



Lecture 1.2 – Introduction to Generative AI

Generative AI Teaching Kit





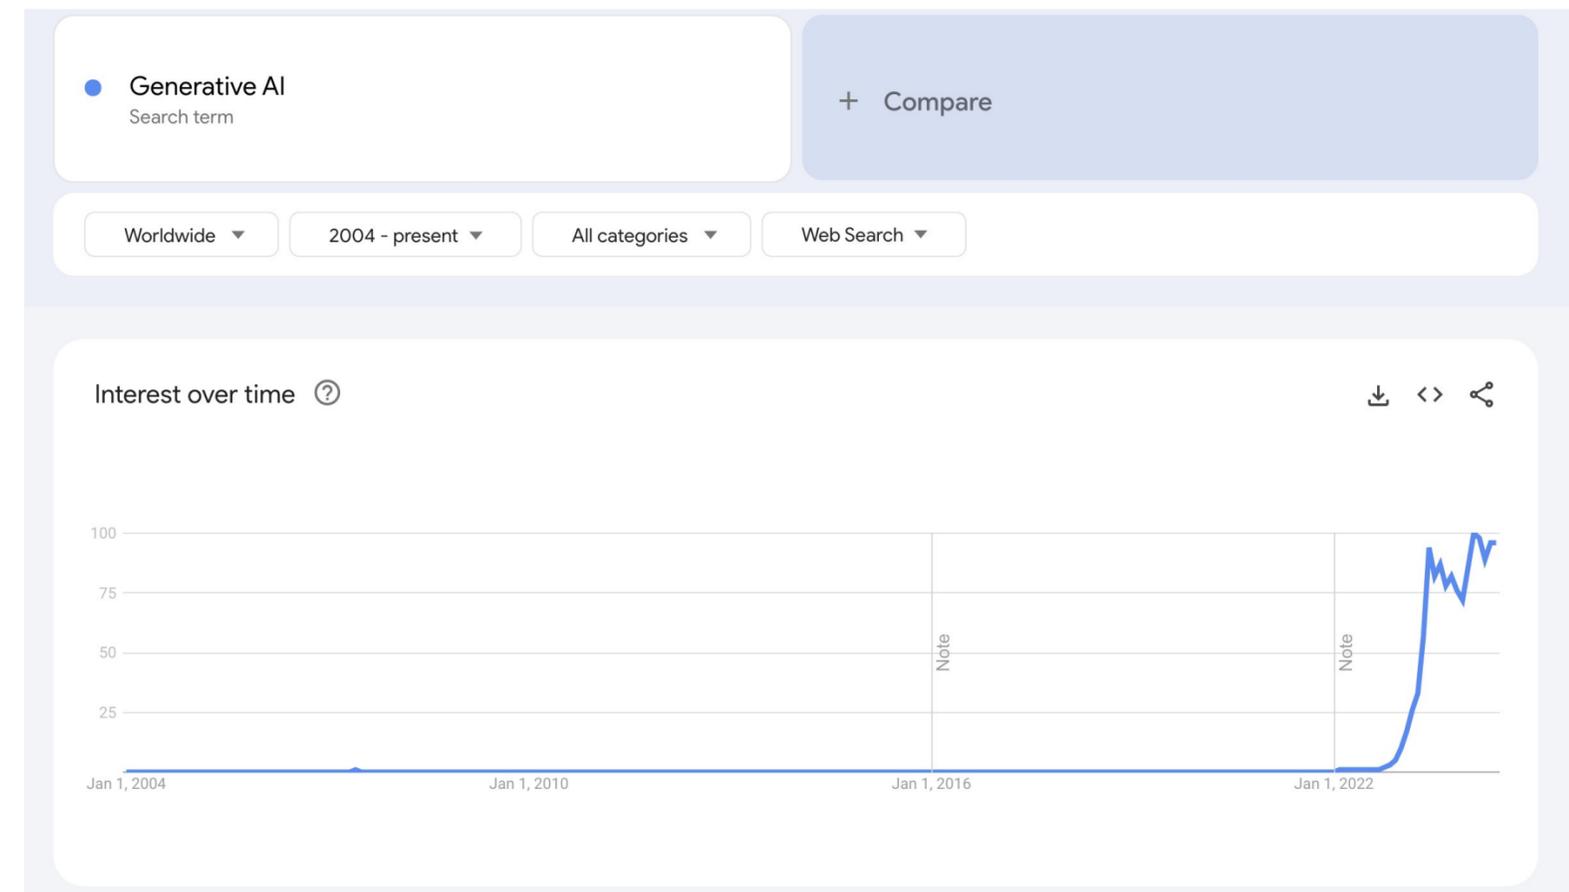
The NVIDIA Deep Learning Institute Generative AI Teaching Kit is licensed by NVIDIA and Dartmouth College under the [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/).

This lecture

- What is “GenAI”?
- History of AI over the last 50 years
- Pitfalls/risks of GenAI

The Rise of Generative AI “GenAI”

- The development of AI reached a new level in November of 2022 with the launch of ChatGPT
- ChatGPT, a chatbot, took the world by storm and was built on top of the Large Language Model GPT-3.
- GPT-3, for the research community, also shook the domain when it was released, showing that models could learn from unsupervised data and then further learn new tasks without further training.
- The race for supremacy began. These models, being able to respond in human-like fashion, have brought the topic of AI into the mainstream.



What is “GenAI”?

GenAI (aka Generative AI) refers to a subset of models and approaches that have been trained specifically to generate new data.

This is different to other, un/supervised machine learning models that are trained to classify data samples, or connect input to output values for regression

Examples of **non-GenAI models**:

- Convolutional Neural Networks
- Logistic Regression
- Random Forest

Examples of **GenAI models**:

- Large Language Models
- Diffusion Models
- Generative Adversarial Networks

Artificial Intelligence

Is the field of study

Machine Learning

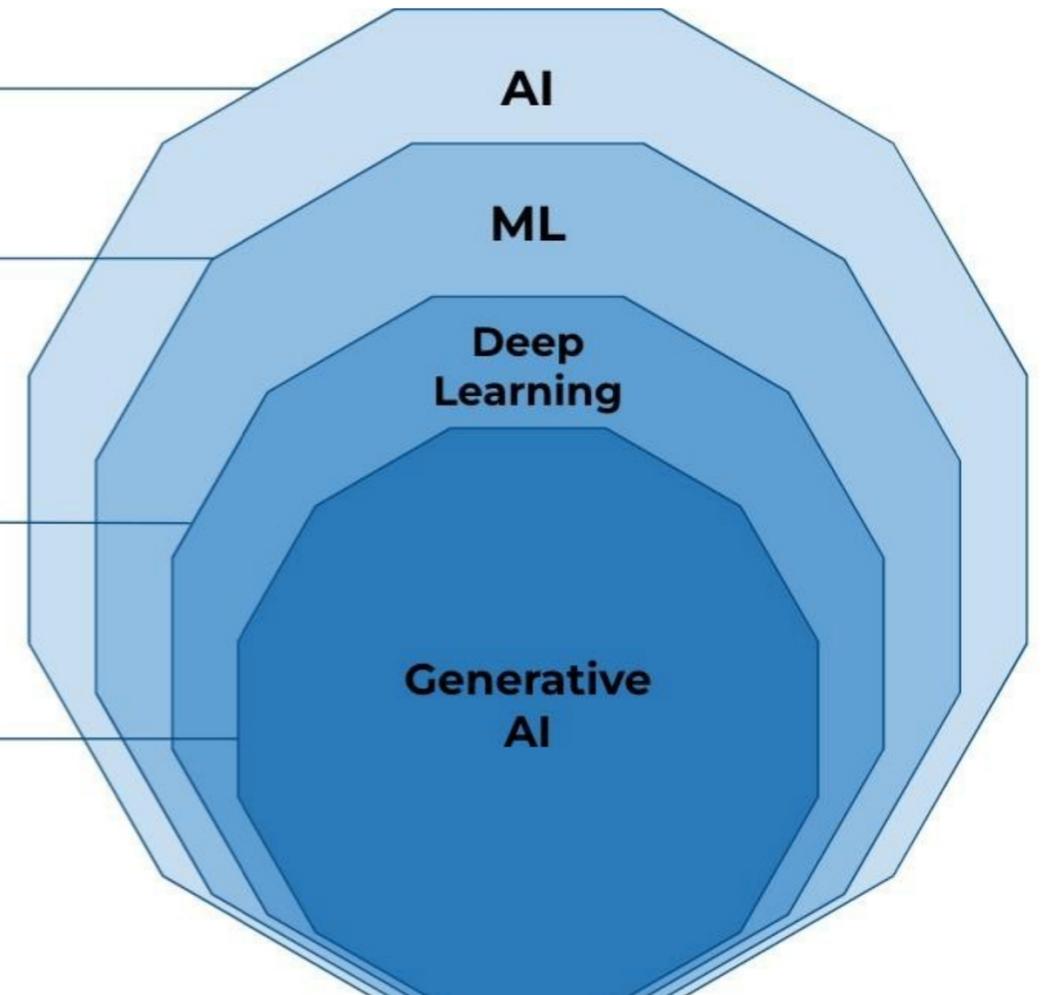
Is a branch of AI that focus on the creation of intelligent machines that learn from data. Another very well know branch inside AI is **Optimization**.

