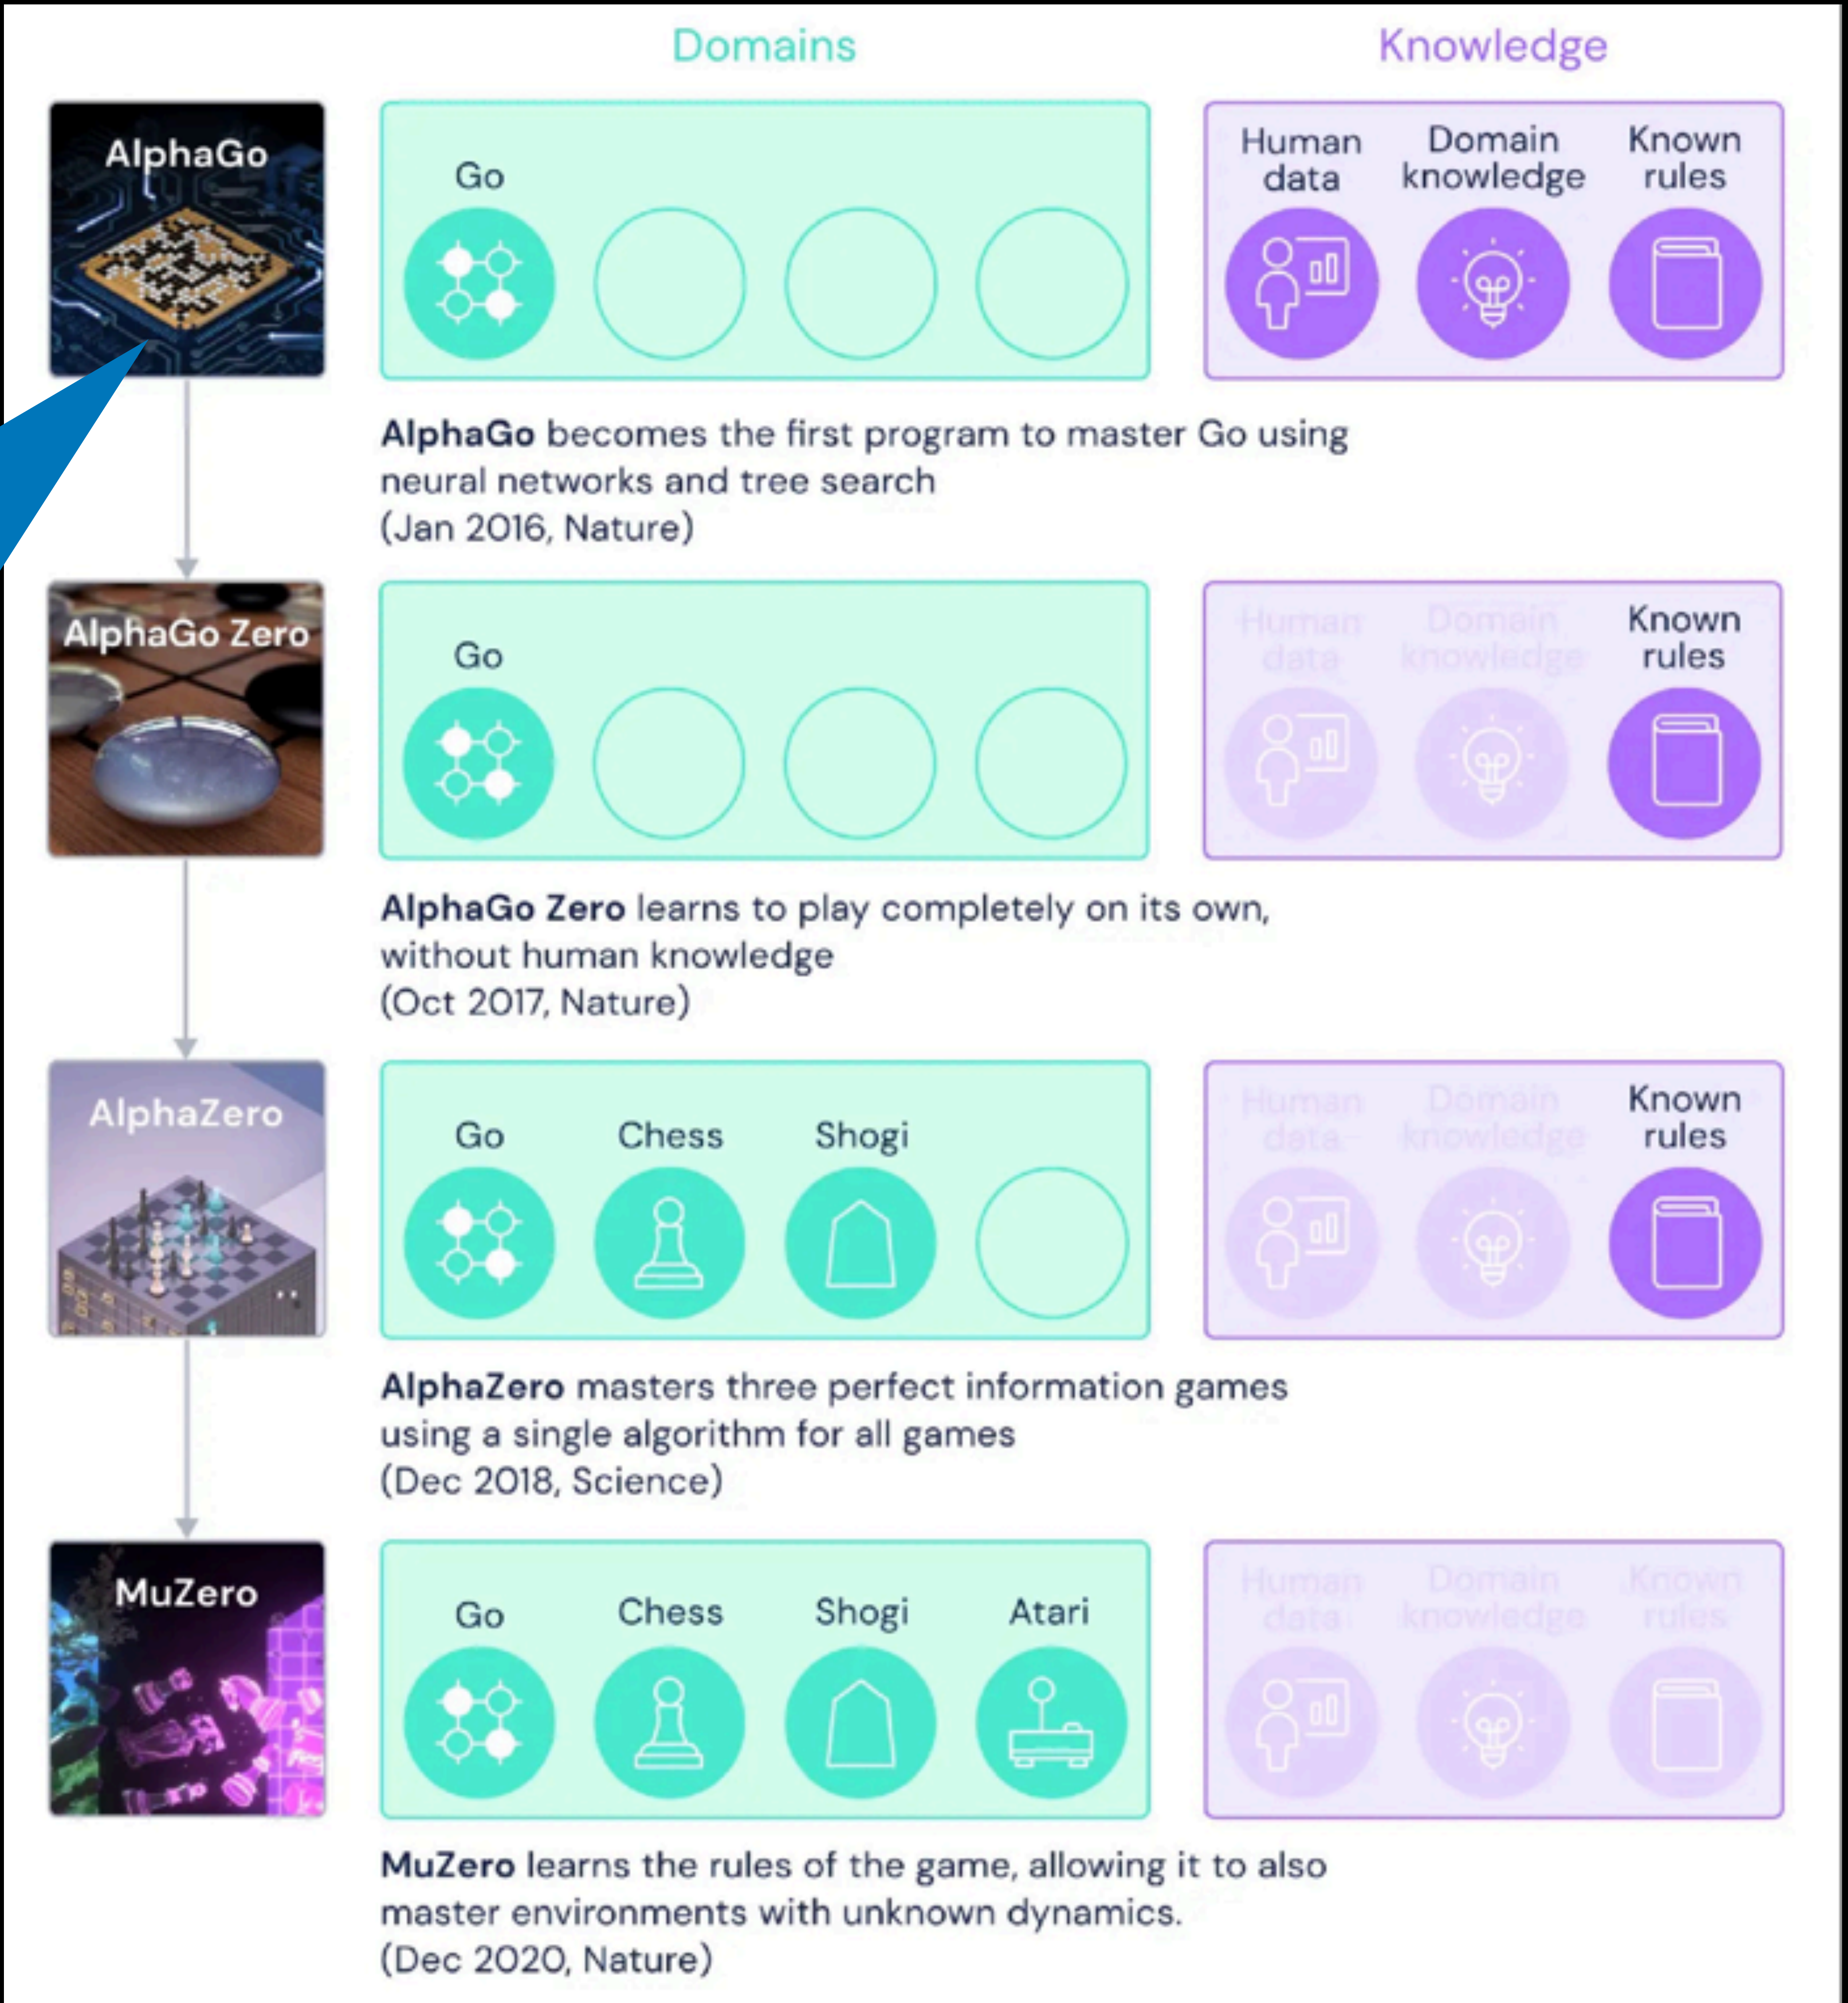

Deep Learning progress example

2017: The algorithms were improved to allow AlphaGo Zero to be able to learn Go only having the rules.

By playing games against itself, AlphaGo Zero surpassed the strength of AlphaGo Lee in three days by winning 100 games to 0, reached the level of AlphaGo Master in 21 days, and exceeded all the old versions in 40 days

