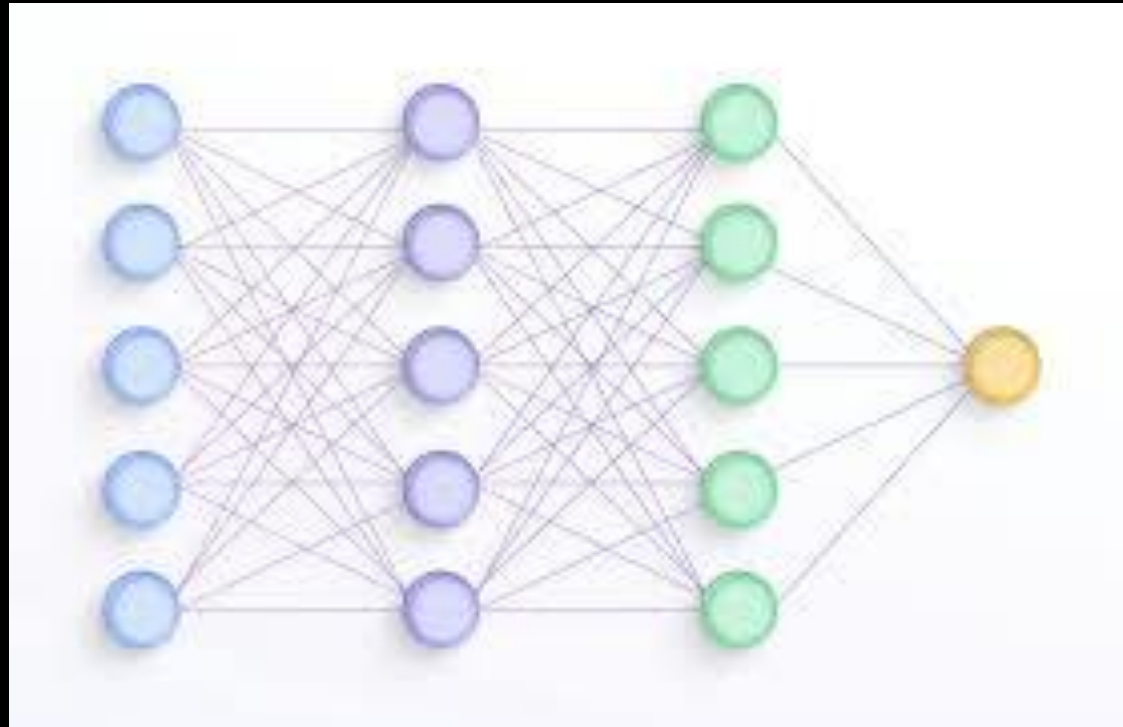


Neural Network lab book



By: Ryan

Problem

Humans have caused a lot of our own problems such as war, pollution, and disease.

Mission

Learn how to build our own robots so they can solve our complex problems.



What is a neural network?

ChatGPT said:

What is a Neural Network (in AI)?

A **neural network** is a computer system **inspired by the human brain** that helps machines **learn from data**.

It's the **core technology** behind things like Large Language Models (LLMs), image recognition, voice assistants, self-driving cars, and more.

How does it work (in simple terms)?

Imagine a **giant network of digital “neurons”**:

- Each "neuron" is like a **tiny calculator**.
- Neurons are organized in **layers**:
- **Input layer**: Takes in raw data (e.g., words, pixels, numbers).
- **Hidden layers**: Where the magic happens—neurons pass information around and learn patterns.
- **Output layer**: Produces the final result (like a word, a category, a prediction).

The network adjusts itself over time by **learning from mistakes**, using something called **backpropagation** and **training data**.

Example:

Let's say you're training a neural network to recognize pictures of cats:

1. **Input**: An image of a cat.
2. The network processes the pixels through many layers.
3. **Output**: “Cat” (if it learned well!).

The more images you show it, the better it gets.

Why it's important:

Neural networks are what make **deep learning** and **modern AI** possible. **They're especially good at tasks that are hard to write rules for, like:**

- Understanding language (like me!)
- Recognizing faces
- Translating languages
- Playing games (like AlphaGo)