Deep Learning

Is a subset of Machine Learning methods, based on **Artificial Neural Networks**.
Examples: CNNs, RNNs

Generative AI

A type of ANNs that generate data that is similar to the data it was trained on.
Examples: GANs, LLMs

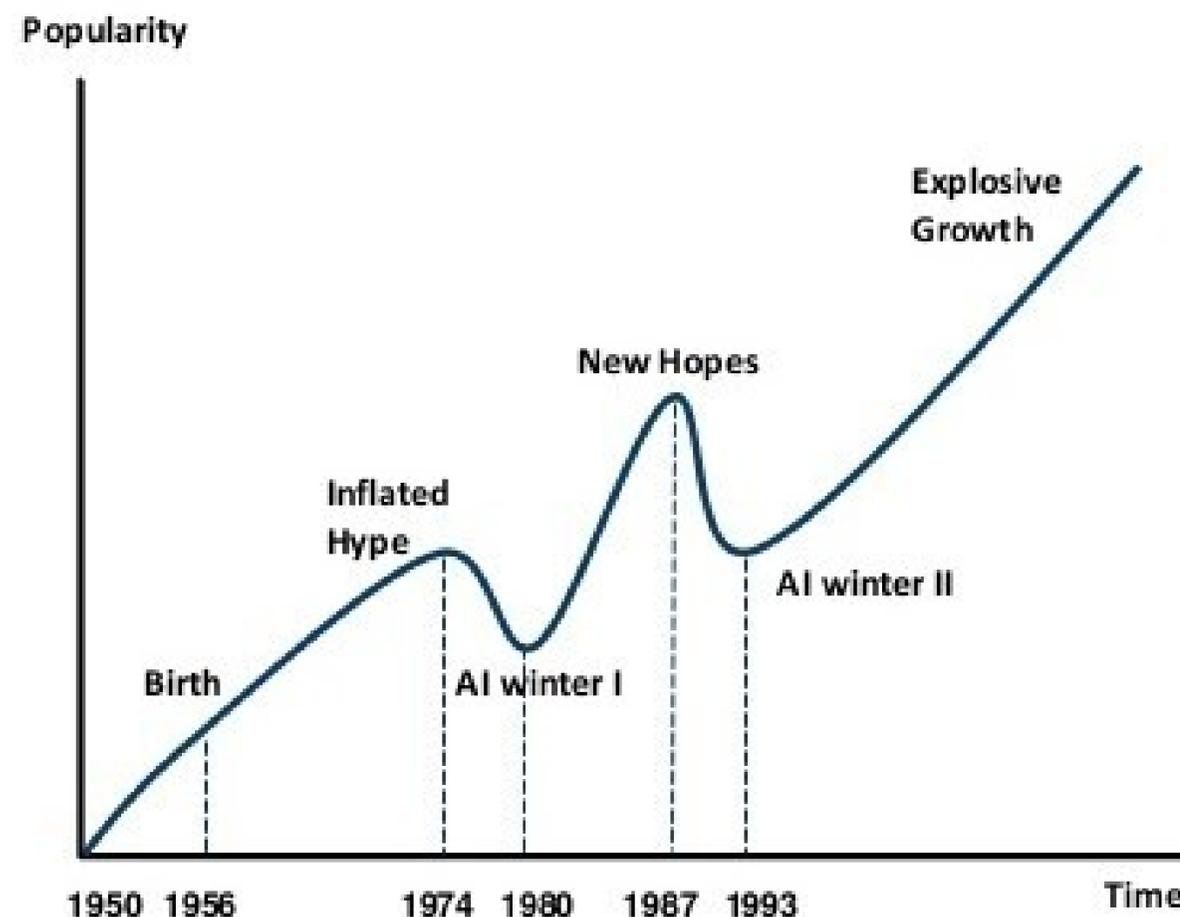


The History of AI – How did we get here?

How did we get here?

Over the last 50+ years, AI popularity and development has ebbed and flowed. Let's take a look into how we got to ChatGPT

AI HAS A LONG HISTORY OF BEING "THE NEXT BIG THING" ...



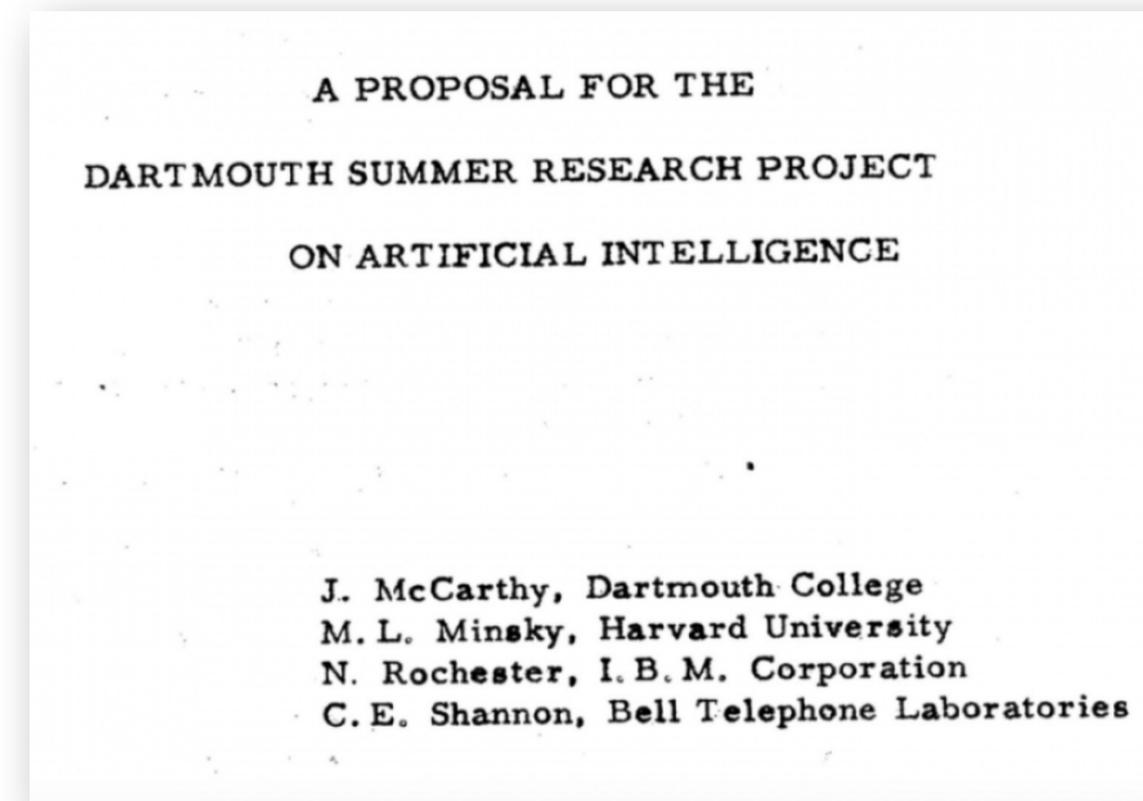
Timeline of AI Development

- **1950s-1960s:** First AI boom - the age of reasoning, prototype AI developed
- **1970s:** AI winter I
- **1980s-1990s:** Second AI boom: the age of Knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1990s:** AI winter II
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