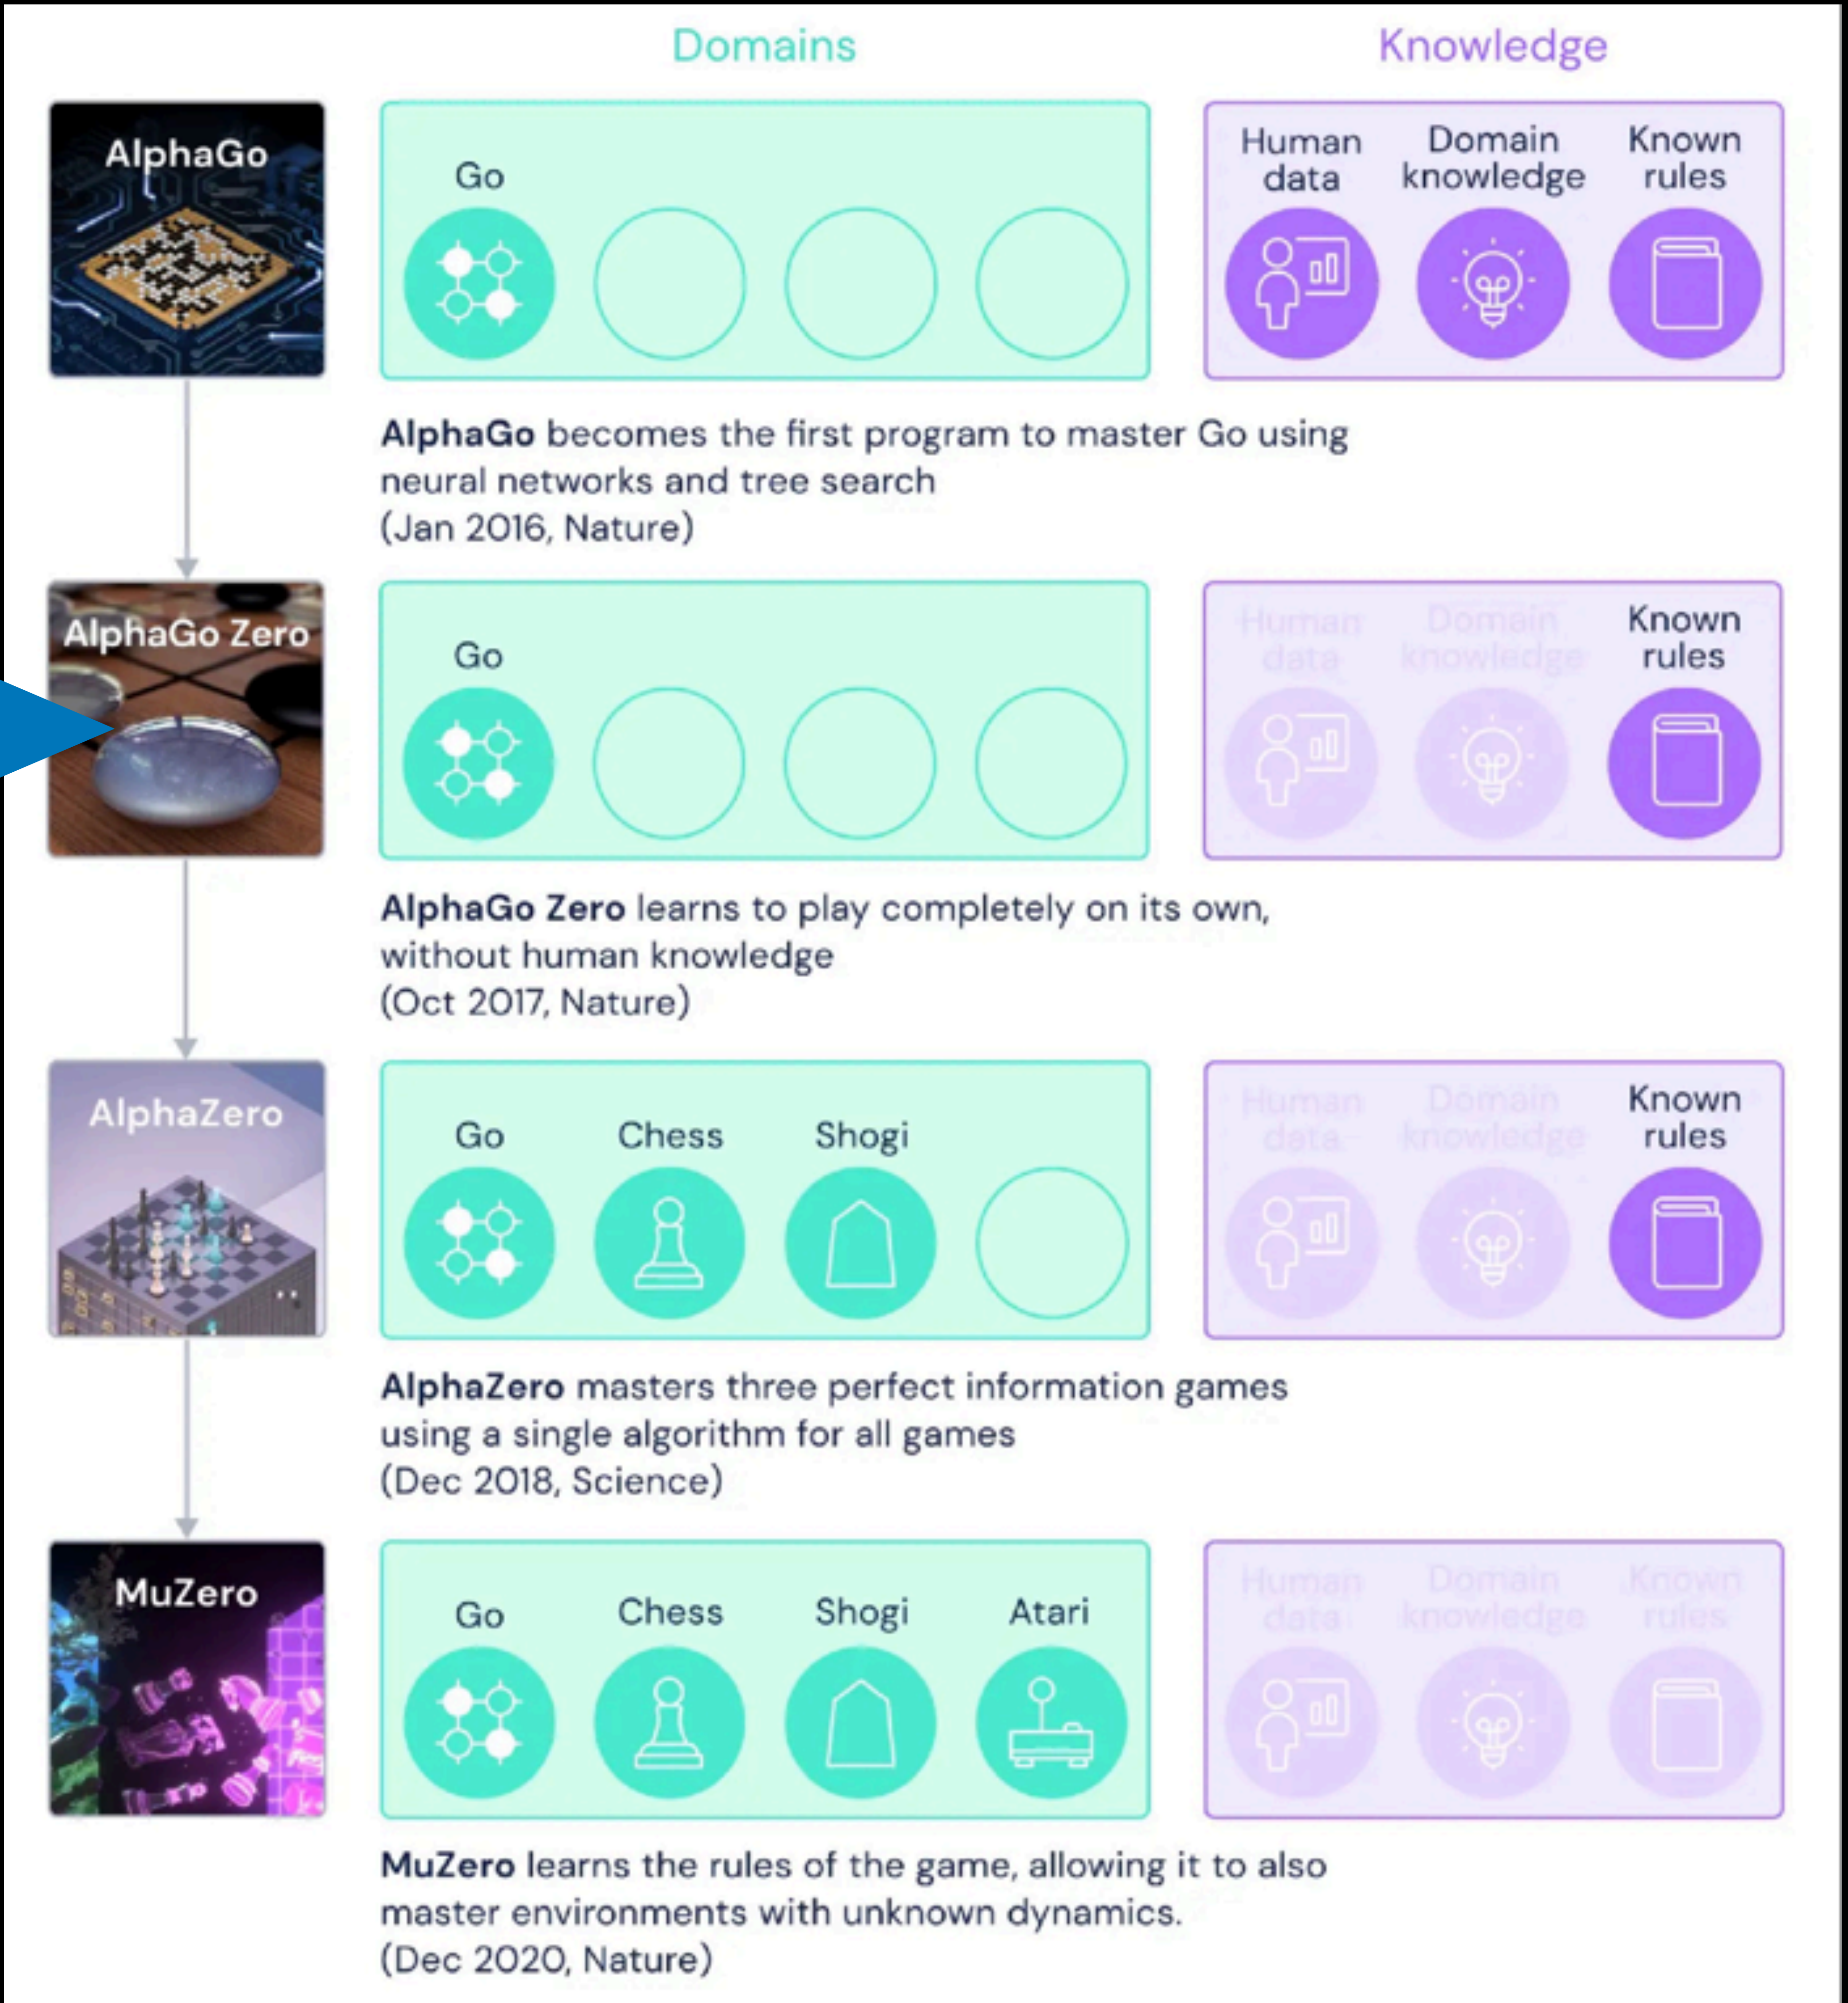

Deep Learning progress example

2017: Alpha Zero was generalised to learn multiple games and was able to beat the most powerful programs playing Go, Chess and Shogi after a few days of playing against itself using reinforced learning.