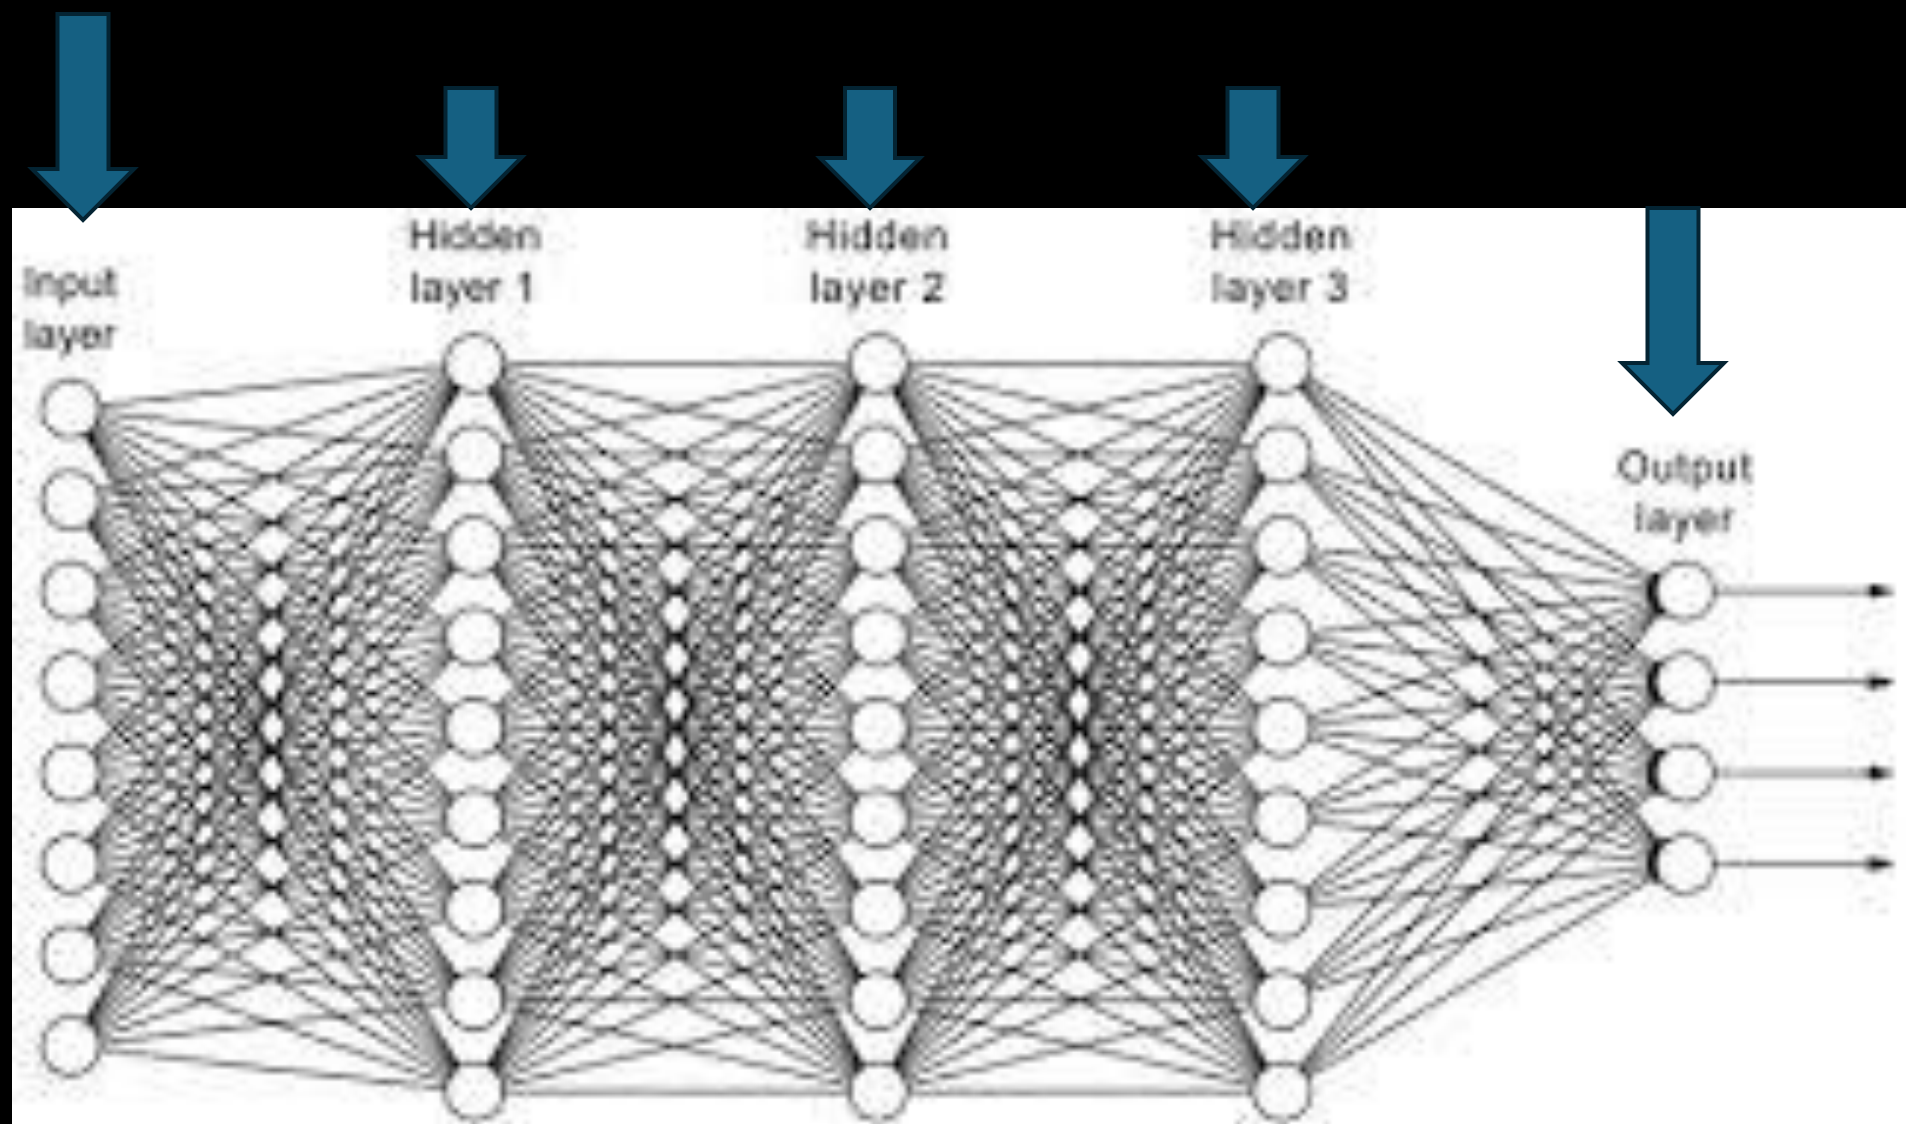
Copilot said:

A **neural network** is a type of machine learning model inspired by the structure and functioning of the human brain. It consists of layers of interconnected nodes, or "neurons," that process and transmit information.

Here's a breakdown:

- **Structure:** Neural networks have an input layer (where data enters), one or more hidden layers (where computations happen), and an output layer (where results are produced). Each connection between neurons has a weight that determines its importance.
- **Functioning:** Data flows through the network, and each neuron applies a mathematical function (called an activation function) to decide whether to "fire" and pass information to the next layer.
- **Learning:** Neural networks learn by adjusting the weights of connections during training. This process, often called "backpropagation," minimizes errors in predictions by comparing outputs to known results.

Neural networks are the foundation of many AI applications, including image recognition, speech processing, and even self-driving cars.



Copilot said:

A neural network is a type of computer system inspired by the structure and function of the human brain. It's designed to recognize patterns, process data, and make decisions based on input it receives.

At its core, a neural network consists of layers of interconnected nodes, or "neurons." These layers typically include:

- **Input Layer:** Takes in raw data (e.g., images, text, or numbers).
- **Hidden Layers:** Process the data by performing calculations and learning patterns. These layers contain most of the network's complexity.
- **Output Layer:** Produces the result, such as a prediction or classification.

Neural networks are a **key part of machine learning and artificial intelligence (AI)**. They **power technologies like voice assistants, facial recognition, and autonomous vehicles**. By "training" a neural network with data, it can learn to improve its accuracy over time. It's a fascinating system that mimics how humans learn and adapt! If you're interested, I can dive deeper into concepts like how training works or specific types of neural networks, like convolutional or recurrent ones.

Copilot said:

How Neural Networks Learn

Neural networks learn by adjusting the "weights" of the connections between neurons during a process called **training**.

Here's how it works:

1. **Forward Pass:** Input data passes through the network, layer by layer, producing an output prediction.
2. **Loss Function:** The network calculates how far off its prediction is from the actual answer (e.g., identifying a cat instead of a dog in an image).
3. **Backpropagation:** It uses a process called backpropagation to send this error backward through the network and adjust weights to minimize errors.
4. **Repetition:** This process repeats thousands, sometimes millions, of times, improving the model's accuracy.

Types of Neural Networks

Neural networks come in various forms, each suited for specific tasks:

1. **Feedforward Neural Networks:** The simplest type, where data moves one way (forward) through the layers. Used for basic tasks like image recognition.
2. **Convolutional Neural Networks (CNNs):** Specialized for image and video processing. They excel at detecting features like edges and patterns in images.
3. **Recurrent Neural Networks (RNNs):** Designed to handle sequential data, like text or time series. They remember previous inputs, making them great for tasks like language modeling or speech recognition.
4. **Generative Adversarial Networks (GANs):** Consist of two networks competing against each other to create data that mimics the original data, like generating realistic images or videos.
5. **Transformers:** Powerhouses behind modern AI models, particularly in natural language processing. They're responsible for models like GPT and others that analyze and generate human-like text.

Neural networks are revolutionizing industries like healthcare (predicting diseases), entertainment (movie

Research What is an LLM (Large Language Model)?

ChatGPT said:

An LLM is a type of artificial intelligence trained to understand and generate human language.

Think of it as a super advanced **text prediction machine**—but instead of just guessing the next word, **it can write essays, explain complex ideas, have conversations, generate code, translate languages, and much more.**

How it works (simplified):

- **Training:** It's **trained on huge amounts of text from books, websites, articles, etc.**
- **Learning patterns:** It learns patterns in how humans use language (grammar, style, facts, reasoning).
- **No consciousness or understanding:** It doesn't *know* things the way humans do—it's all about patterns and probabilities.