1956 Dartmouth Conference: The Birth of Artificial Intelligence

Significance

- Coining of "Artificial Intelligence": The Dartmouth Conference is where the term "Artificial Intelligence" (AI) was first officially used. John McCarthy, one of the conference's organizers, is credited with coining the term.
- Foundational Moment for AI: This event is often considered the founding moment of AI as a field of study. It was the first time researchers from various disciplines came together to explore the concept of machine intelligence.



Key Participants

- John McCarthy: Known for later developing the LISP programming language, a cornerstone of AI research.
- Marvin Minsky: A cognitive scientist and AI pioneer who would later co-found the MIT AI Lab.
- Claude Shannon: Known as the father of information theory, he brought a deep understanding of communication and information processing to the discussion.
- Nathaniel Rochester: An IBM researcher who contributed to the development of the first computers and was interested in the possibilities of programming machines to think.

1958 Perceptron by Rosenblatt at Cornell

Foundational Model for Neural Networks:

The perceptron, introduced by Frank Rosenblatt in 1958, is one of the earliest models of an artificial neuron, laying the groundwork for the development of modern neural networks. It was inspired by biological neurons and aimed to mimic the way the brain processes information.

Limitations and Influence on Future Research:

While the perceptron was a breakthrough, it had significant limitations, most notably its inability to solve problems that are not linearly separable, such as the XOR problem.

FIG. 1 — Organization of a biological brain. (Red areas indicate active cells, responding to the letter X.)

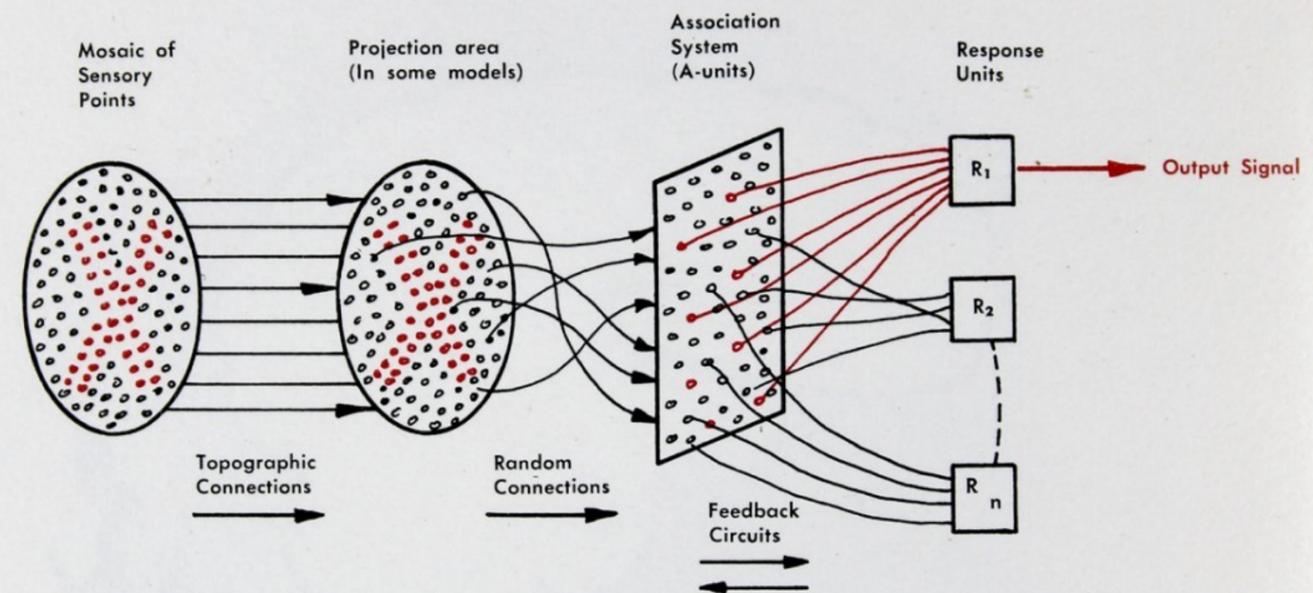


FIG. 2 — Organization of a perceptron.

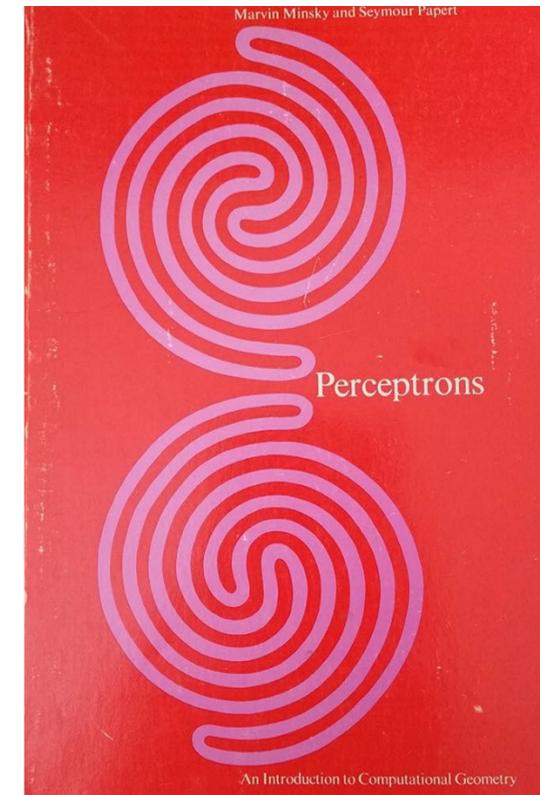
1969 Analysis of Perceptrons Minsky & Papert

Critical Examination of Perceptrons:

In their 1969 book "Perceptrons: An Introduction to Computational Geometry," Marvin Minsky and Seymour Papert critically analyzed the capabilities and limitations of perceptrons, a type of artificial neuron model introduced by Frank Rosenblatt. Their work highlighted fundamental challenges in the perceptron model, particularly its inability to solve problems that are not linearly separable, such as the XOR problem.

Impact on the Field and AI Winter:

The book's critical analysis had a significant impact on the AI community. It led to a period of reduced funding and interest in neural network research, often referred to as the "AI Winter." However, the insights from Minsky and Papert's work also set the stage for the later development of more advanced neural network models, including the backpropagation algorithm in the 1980s, which allowed for the training of multi-layer networks and revitalized interest in the field.



1st AI Winter: 1974 - 1980

An AI Winter refers to a period of reduced funding, interest, and activity in artificial intelligence research. During these times, enthusiasm for AI wanes due to unmet expectations, and progress in the field slows down significantly.

The first AI Winter led to a significant slowdown in research activity, with fewer new developments and innovations. Many AI labs were closed or repurposed, and interest in AI declined across the academic and industrial sectors.