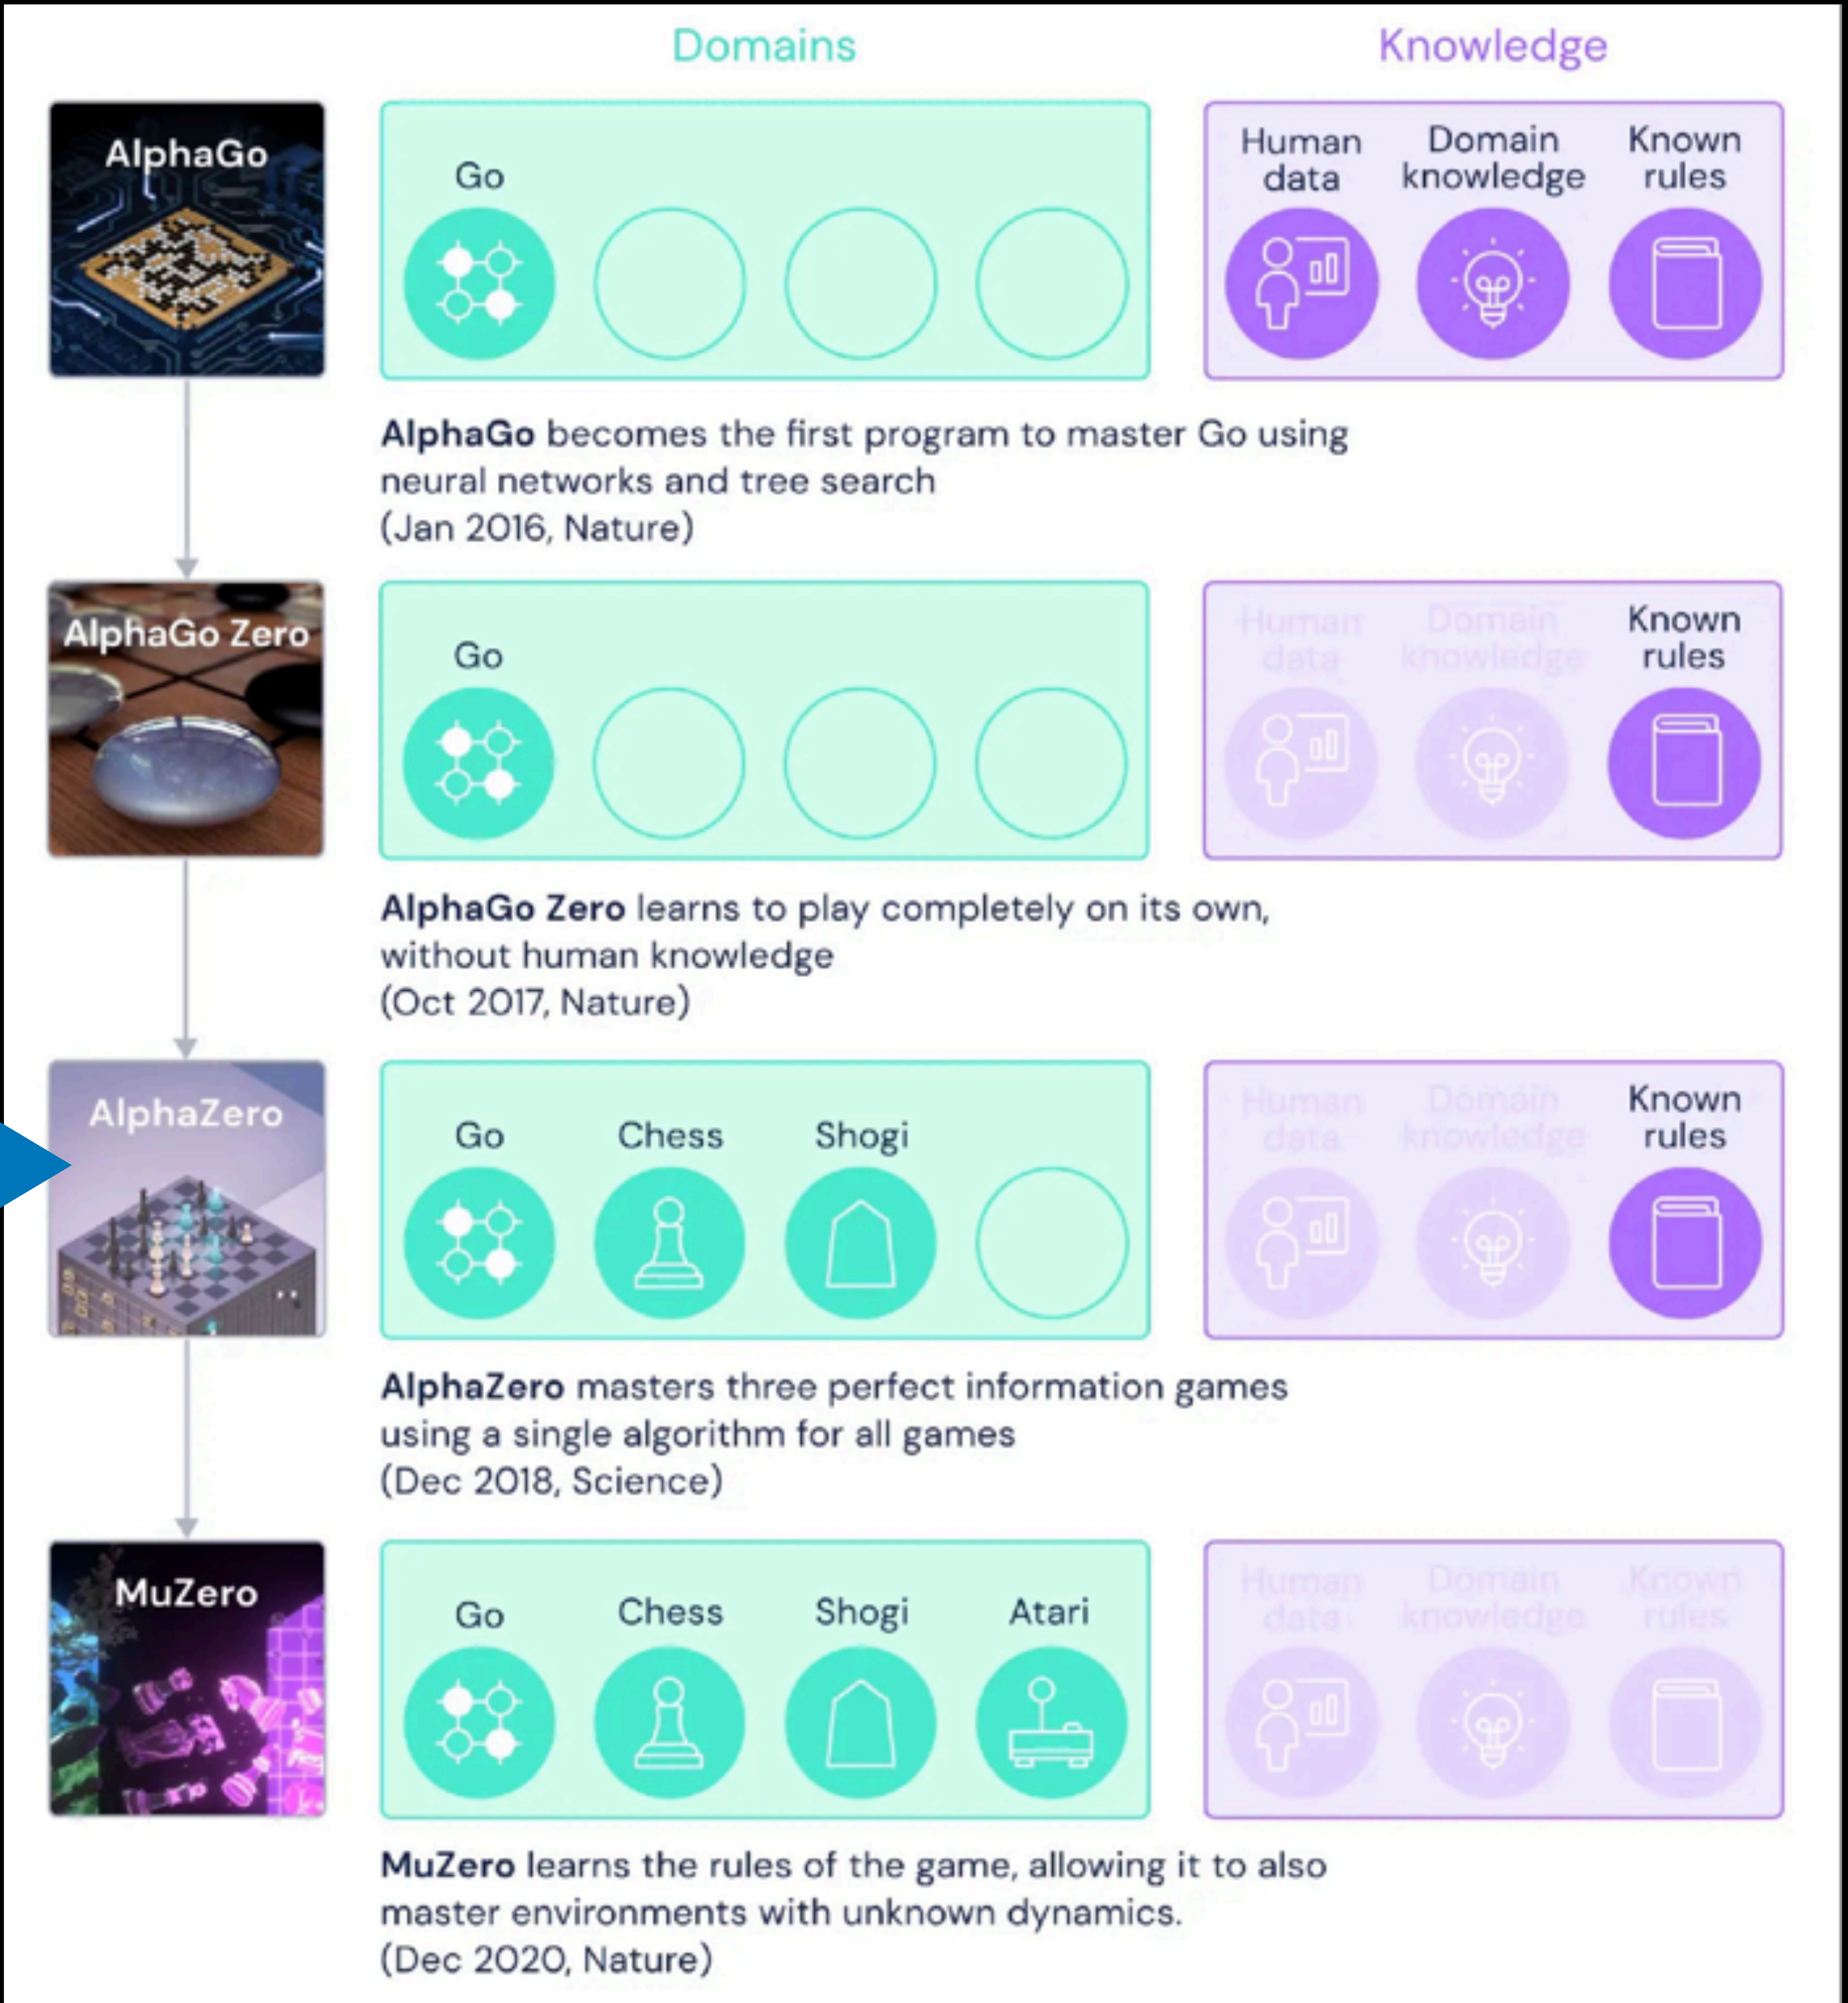

Deep Learning progress example

MuZero was developed to master games without knowing the rules.

was trained via self-play, with no access to rules, opening books, or endgame table bases.