Examples of LLMs:

- **GPT-4** (that's me!)
- **Claude** by Anthropic
- **Gemini** by Google DeepMind
- **LLaMA** by Meta
- **Mistral, Command R**, and more

What can LLMs do?

- Chat and answer questions (like this!)
- Summarize text
- Write emails, stories, scripts
- Translate languages
- Write and debug code
- Power chatbots, search engines, virtual assistants, etc.



Copilot said:

Large Language Models (LLMs) like me are **powered by a type of deep learning architecture called transformers**, which are excellent at processing sequences of data—like text—while understanding context, meaning, and patterns. The key idea is that these **models learn relationships between words, phrases, and concepts in language by analyzing enormous datasets, often consisting of books, articles, and online content.**

Here are a few fascinating aspects of LLMs:

- **Training:** LLMs are trained on billions, sometimes trillions, of text examples. They predict what comes next in a sequence based on context, learning patterns of human-like language over time.
- **Multitasking Ability:** Once trained, LLMs can perform many tasks beyond basic language understanding, like answering questions, drafting emails, writing stories, creating code, or even solving math problems.
- **Fine-Tuning:** Models can be adjusted for specific tasks, making them even more efficient. For example, an LLM can be fine-tuned to focus solely on legal language or medical terminology.
- **Applications:** LLMs are used in chatbots, content creation tools, search engines, translation systems, educational aids, and much more.

Despite their capabilities, LLMs do have limitations. They don't truly "understand" language like humans do—they rely on statistical probabilities to guess the next word. This can sometimes lead to mistakes or responses that sound confident but are factually incorrect. If you want, I can explain more technical aspects, like how "attention mechanisms" work in transformers, or share real-world examples of LLM applications. What intrigues you the most about them?

Research "What is a transformer?"

ChatGPT said:

What is a Transformer (in AI)?

A **transformer** is a **type of neural network architecture** that completely changed the game for how machines understand and generate language (and now, images, audio, etc.).

It was introduced in 2017 by researchers at Google in the paper:

"Attention Is All You Need"

Before transformers, models like RNNs and LSTMs were used — but transformers blew past them in performance, speed, and scalability.

Core idea: Self-Attention

The magic sauce is called **self-attention** (or "attention mechanism").

It lets the model:

- **Look at all the words in a sentence at once**, not just one at a time.
- Focus on the **most relevant words**, even if they're far apart in the sentence.

For example:

"The cat, which was hiding under the table, suddenly jumped."

To understand "*jumped*", the model needs to know "**cat**" is the subject.

Self-attention lets it **zoom in on that** — no matter how far apart the words are.

How a Transformer is Built (at a glance):

- **Input embeddings:** Words turned into numbers (vectors)
- **Positional encoding:** Adds info about word order (since the model sees everything at once)
- **Encoder layers:** Process the input (used in translation, summarization, etc.)
- **Decoder layers:** Generate output (used in text generation, like ChatGPT)
- **Multi-head attention:** Looks at information from different "angles" at once

Famous models based on transformers:

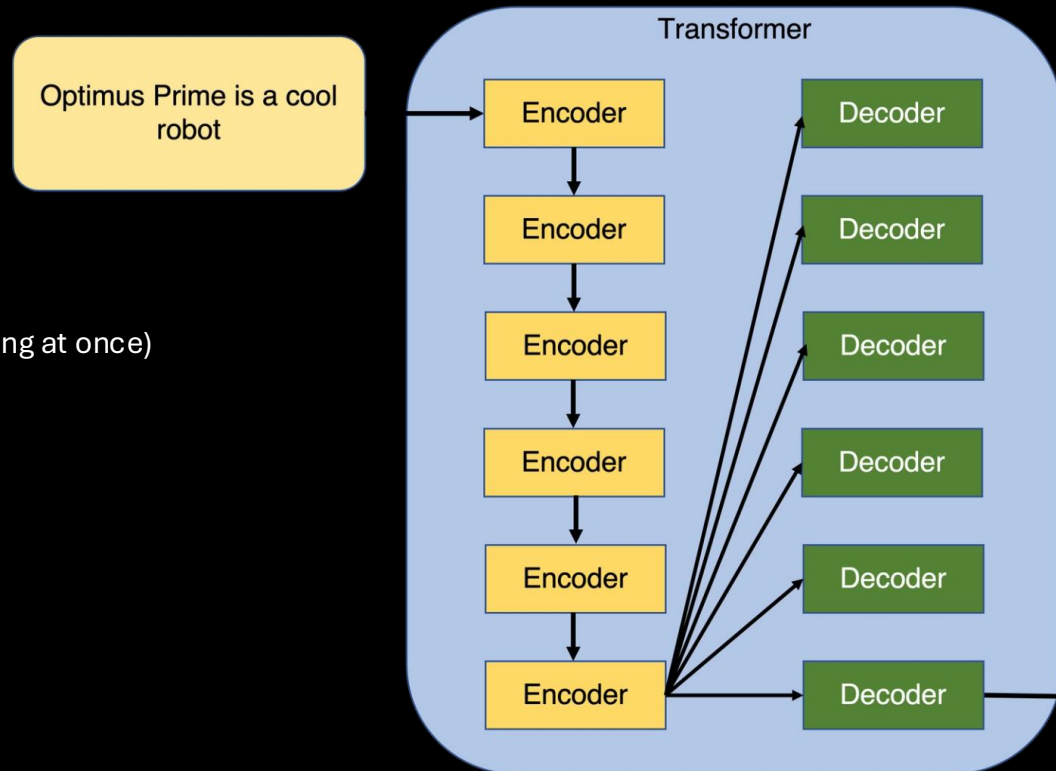
- **GPT (OpenAI)** — Just the decoder half (focused on generation)
- **BERT (Google)** — Just the encoder half (focused on understanding)
- **T5 / PaLM / LLaMA / Claude / Gemini / etc.** — All transformer-based

