

## 1. Explainable AI (XAI)

Think back to the days of high school math class. It wasn't enough to write down the final answer. You had to show your work, explain your steps, and make your thought process clear. That's the essence of Explainable AI (XAI).

XAI is all about creating artificial intelligence models that provide clear explanations for their decisions. It's not enough for the model to make a prediction or decision; it should also be able to show us how it got there.

### Example

Imagine an AI model that approves or denies loan applications. It's not sufficient for the model just to say "approved" or "denied." We need to know why:

- Was it the applicant's credit history?
- Their income?

XAI requires the model to explain its reasoning. This is crucial because it brings transparency, trust, and fairness to AI. It helps us understand and verify the decision-making process of AI models, which is particularly important in high-stakes industries like healthcare or finance.

It's like showing your work in math class — it validates your solution and helps others understand your process.

## 2. Generative AI

Picture this: a virtuoso painter, a spellbinding novelist, and a visionary architect all fused into one entity. That, in a vibrant splash of creativity, is the essence of Generative AI.

Generative AI is like the magic paintbrush of the artificial intelligence world. It doesn't just analyze or predict; it creates. From images to music, prose, and even synthetic voices, Generative AI crafts content from scratch, producing outputs that can mirror, and at times even surpass the creative prowess of humans.

### Examples

Let's dip our toes into a few real-world examples:

- **Deepfake videos:** You may have come across those eerie [deepfake videos where a celebrity appears](#) to be saying or doing something they never actually did. That's the handiwork of Generative AI! It uses an array of images and videos to fashion a realistic but completely artificial rendering of the celebrity.
- And then there's **ChatGPT**, developed by OpenAI. You type a message to it, and it responds, generating conversation as a human would.

- In the realm of **art and design**, there are AI models that can sketch entirely new fashion designs or generate unique pieces of art, just from being fed a mix of styles and patterns.
- **The healthcare industry:** Generative AI can generate synthetic data for research, useful for scenarios where actual data is scarce or privacy concerns are high.
- **The entertainment industry**, too, isn't untouched. Generative AI is used in video games and virtual reality to create rich and immersive environments that adapt to players' actions in real-time.

### 3. Large Language Model (LLM)

When you're in a new city, you're probably grateful for your GPS, right?

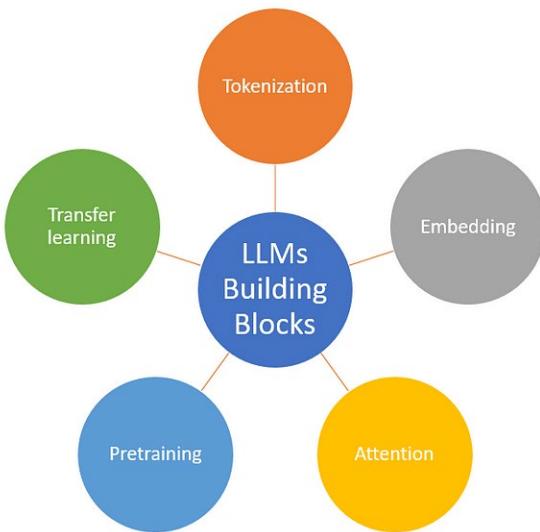
It helps you find your way, pointing out turns and landmarks, ensuring you reach your destination without getting lost. Well, in the vast landscape of AI, the Large Language Models (LLMs) are like the powerful GPS guiding us through the complexities of human language.

LLMs are the linguists of the AI world, designed to understand, generate, and engage with human language in a way that feels natural and intuitive. And when we say 'large,' we mean it — these models are trained on enormous amounts of text data, enabling them to grasp nuances, context, and even the idiosyncrasies of the language.

#### Examples

**ChatGPT** is a prime example of LLM that can generate human-like text, answer queries, and even write poetry. But the reach of LLMs extends far beyond just generating text. They're being used in:

- **Translation services** to improve the fluidity and accuracy of translated texts.
- **Customer service** to powering chatbots that can handle complex inquiries.
- They're even **helping students learn**, by providing tutoring in various subjects.



## LLM Building Blocks

### 4. Large Language Model Meta AI (LLaMA)

Have you ever tried to assemble a piece of furniture using a manual, but the instructions only came in a language you didn't understand? Frustrating, right?

Now, imagine having a tool that could translate the manual into your language and even explain it in the context of your specific furniture. This is quite similar to what LLaMA — Language, Logic, and Mathematics — does in the world of AI.

LLaMA is a versatile craftsman in the artificial intelligence toolbox. It's designed to enable AI to understand and work with not just human languages, but also formal languages such as logic and mathematics. It extends the scope of AI from understanding everyday text to deciphering and creating complex mathematical proofs or logical reasoning.