Impact on the field:

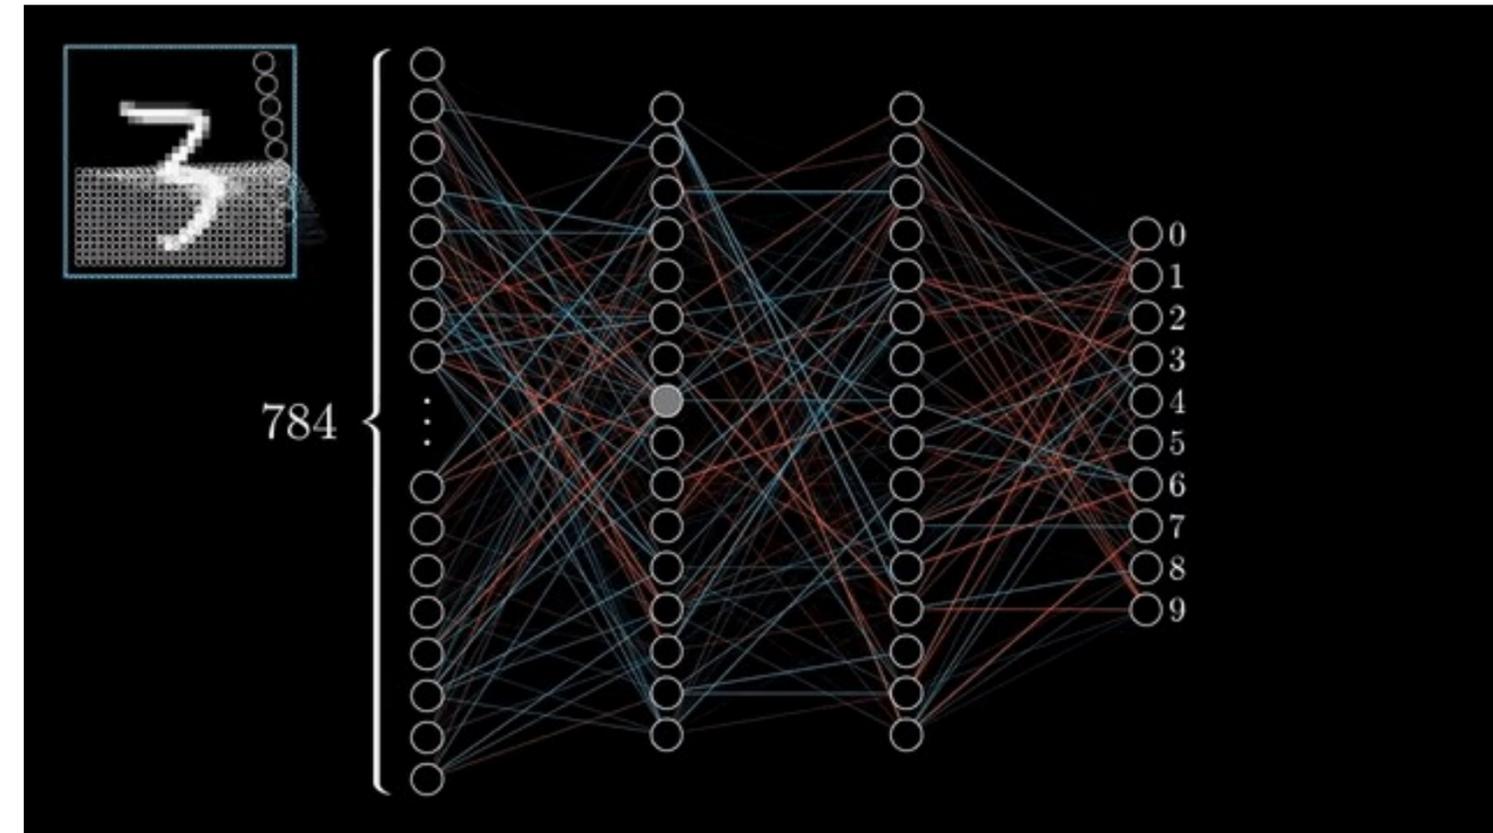
- Many researchers shifted their focus to related fields.
- Some subfields of AI, like machine learning, continued to develop quietly, laying the groundwork for future breakthroughs.
- The AI Winter underscored the importance of managing expectations and the challenges of translating research breakthroughs into practical applications, lessons that continue to influence AI research today.



Backpropagation – Popularized in **1986** by Hinton et.al

Backpropagation, or “*backward propagation of errors*,” is an algorithm used to train neural networks by efficiently computing the gradient of the loss function with respect to the network’s weights. This process allows the network to adjust its weights to minimize errors, thereby improving its performance.

- It uses the chain rule from calculus to compute gradients layer by layer, starting from the output layer and moving backward through the network.
- Before backpropagation, neural networks were primarily limited to linear models with simple learning capabilities.
- The introduction of backpropagation demonstrated that neural networks could approximate any continuous function, given sufficient data and appropriate network architecture.



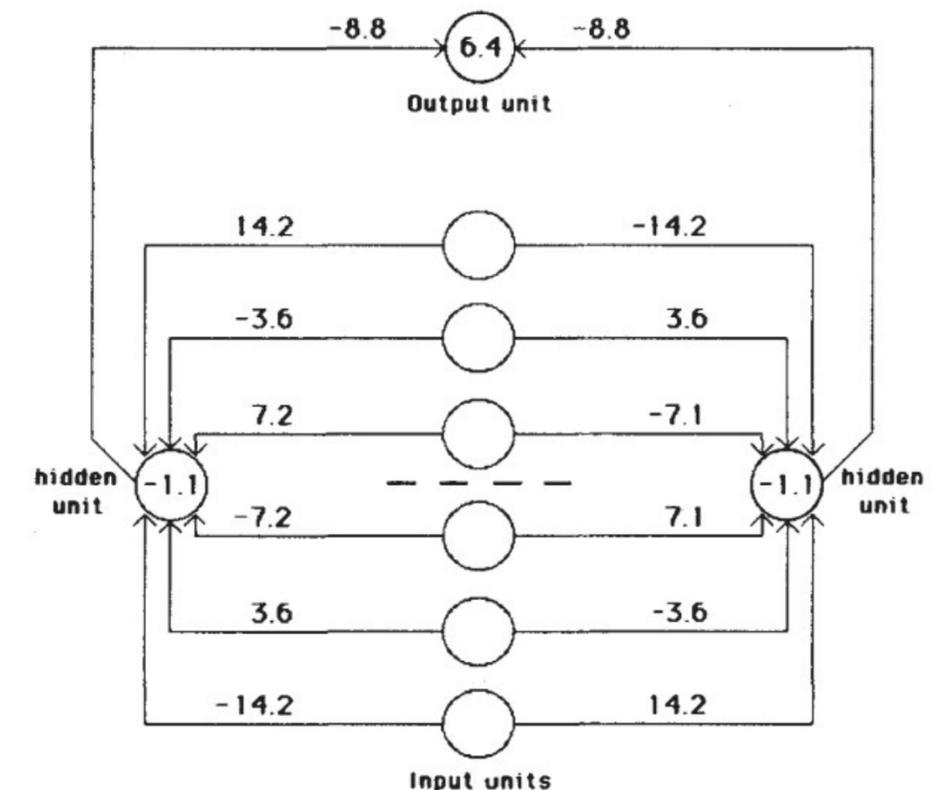
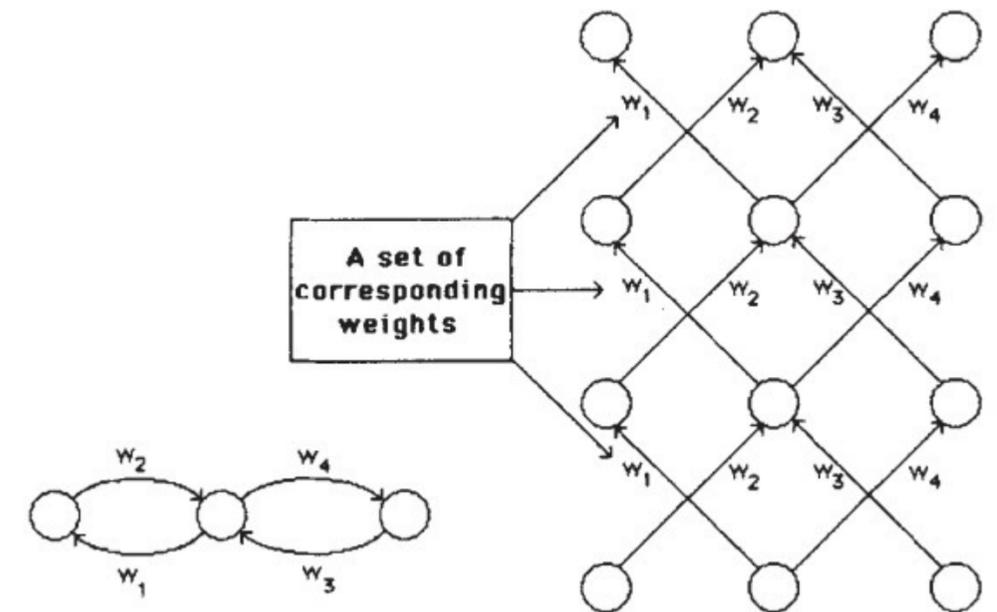
1986 - Multi-layer Perceptron