Why it's a big deal:

- **Much faster to train** (no need to process sequences step-by-step)
- **More accurate** on language tasks
- **Scales well** — transformers power today's largest and most capable AIs

So basically:

Transformers = **the architecture that made modern AI possible.**



Copilot said:

In AI, a **transformer** is a type of deep learning model that has revolutionized the field of natural language processing (NLP). Introduced in a 2017 paper titled *Attention Is All You Need* by Vaswani et al., **transformers use a mechanism called self-attention, allowing them to process words in a sentence in parallel rather than sequentially.**

Here's why they're important:

- **Parallel Processing:** Unlike older models that processed words one at a time, transformers can analyze an entire sentence at once.
- **Self-Attention:** They weigh the importance of different words relative to each other, helping them understand context better.
- **Scalability:** They have enabled large-scale AI models like GPT and BERT, which power many advanced language-based applications today.

Essentially, transformers have been the backbone of modern AI advancements in generating, understanding, and translating text.

Timeline of AlphaGo!!

Here's the timeline:

2015 – First pro win:

- AlphaGo **beat Fan Hui**, a European Go champion, **5 games to 0**.
- This was the **first time a computer had ever beaten a pro human player** without handicaps.

2016 – Big match vs. Lee Sedol:

- AlphaGo played against **Lee Sedol**, one of the world's top Go players.
- The match was **best of 5**.
- AlphaGo **won 4-1**.
 - Lee Sedol won **Game 4**, which became legendary.
 - It was the only game AlphaGo lost during that series.

Summary:

- AlphaGo beat a professional human **right away**.
- Against a world champion, it won **4 out of 5 games**.
- After that, newer versions (like AlphaGo Zero) got even stronger—**training without any human data at all**.



Research How does AlphaGo Zero train?

Copilot said:

AlphaGo Zero trains itself using a technique called **reinforcement learning**, where it plays games against itself and learns from the outcomes. Unlike previous versions of AlphaGo, which relied on human gameplay data, AlphaGo Zero starts with no prior knowledge beyond the basic rules of Go.

Here's how it works:

- **Self-Play:** AlphaGo Zero plays millions of games against itself, making moves based on a combination of deep neural networks and a powerful search algorithm.
- **Neural Network Updates:** After each game, the system updates its neural network to improve its ability to predict moves and game outcomes.
- **Monte Carlo Tree Search (MCTS):** It uses MCTS to evaluate possible moves more efficiently, focusing on the most promising ones rather than considering every possible move.
- **Continuous Improvement:** Over time, AlphaGo Zero refines its strategies, surpassing human-level play and even previous AI versions.

This approach allows AlphaGo Zero to develop superhuman skills without human intervention, making it one of the most advanced AI systems in game-playing history.

Research "How will neural networks shape the future?"

Copilot said:

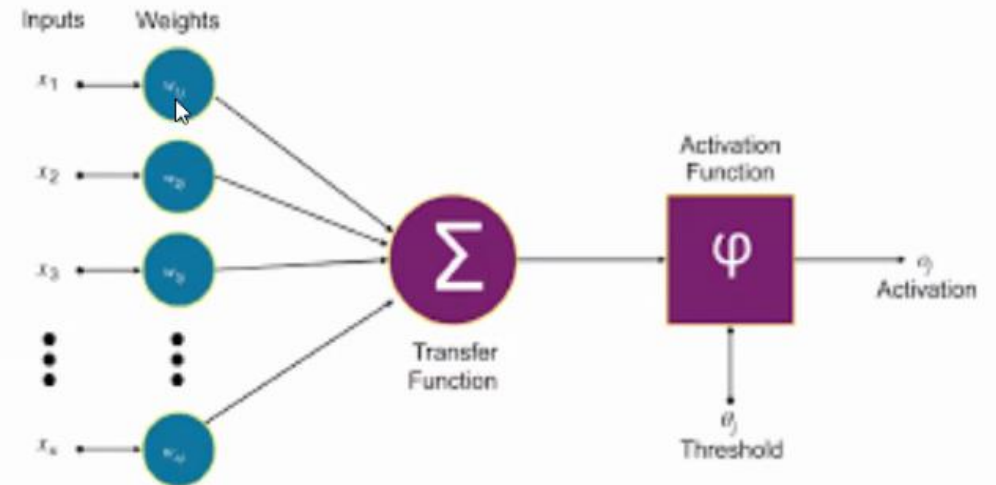
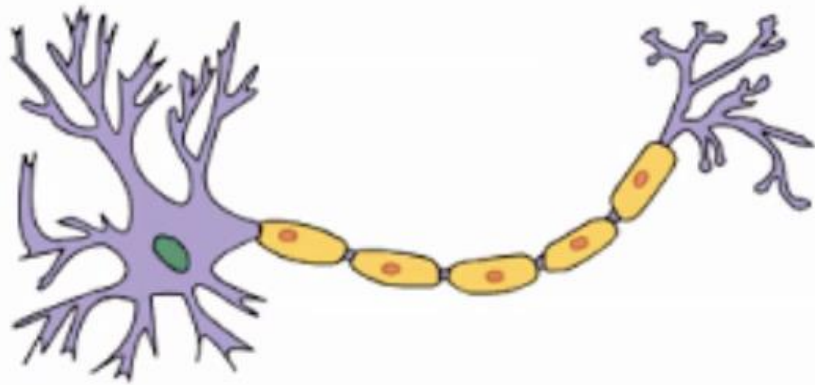
Neural networks are already reshaping industries and will likely continue to revolutionize the future in ways we can't yet fully predict. Here are a few keyways they're set to make an impact:

- **Smarter AI Assistants** – Virtual assistants and chatbots will become more intuitive, able to grasp context, tone, and even emotions, making interactions feel more natural.
- **Breakthroughs in Healthcare** – From early disease detection to personalized medicine, neural networks will improve diagnostics and treatment strategies with unparalleled accuracy.
- **Self-Driving Technology** – Autonomous vehicles will rely heavily on neural networks to process vast amounts of sensor data, improving safety and efficiency on the roads.
- **Creative Content Generation** – Whether in art, music, writing, or even game development, neural networks will aid human creativity and automate various aspects of content creation.
- **Advanced Scientific Discovery** – Simulating molecular interactions for drug discovery or predicting climate patterns with neural networks will unlock new scientific frontiers.
- **Personalized Education & Learning** – AI-powered learning platforms will adapt lessons to individual students, making education more effective and accessible worldwide.

While neural networks will create groundbreaking opportunities, they'll also raise ethical concerns, including bias in AI, data privacy, and job displacement. Striking a balance between innovation and responsible development will be key.

How can a computer think like a brain?

By programming artificial neurons!



Bunny

10 Image Samples

Webcam

Upload

Dog

10 Image Samples

Webcam

Upload

Add a class

Training

Model Trained

Advanced

Preview

Export Model

Input ☒ ON

File

Choose images from your files,
or drag & drop here

Import images from
Google Drive




Output

Bunny

100%


Dog


Bunny 


10 Image Samples


Webcam


Upload




















Dog 


10 Image Samples


Webcam


Upload


















 Add a class

Training


Model Trained

Advanced 

Preview

 Export Model

Input ☒ ON

File 


Choose images from your files,
or drag & drop here


Import images from
Google Drive



Output

Bunny

100%

Dog

Confusion Matrix To The Rescue!

Answer

		Prediction	
		Healthy	Sick
Answer	Healthy	98	0
	Sick	2	0

Oops! Not so good after all...