### Examples

- LLaMA could be used to develop an AI system that **helps students** solve complex mathematical problems. It could understand the problem stated in natural language, translate it into a formal mathematical problem, solve it, and then explain the solution back in human language.
- In the realm of **software development**, a LLaMA-based AI could assist programmers in understanding and writing complex code, spotting logical errors, or even generating pieces of code based on specific requirements stated in ordinary language.

### 5. Neural Networks

Imagine you're in a vast forest, and each tree represents a tiny piece of information. Alone, these trees might not mean much, but together, they form an intricate, interconnected ecosystem.

This is the essence of Neural Networks in the realm of artificial intelligence.

Each 'tree' or node in a Neural Network takes in data, processes it, and passes it on, much like trees sharing nutrients in a forest. And just like a forest that adapts to its environment over time, these nodes learn from the data they process, constantly improving their ability to recognize patterns, make predictions, and so much more.

Neural Networks are a vibrant echo of our brain's own web of neurons, striving to mimic our biological learning processes.

## 6. Natural Language Processing (NLP)

Picture this: You're at a bustling international airport, surrounded by people from all corners of the world speaking in languages you might not understand.

But what if you had a friendly translator at your side, seamlessly interpreting those languages for you, helping you communicate and even understand the nuances and emotions in each conversation? That's pretty much what [Natural Language Processing \(NLP\)](#) does in the world of artificial intelligence.

But instead of an airport filled with people, you have a digital world brimming with data. This data can be in the form of text, voice recordings, social media posts, customer reviews, and so much more. NLP is like your personal AI translator, decoding this data, and making sense of human language with all its complexities, idiosyncrasies, and even sarcasm.

This extraordinary technology allows machines to read, understand, and derive meaning from human language in a valuable and structured way. It's what powers your digital assistant when it understands your voice commands, or when your email filters out spam.

## 7. Reinforcement Learning (RL)

Do you remember the thrill of mastering a new skill as a child, perhaps riding a bike or baking cookies?

There was a pattern:

Try, stumble, learn from the mistake, and then try again with your newfound knowledge.

That's exactly the principle behind Reinforcement Learning (RL) in artificial intelligence.

In RL, an AI agent is dropped into an environment where it's free to roam, stumble, and learn. The AI is like a kid on a playground, learning the ropes of a new game. It takes an

action, sees the result, and then receives a reward (a virtual pat on the back) or a penalty (a gentle nudge in the right direction) based on the outcome.

Over time, just like a kid learning to master the monkey bars, the AI agent starts figuring out which actions lead to rewards and which ones result in penalties. It's continuously learning from its interactions with its environment, with the goal of maximizing its total reward over time.

This learning model is at the heart of training self-driving cars, game-playing AI like [AlphaGo](#), and much more!

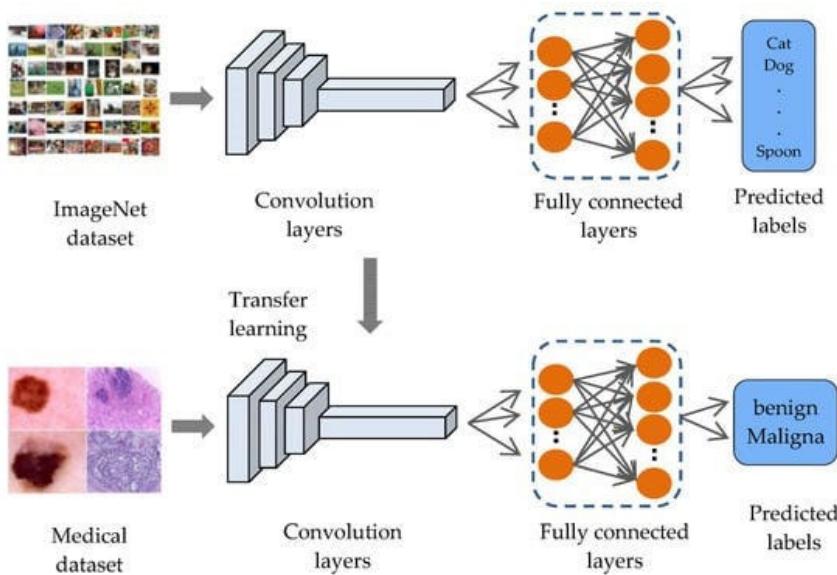
## 8. Transfer Learning

Remember how you once learned to ride a bike and then found it surprisingly easy to learn to ride a scooter?

That's because you were able to transfer the skills you learned from one task to another. This concept forms the crux of Transfer Learning in the world of artificial intelligence.