In the same backpropagation paper from 1986, the multi-layer perceptron was demonstrated to train to classify classes

Learning Internal Representations: The authors demonstrated that MLPs could learn internal representations in hidden layers, which are crucial for capturing complex features in data. This ability to learn hierarchical representations was a significant advancement over previous models.

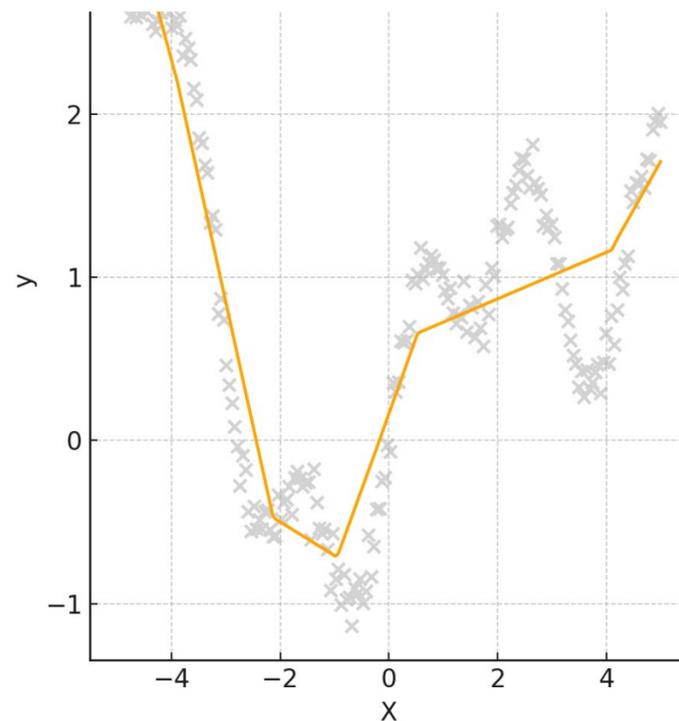
Non-Linear Capabilities: The paper highlighted that MLPs with non-linear activation functions in the hidden layers could solve problems that linear models could not, such as XOR, which was a known limitation of single-layer perceptrons.

Empirical Validation: The 1986 paper provided empirical results showing that MLPs, trained using backpropagation, could successfully perform tasks like pattern recognition, proving the effectiveness of the approach in practical applications.

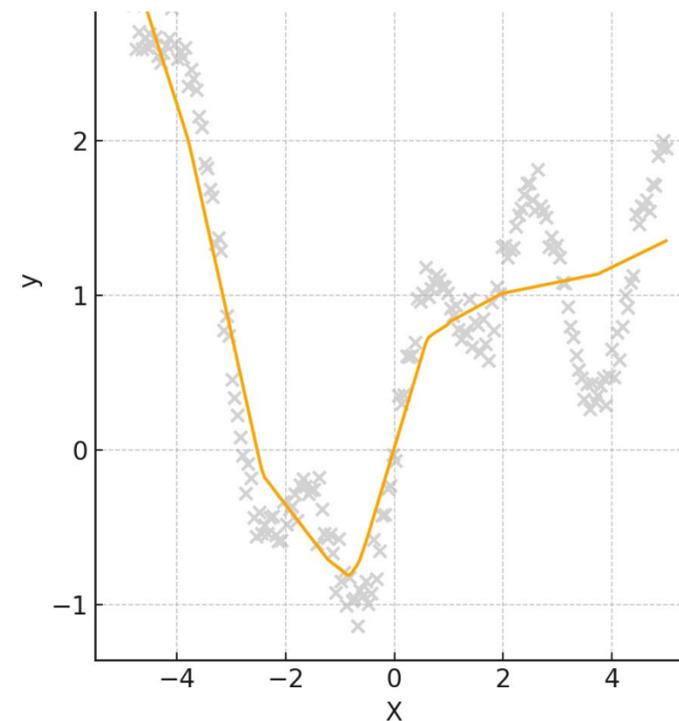


1989 – Universal Function Approximation Theory

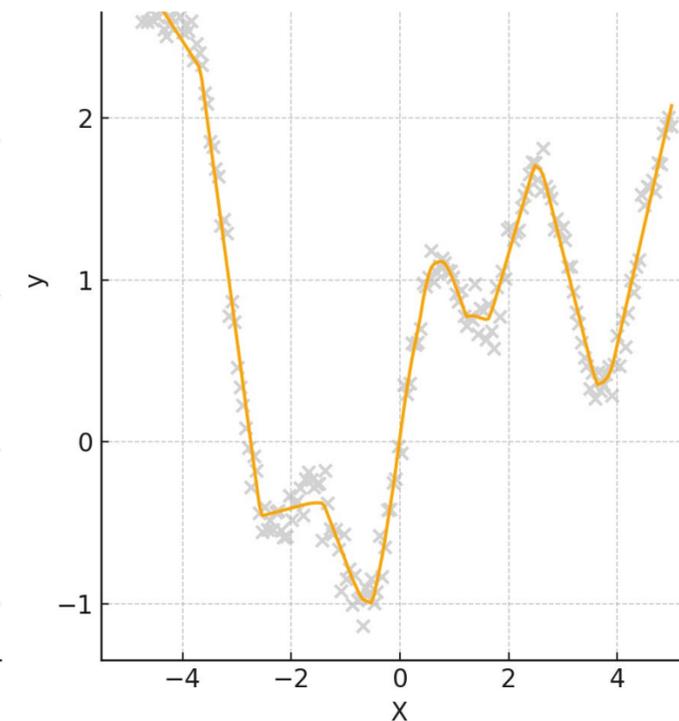
The Universal Approximation Theorem, proposed by George Cybenko in 1989 states that a feedforward neural network with a single hidden layer, containing a finite number of neurons and using a non-linear activation function (like the sigmoid function), can approximate any continuous function, given sufficient neurons in the hidden layer.