Bunny

10 Image Samples

Webcam

Upload

Dog

10 Image Samples

Webcam

Upload

Add a class

Training

Model Trained

Advanced

Preview

Export Model

Input ☒ ON

File

Choose images from your files,
or drag & drop here

Import images from
Google Drive



Output



Robot Learner's Rulebook



Use **LOTS** of examples




Have a lot of **VARIETY**



DON'T REPEAT










CLEAN and CLEAR


Bunny 

12 Image Samples

Webcam

Upload











Dog 

12 Image Samples

Webcam

Upload



 Add a class

Training


Train Model

Advanced 

Preview

 Export Model

Input ☐ ON

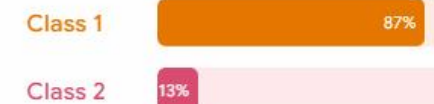
File 


Choose images from your files,
or drag & drop here

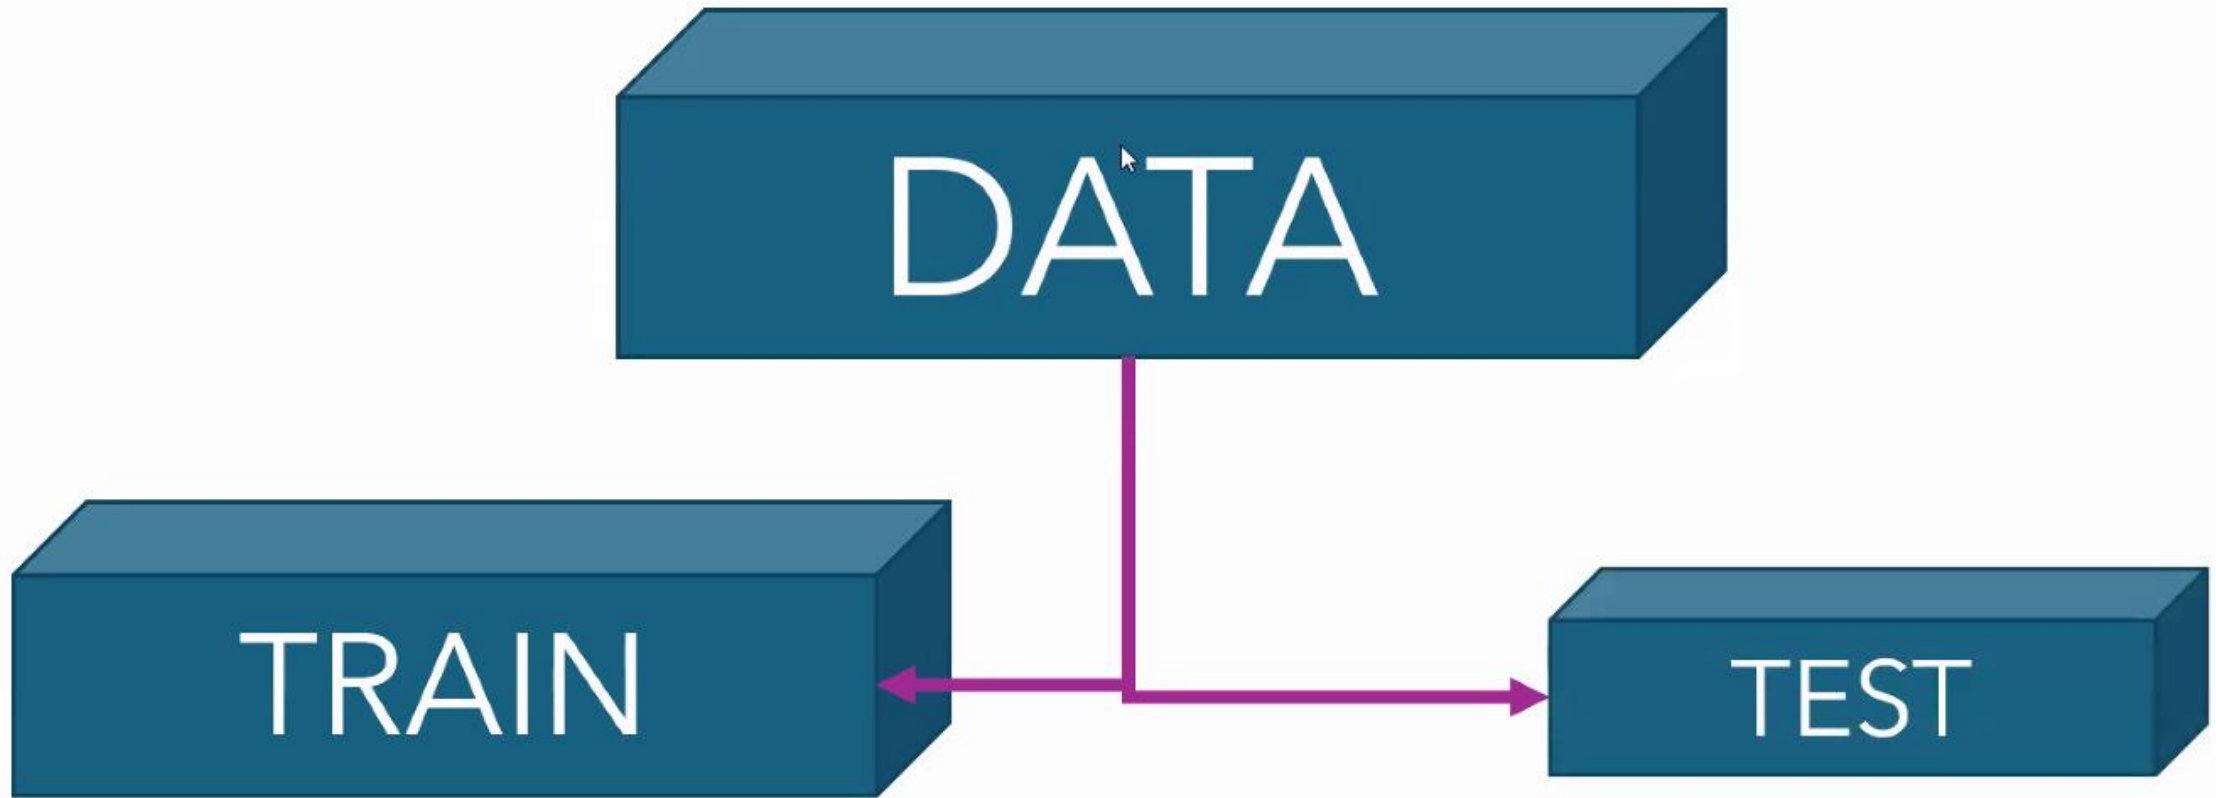

Import images from
Google Drive



Output



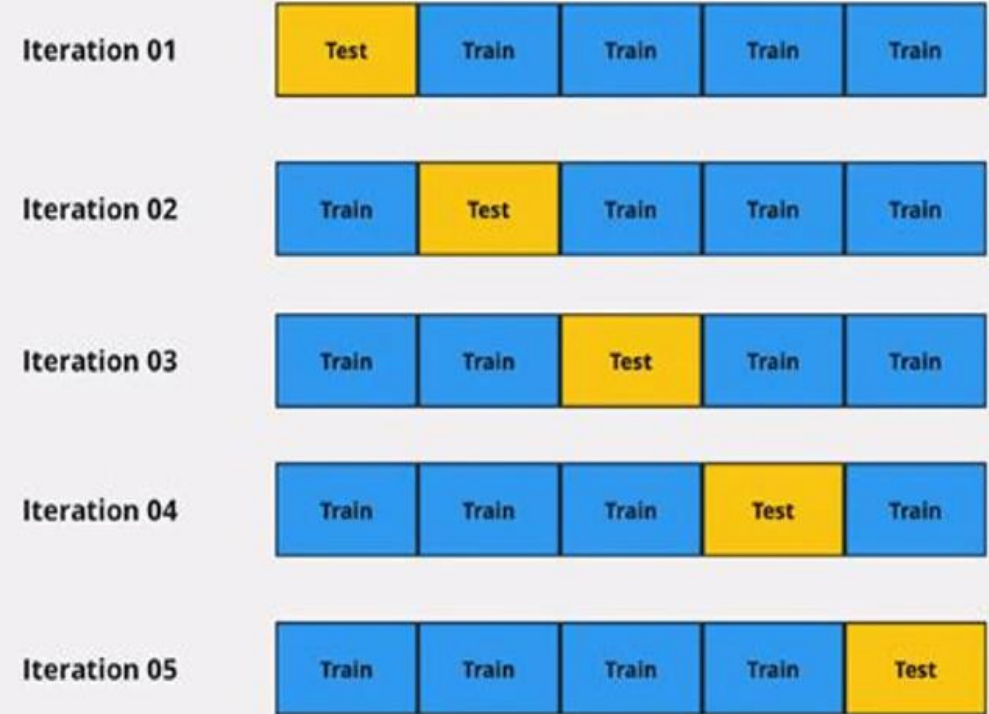
How can we test our robot many times?



How can we test our robot many times?

- 1) Put your data into groups
- 2) Use 1 group to test your model
- 3) Use the other groups to train your model
- 4) Choose another test group and repeat!

K-Fold Cross Validation



Testing the AI to the top of its ability

	Test 1	Test 2
Dog:		
Score:	100%	98%
Score:	100%	99%
Bunny:		
Score:	99%	100%
Score:	100%	100%

Key Lesson: Accuracy Isn't Everything!

It's not about *how* many answers your robot gets right; its about *which* mistakes it makes!

A smart robot should try to get the important things right



And when wrong, it should try not to cause problems





Phoenixes are bright red birds that can live for thousands of years. When it is about to die it bursts into a scorching hot fire. From the ashes of the fire the phoenix is reborn. Phoenixes also have dark black beaks. The author of Harry Potter; J.K. Rowling added a phoenix in her book. The