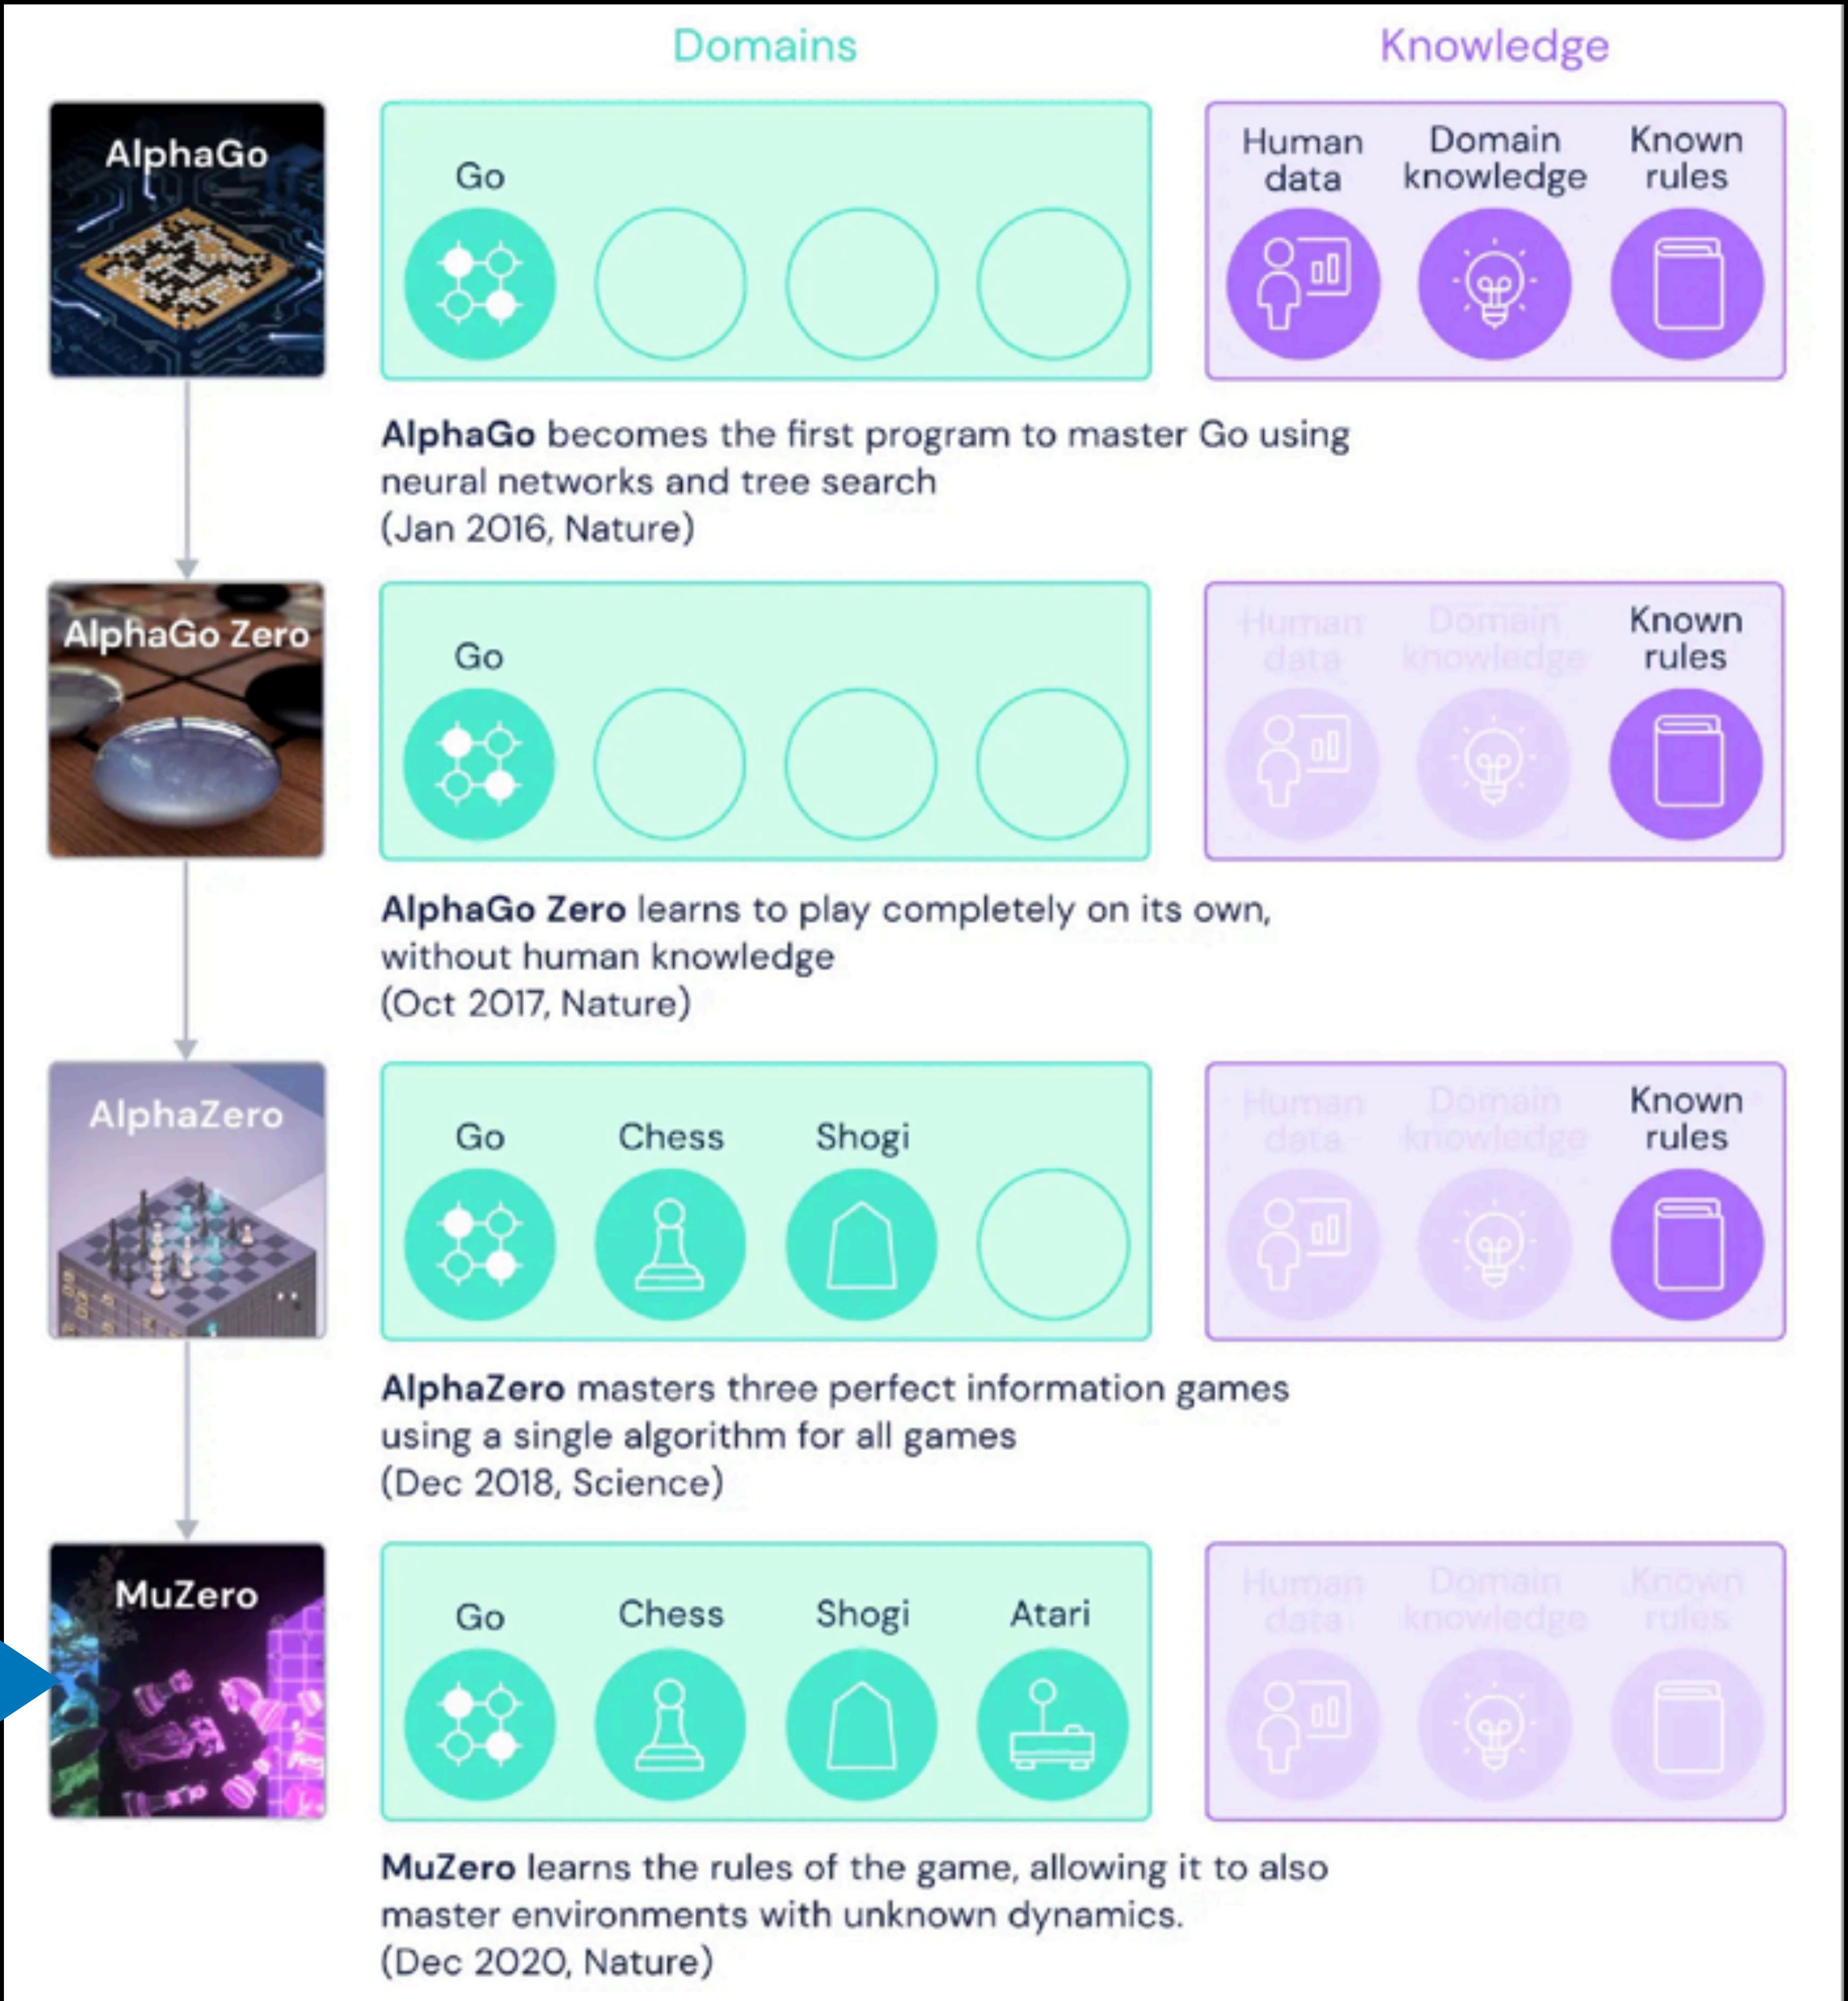

Deep Learning progress example

MuZero was given the objective of playing 7 Atari games.

To learn Atari games, MuZero was fed the raw pixels.

The video shows the learning path for the Breakout game.