Neurons: 5



10



100

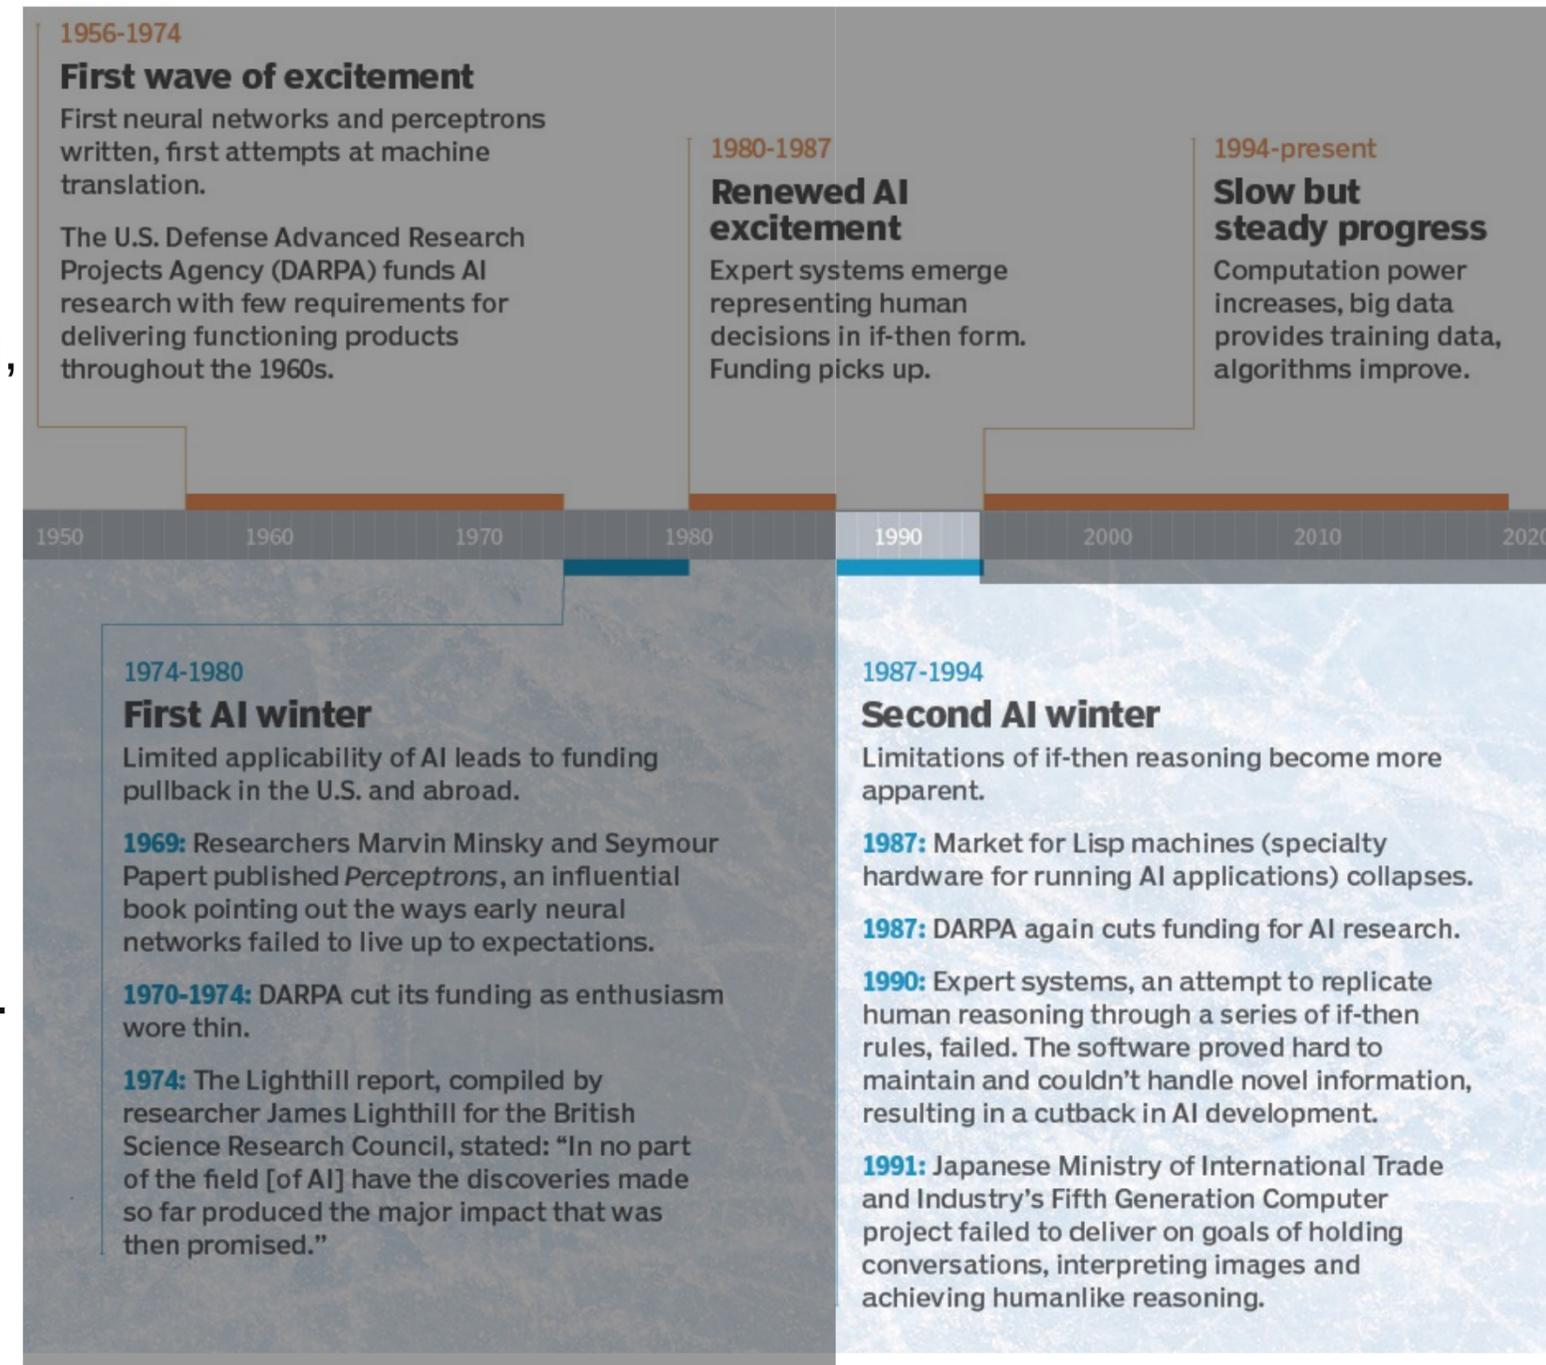
This theorem provided a theoretical foundation for the power of neural networks, showing that even simple networks could, in theory, approximate complex functions, which justified their use in a wide range of applications.

The 2nd AI Winter 1980s-1990s

What led to the second AI winter?

- 1. Unmet Expectations:** Expert systems, initially promising, failed to generalize beyond narrow domains, leading to disillusionment.
- 2. Economic Recession:** The late 1980s recession led to budget cuts in AI research, shifting focus to more immediate technologies.
- 3. Technological Limits:** Insufficient computational power and algorithmic challenges hindered AI development.
- 4. Reduced Interest:** Funding and interest in AI dropped, leading to a slowdown in research and fewer innovations.
- 5. Legacy and Recovery:** Lessons from this period set the stage for the AI resurgence in the late 1990s, paving the way for modern AI advancements.

AI winters freeze progress



2012 AlexNet – Computer Vision and GPUs Dominate AI

AlexNet was a groundbreaking neural network that significantly advanced the field of AI by demonstrating the power of deep learning and GPU acceleration

Revolutionary Performance:

AlexNet, developed by Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, won the 2012 ImageNet competition with a large margin.

GPU Acceleration:

AlexNet utilized GPUs (specifically NVIDIA GPUs) to dramatically reduce the training time for deep neural networks. This approach allowed the model to be trained on a much larger scale.

Catalyst for AI Growth:

The success of AlexNet spurred widespread adoption of deep learning and GPUs in AI research and industry.