In Transfer Learning, an AI model applies knowledge gained from one problem to solve a different but related problem. It's as if the AI is a quick-learning student who can cleverly apply the geometry learned in math class to solve problems in physics class.

For example, a model trained to recognize cars in images might not need to start from scratch to recognize trucks. It can transfer what it has learned about shapes, edges, and wheels from the car dataset and use this knowledge to recognize trucks, making the learning process faster and more efficient.



[Incorporating a Novel Dual Transfer Learning Approach for Medical Images](#)

## 9. Federated Learning

Imagine a potluck dinner, where everyone brings their own dish to contribute to the meal. Each person prepares their dish at home, but the end result is a combined feast that everyone can enjoy.

Federated Learning works in a similar way, but instead of food, we're dealing with [data](#) and machine learning models:

1. In this setting, your device, let's say your smartphone, holds valuable data that can help train a machine learning model.
2. Instead of sending your data to a central server — which raises *privacy concerns* — the learning comes to you. **The model is sent to your device**, and learns from your data.
3. Only the learnings (or the updated model) are sent back to the server. Your actual data never leaves your device.

### Example

Consider this example: you're using a predictive text app. Your smartphone contains unique ways you use language, which could be invaluable in training the app.

- With Federated Learning, the app can **learn from your typing style** without your personal data ever leaving your smartphone.
- The updates it makes from learning your style are then shared back with the central model, improving its understanding and making predictions even better.

So, the next time you hear “Federated Learning,” remember the potluck dinner — everyone brings something to the table, contributing to the collective intelligence **without giving away their secret recipe!**

In the same way, Federated Learning brings the power of collective learning to AI while preserving privacy.

## 10. Computer Vision

Imagine you're at a bustling city park. You effortlessly take in the sights — children playing, kites soaring in the sky, dogs chasing frisbees. Your brain processes these images with ease, understanding the scenes unfolding around you.

What if we could teach computers to ‘see’ and understand images and scenes just like we do?

That's the magic of Computer Vision in the realm of artificial intelligence.

Computer Vision is like giving a pair of intelligent eyes to a computer. It's about teaching machines to “see” and interpret visual data from the world around them.

Whether it's recognizing a face in a crowd, detecting obstacles on a self-driving car's path, or analyzing medical images for signs of disease, Computer Vision enables computers to understand and make sense of visual information.

It's this capability that's driving innovations in everything from autonomous vehicles to augmented reality!

## 11. Autoencoders

Ever played a game of charades, where you had to distill the essence of a movie into a few key gestures?

Now, what if we could train an AI to play a similar game with data?

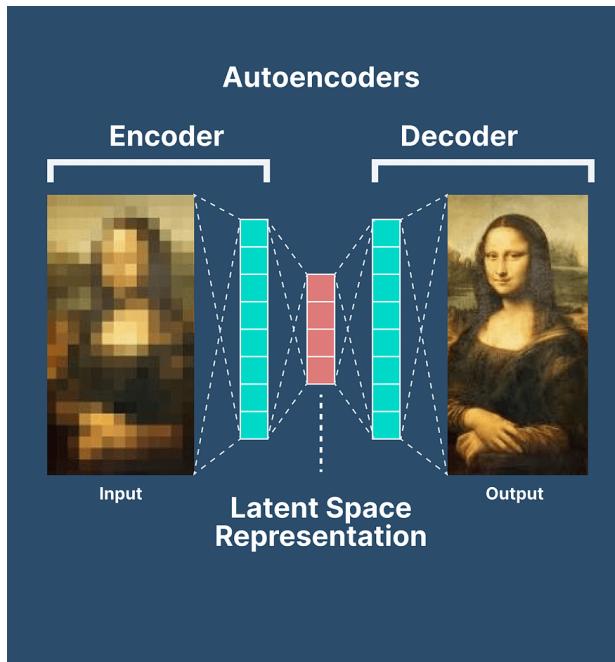
That's where Autoencoders, a special kind of neural network, come into the picture.

An Autoencoder is like an expert minimalist artist, tasked with capturing the essence of a complex scene in a few simple strokes. It works in two steps: encoding and decoding:

1. In the encoding phase, the model compresses the input data, capturing its most important features into a condensed representation, much like the artist distilling a scene into a few lines.
2. The magic happens in the decoding phase. Here, the model takes the compact representation and reconstructs the original data as closely as possible. It's like the artist filling in the details to recreate the scene from the minimalistic sketch.

The result?

A powerful tool for data compression and noise reduction. Autoencoders shine in tasks such as anomaly detection, where they can learn to reconstruct "normal" data, and anything that deviates from this reconstruction is considered an anomaly.