name of the phoenix is Fawkes. J.K. Rowling named the phoenix Fawkes after Guy Fawkes was associated with explosives and a plot to destroy London's parliament and the assassination of King James I. It is hard to tame phoenixes and only a couple wizards in Harry Potter have done it such as Albus Dumbledore.

PHOENIXES RULE!!!!!!



Caution stay 5 kilometers away!!

This is the Datre (Pronounced: Dat-er). It is the foulest of foul beasts. The Datre is a dragon. It has the power to make people and things do what it wants. This beast was made in a lab and is born when a dragon egg is hatched from under a tree. It has powers nobody has even dreamed of. The Datre is found in Greenland, Iceland, Russia, and in northern Canada. The death count of this beast is unknown.



Caution stay 5 kilometers away!!

This dragon may look adorable; however, it can breathe fire up to 5 football fields! The longest it has been in captivity is currently 10 minutes as it can destroy everything, even carbon nanotubes, the strongest substance known to man. It is only about 25 cm long as an adult. When it is born it is only 3 cm long. And it lives for around 10 years. It is found in Asia, Europe, and North America. The name of the species is called the Emberclaw.



Dangerous! Do not go near!!!!

This is the CBJF otherwise known as the Cyborg Jellyfish. This jelly fish was made in a lab on the planet PA-99-N2 in the Andromeda galaxy. The scientists that made this jellyfish were killed because of it. It now lives in the ocean attacking vessels that come across it. Several attempts have been made to kill it however all have failed as the jellyfish set off the self-destruct sequence on the boat and killed everyone onboard. Their bodies were never found.