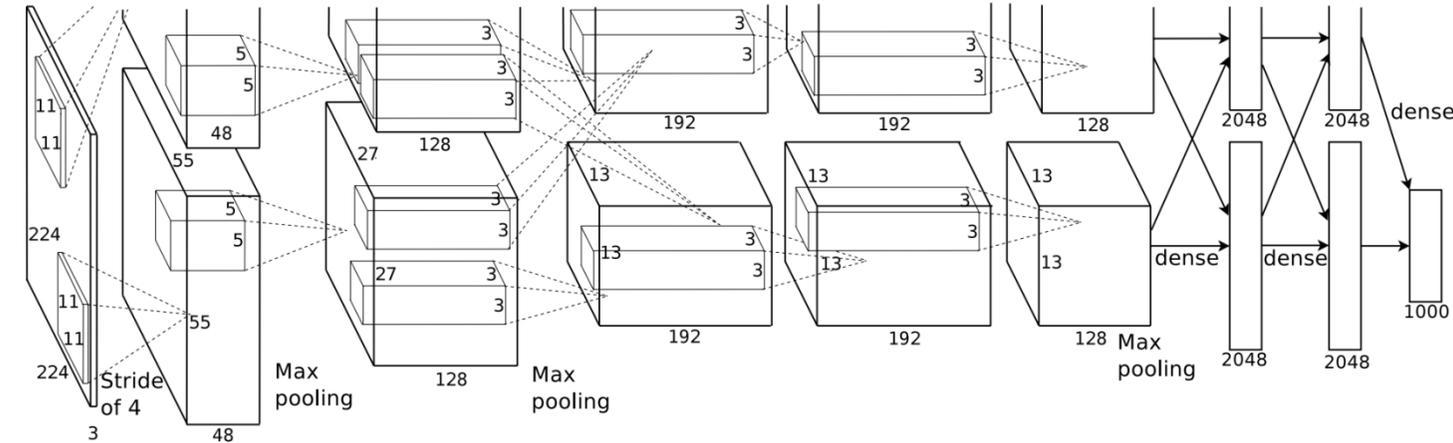
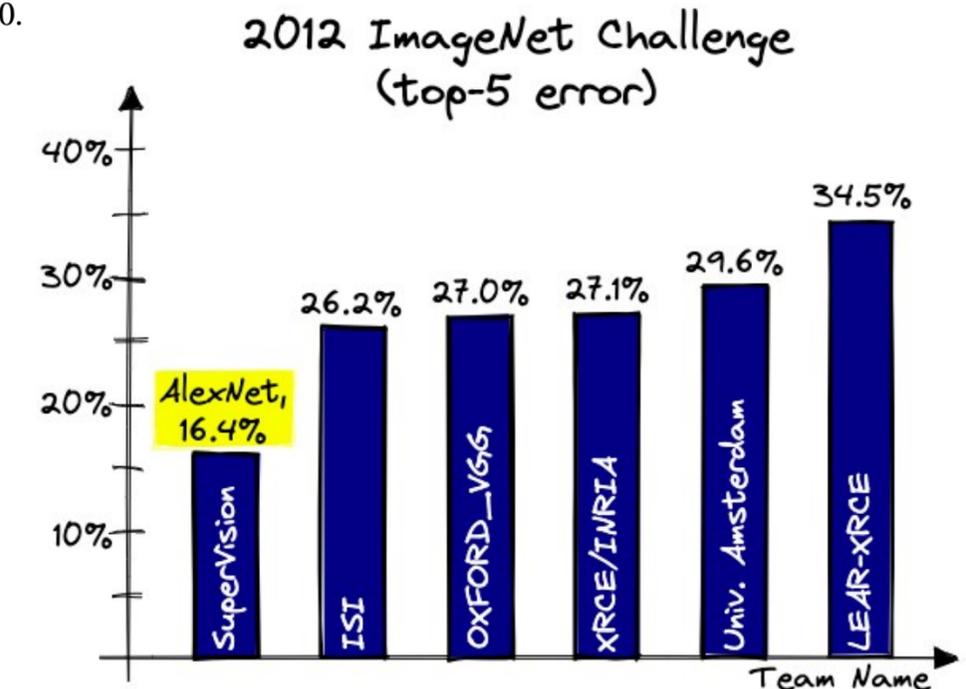


Figure 2: An illustration of the architecture of our CNN, explicitly showing the delineation of responsibilities between the two GPUs. One GPU runs the layer-parts at the top of the figure while the other runs the layer-parts at the bottom. The GPUs communicate only at certain layers. The network's input is 150,528-dimensional, and the number of neurons in the network's remaining layers is given by 253,440–186,624–64,896–64,896–43,264–4096–4096–1000.