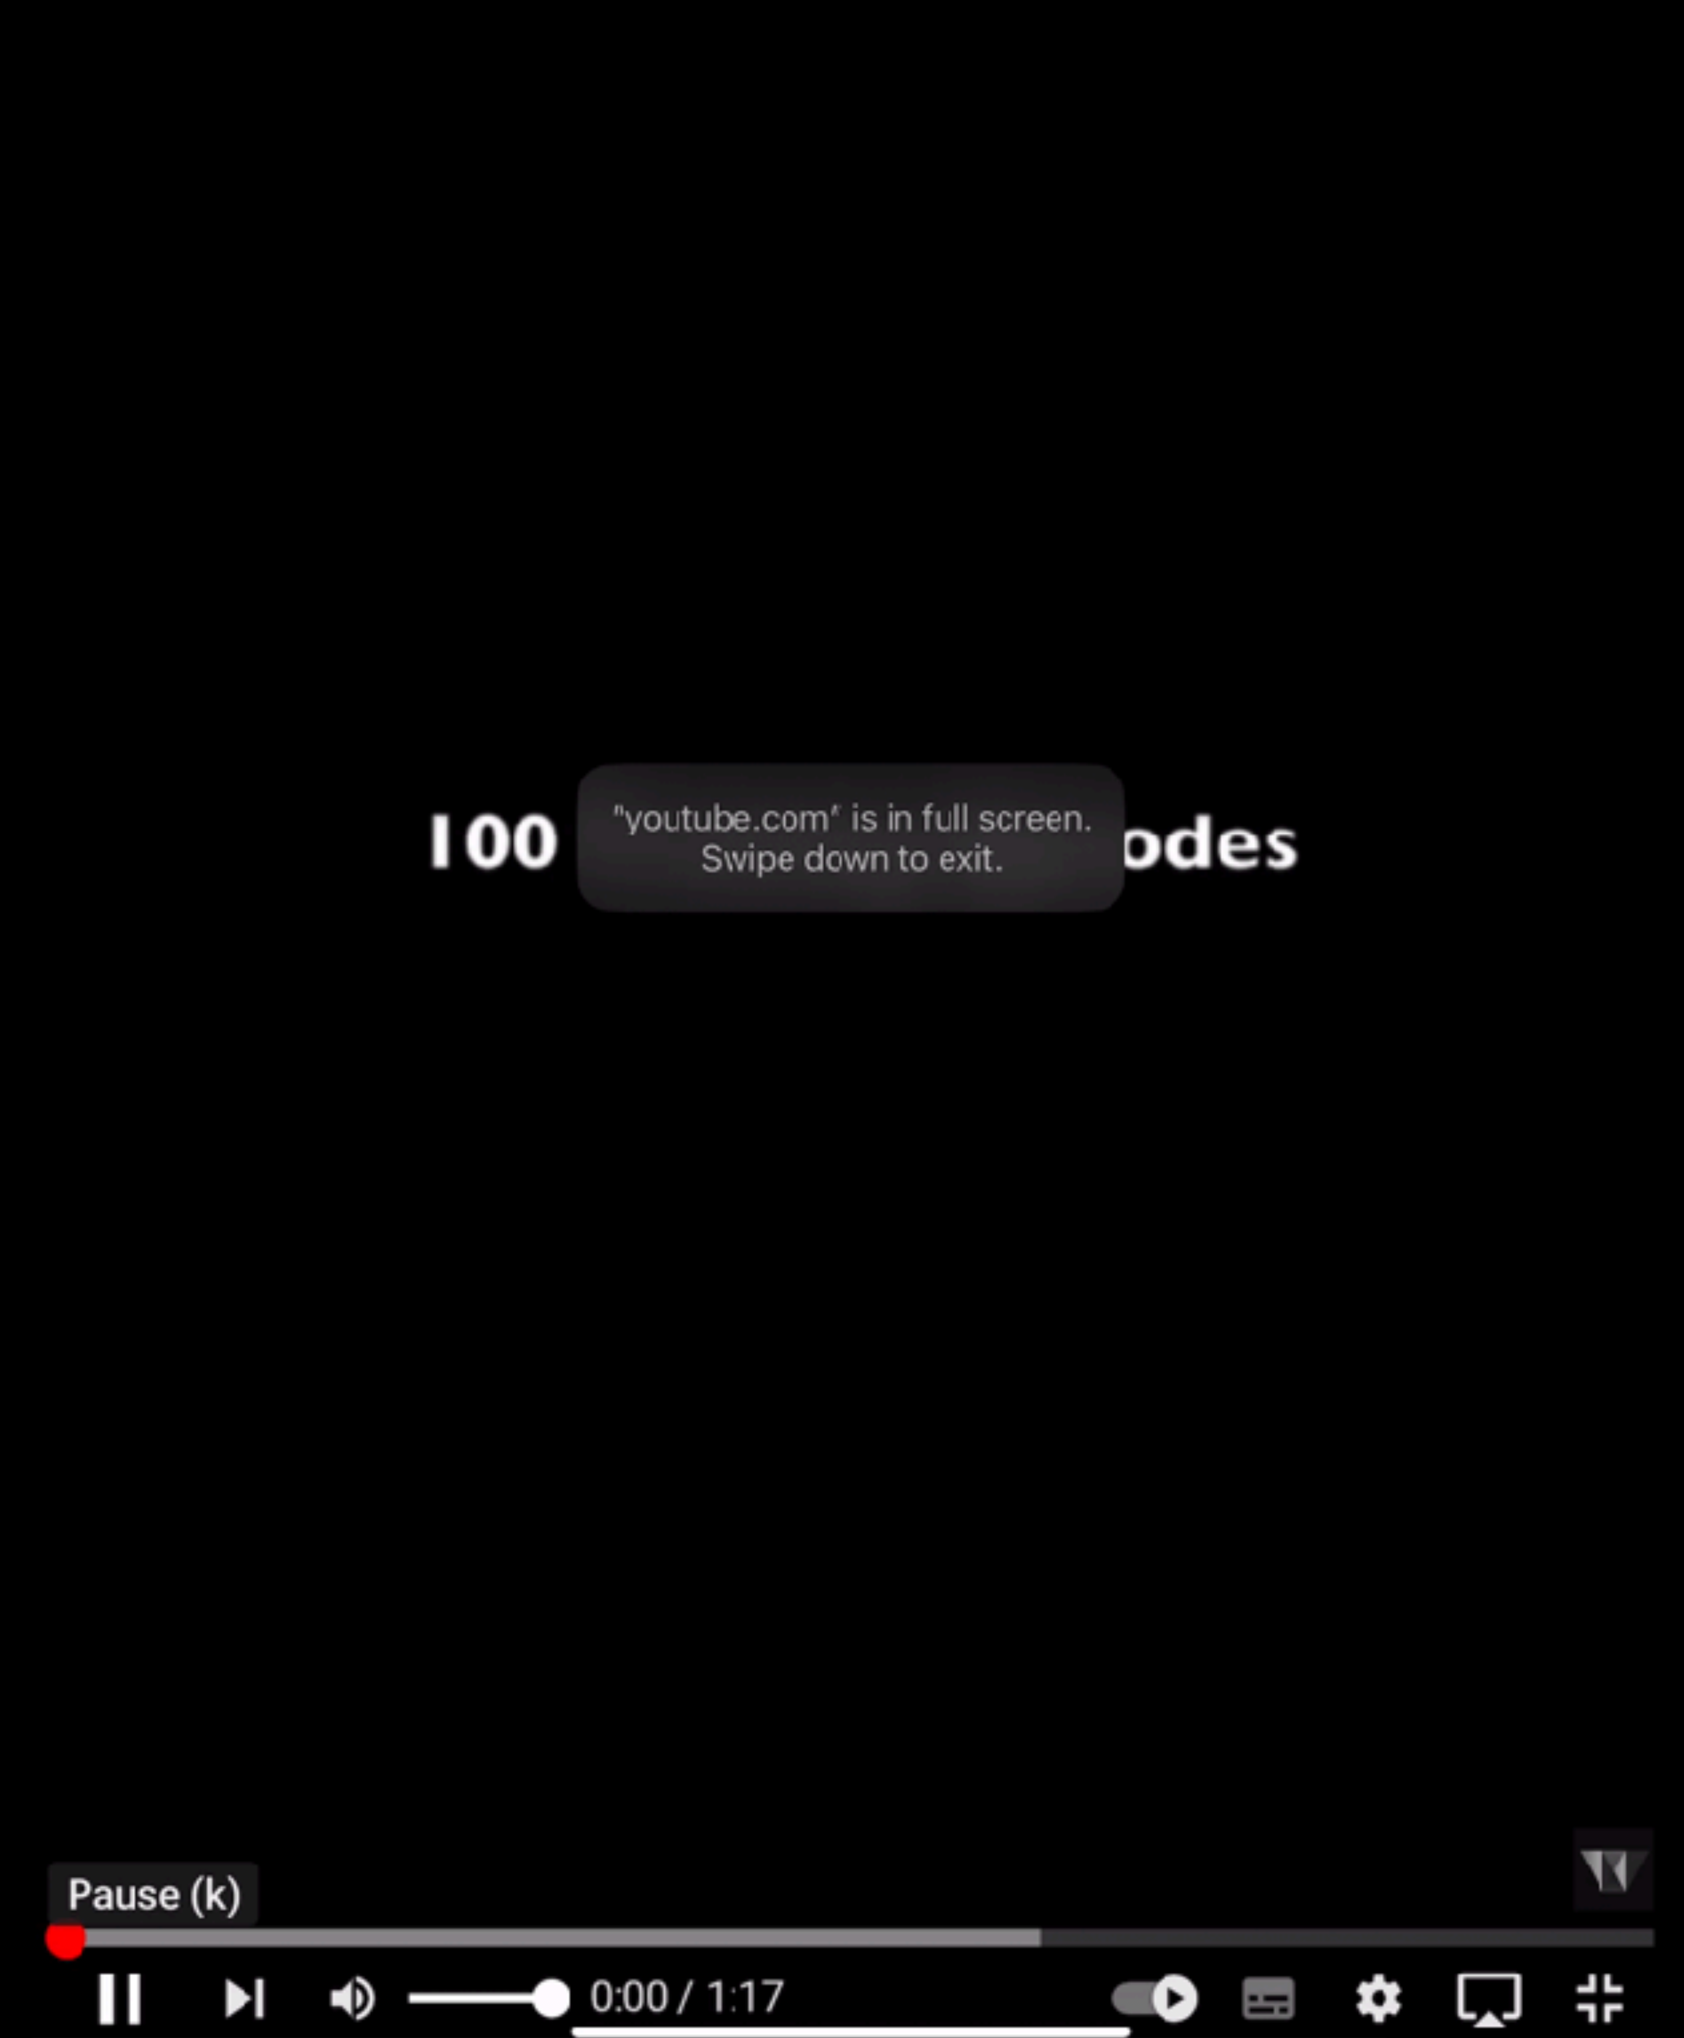

The AI Frontier

Artificial Intelligence

Generative AI: Any AI system whose primary function is to generate Content (as opposed to classifying data or choosing actions etc.)

Large Language Models: AI system that works with Language

Foundation Model: AI system with broad capabilities that can be adapted to a range of different, more specific purposes. (a foundation model GPT-3.5 was adapted and built into ChatGPT)

Large Language Model

Artificial Intelligence

- A language model is a probabilistic model of a natural language that can generate probabilities of a series of words, based on text it was trained on.
- a combination of feedforward neural networks and transformers.
  - Translate text into tokens representing words and sequences of words with weighted relationships.
- Models can have up to billions of weights.
- Are thought to acquire embodied knowledge of the “corpora”, *but also its inaccuracies and biases.*

LLM Training

Artificial Intelligence

The training process of a large language model involves:

- Pre-processing the text data to convert it into a numerical representation that can be fed into the model.
- Randomly assigning the model's parameters.
- Feeding the numerical representation of the text data into the model.
- Using a loss function to measure the difference between the model's outputs and the actual next word in a sentence.
- Optimising the model's parameters to minimise loss.
- Repeating the process until the model's outputs reach an acceptable level of accuracy.

Training methods - GAN

Artificial Intelligence

- Generative adversarial networks (GANs) are one way of managing neural network training
- GANs pit two neural networks against each other:
  - a generator that generates new examples
  - a discriminator that learns to distinguish the generated content as either real (from the domain) or fake (generated).
- The two models are trained together and get smarter as the generator produces better content and the discriminator gets better at spotting the generated content.
- This procedure repeats, pushing both to continually improve after every iteration until the generated content is indistinguishable from the existing content.

Training methods - Diffusion

Artificial Intelligence

Diffusion Model:

- Determine vectors in latent space through a two-step process during training.
- The two steps are forward diffusion and reverse diffusion.
- The forward diffusion process slowly adds random noise to training data
- The reverse process reverses the noise to reconstruct the data samples.

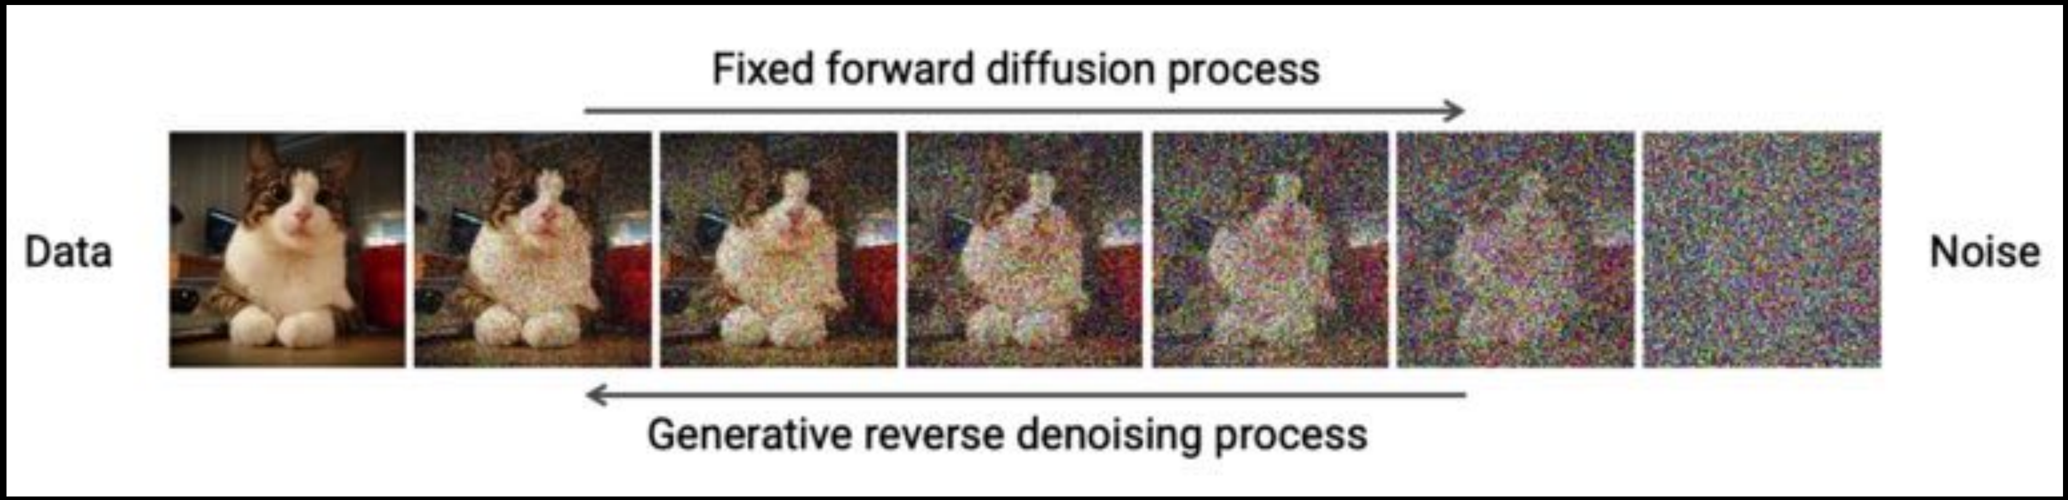

Generative AI

Artificial Intelligence

Generative AI models learn the patterns and structure of their input training data and then generate new data that has similar characteristics.

- Take raw data — say, all of Wikipedia or the collected works of Rembrandt — and encode a simplified representation of the training data
- “learn” to generate statistically probable outputs when given a “prompt”
- prompt might be in the form of a text, an image, a video, a design, musical notes.
- Return new content in response to the prompt, creating a new work that’s similar, but not identical, to the original data.
- Content can include essays, solutions to problems, or realistic fakes created from pictures or audio of a person.

Back to Reality - Limitations

Artificial Intelligence

Symbolic AI

- Very difficult to capture implicit knowledge
- Complexity of knowledge storage

Neural Network systems:

- Can be fragile: a small change of input may produce a very different and unexpected output
- Not transparent: Cannot see the logic from input to output
- Does not always identify the source of content.
- Training data can be biassed or include inaccuracies (e.g. “scraped” from the internet)
- Available training data is limited compared to “real world”
- Realistic-sounding content makes it harder to identify inaccurate information.
- Results can gloss over bias and inaccuracy
- It can be difficult to understand how to tune for new circumstances.

AI Limitations

Artificial Intelligence

Consider this questions, a test for comprehension called a Winograd schema:

- Statement 2a: The trophy doesn’t fit into the brown suitcase because it is too small.
- Statement 2b: The trophy doesn’t fit into the brown suitcase because it is too large.

- Question: What is too [small/large]?

AI Limitations

Artificial Intelligence

- Another similar challenge for AI involves understanding of the human world, and the unwritten rules that govern our relationships within it.
- Consider the following short dialogue from the psychologist and linguist Steven Pinker:
  - Bob: ‘I’m leaving you.’
  - Alice: ‘Who is she?’
- Can you explain this dialogue? Of course you can.
- The knowledge base required to understand this dialogue is immense

The Brain as an AI system

Artificial Intelligence

- Billions of cells already arranged by genetic design in specialised but interconnected models covering:
  - Vision
  - Sensation, perception, spatial reasoning
  - Decision, problem solving, consciousness
  - Balance & coordination
  - Memory, hearing, language
  - Thirst, hunger, sleep, mood
- Primary learning based on continuous input of 5 senses - sight (3d colour), hearing, smell, taste, touch from the real world
- Together with control of motion.
- Secondary learning based on continuous aural language aligned with real world meaning….

AI has a way to go - but amazing for specific jobs

Thank You