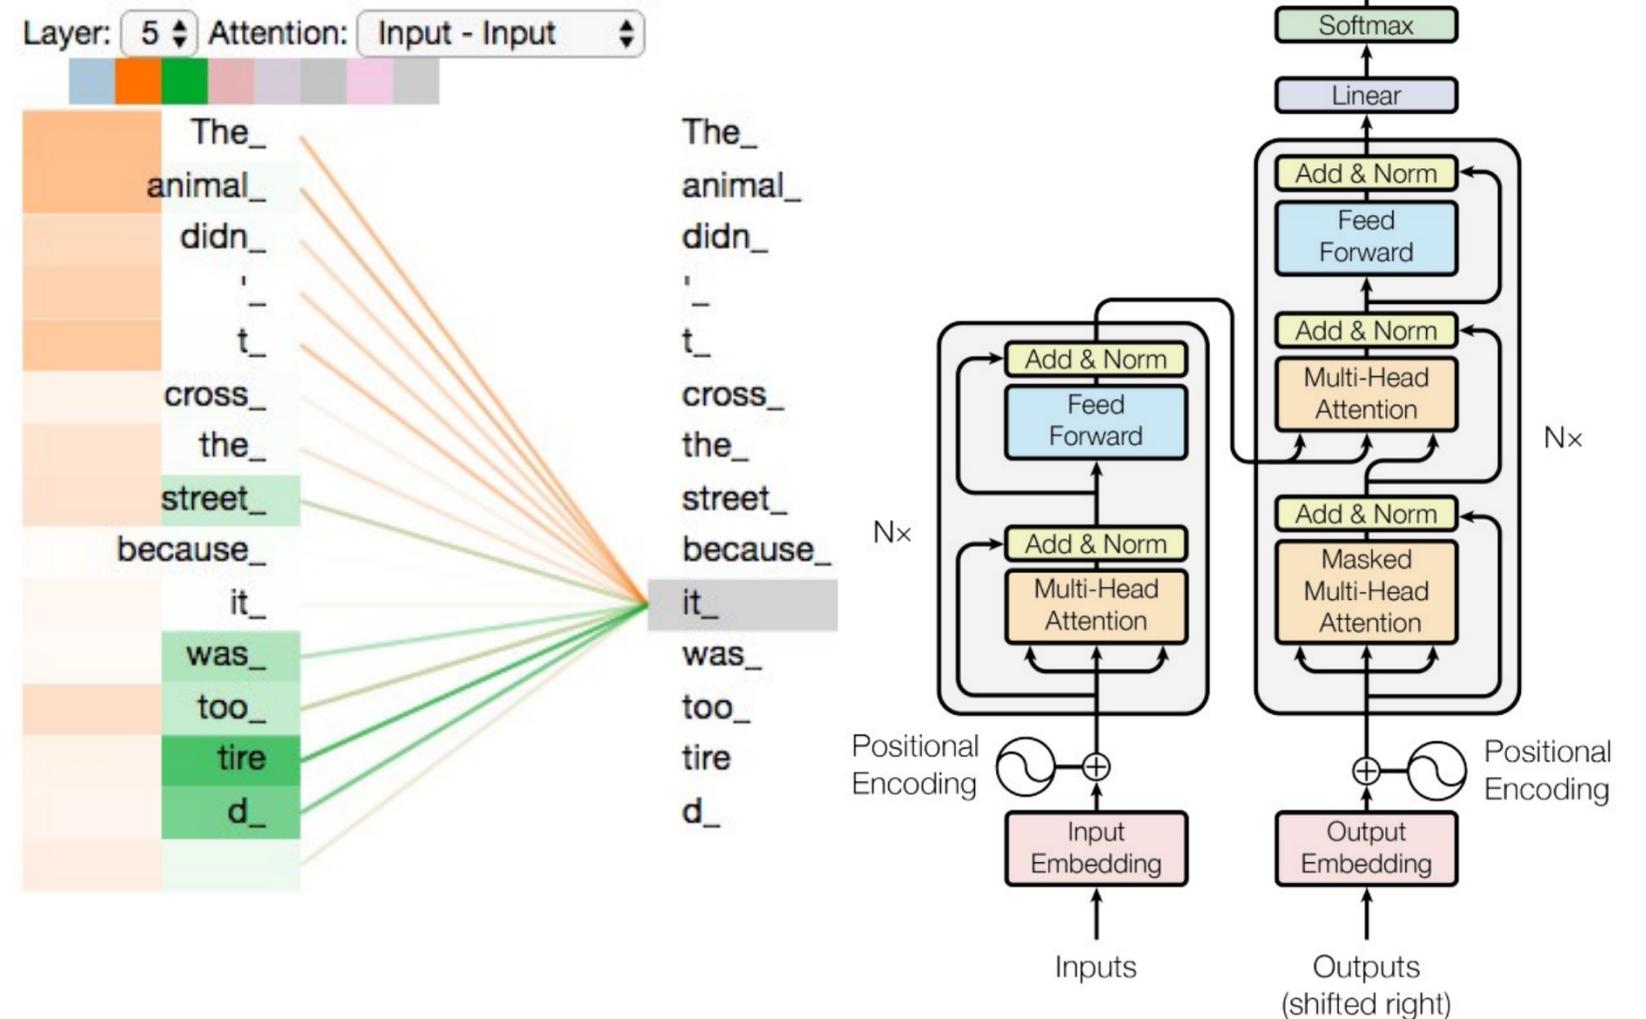


2017++: Attention is all you Need – Everything Changes

In the 2017 paper, Attention is all you need, a new deep learning architecture, the Transformer was introduced.

This model, and its variants: BERT and GPT have been found to solve traditional natural language problems and are capable of generating complex semantic outputs

These text generation models, together with the image generation models constitute Generative AI and will be the focus moving forward.



The Risks of GenAI

Risks of GenAI

For all its technological wonder, GenAI poses a serious challenge some fundamental parts of modern society:

- Trust and Information
- Digital records of events and people
- Jobs



Deepfakes

Deepfakes are AI-generated synthetic media, typically videos or images, where someone's likeness is realistically replaced with another's, often leading to convincing yet false representations.



These can be used for entertainment, but also to spread misinformation, create fake news, or manipulate public perception, posing significant ethical and security risks.

Job losses due to GenAI

Automation of Routine Tasks:

AI is automating repetitive jobs like data entry and manufacturing, leading to job reductions in these areas.

Impact on Skilled Jobs:

AI is increasingly capable of handling complex tasks, potentially displacing roles in fields like law, medicine, and finance.

New Job Opportunities:

While some jobs are lost, AI also creates new roles in areas like AI development and data science, but workers need reskilling to adapt.



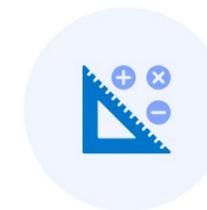
The careers most likely to be affected by AI



Software Development



IT Operations and Helpdesk



Mathematics



Information Design & Documentation



Legal

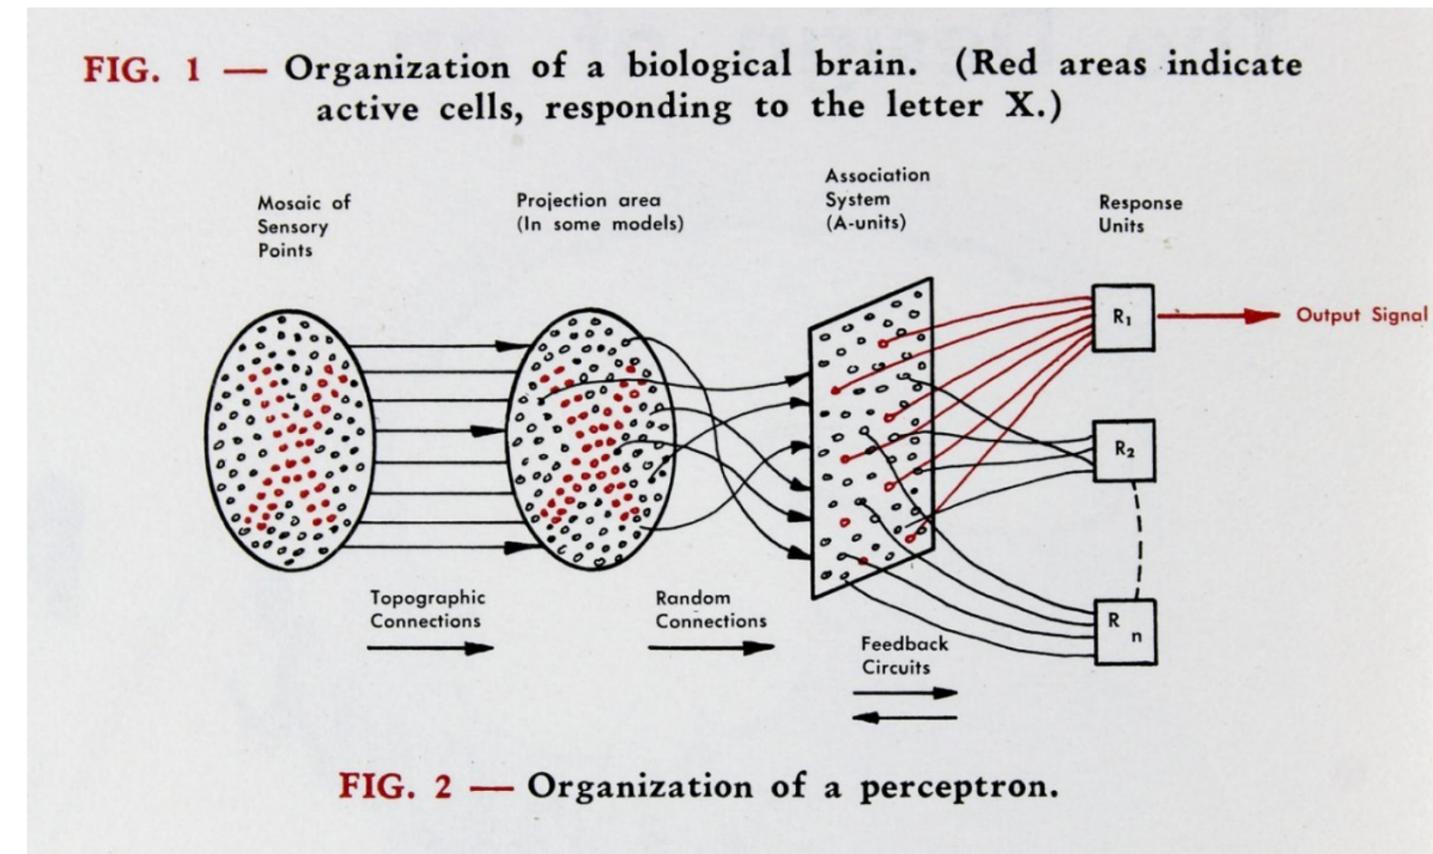


Accounting

Wrap Up

History and Risks of GenAI

- Today we introduced Generative AI (GenAI)
 - We walked through the history of important developments in AI on the road to GenAI
 - We introduced some of the risks of GenAI in modern society
-





Thank you!