## [Autoencoders in Deep Learning](#)

### Examples of use cases

Autoencoders are widely used in the field of artificial intelligence and machine learning for their ability to reduce dimensionality and remove noise from data. Here are some real-world examples of their usage:

1. **Anomaly Detection:** Autoencoders can be used to detect unusual patterns or outliers in data. For instance, in credit card fraud detection, an autoencoder can be trained on normal transactions. When it encounters a fraudulent transaction, the reconstruction error would be high since the transaction is different from what it has learned, signaling a potential anomaly.
2. **Image Denoising:** Autoencoders can be used to remove noise from images. They can be trained to reconstruct the original image from the noisy input. This has applications in fields like medical imaging, where clear images are critical.
3. **Data Compression:** Autoencoders are also used for reducing the dimensionality of data, a process known as data compression. For example, they can be used to reduce the dimensionality of image data for faster processing in computer vision tasks.
4. **Recommendation Systems:** Some recommendation systems, such as those used by online streaming or shopping platforms, utilize autoencoders to capture the underlying factors that drive user preferences, thereby improving the quality of recommendations.

5. **Generative Modeling:** Autoencoders, particularly Variational Autoencoders (VAEs), can generate new data that resemble the input data. This is used in fields such as deepfake creation or synthesizing new pharmaceutical molecules in drug discovery.

## 12. Transformer Models

Have you ever attended a lively party where you're juggling multiple conversations at once?

You need to understand who's saying what, keep track of different topics, and decide which conversation requires your attention the most.

Now, imagine an AI model that can handle data in a similar way. Welcome to the world of Transformer Models.

A Transformer Model is like the life of the AI party, juggling multiple pieces of data at once. It pays attention to different parts of the data depending on their importance, a process known as the “attention mechanism.”

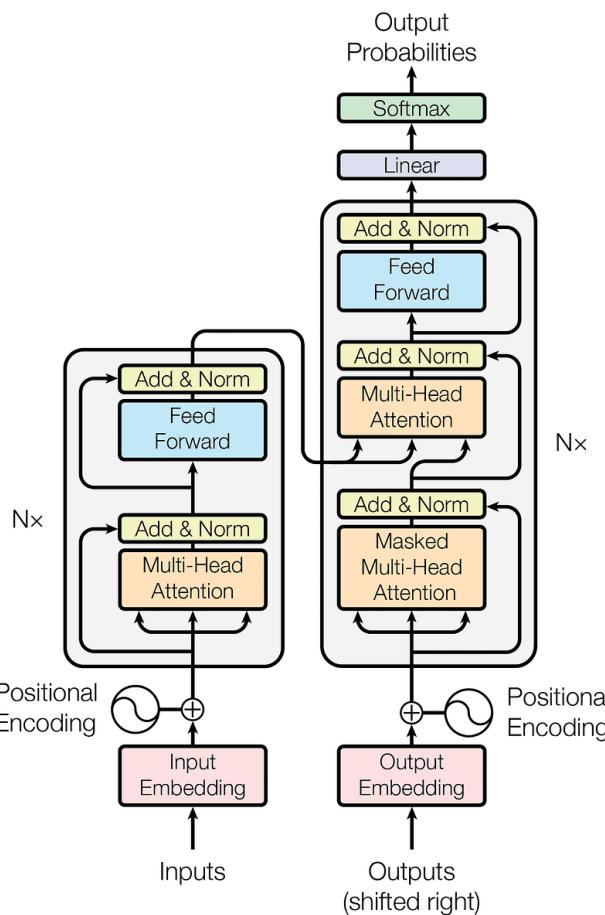
This ability is what makes Transformer Models incredibly good at understanding sequences and relationships in data, especially when it comes to natural language processing tasks.

Take, for instance, a sentence where the meaning of a word depends on another word that appears much later in the sentence. A Transformer Model can directly focus on the dependency between these two words, no matter how far apart they are, allowing it to understand the context better.

Transformer Models form the backbone of many state-of-the-art language models like **GPT-3** and **BERT** ([Bidirectional Encoder Representations from Transformers](#)), which are used for tasks such as translation, summarization, and sentiment analysis.

So, the next time you hear ‘Transformer Models,’ think of this sociable AI model, adept at handling the intricacies of data just like a party-goer deftly navigating multiple conversations.

It's this ability to “pay attention” that's propelling AI's understanding of language to new heights!



## The Transformer Model

### 13. Attention Mechanisms

Picture yourself in a crowded, noisy room trying to focus on one particular conversation. Amidst all the noise, your brain has this amazing ability to ‘tune out’ the irrelevant chatter and “tune in” to the conversation that matters most.

A similar concept exists in the world of AI, and it’s known as Attention Mechanisms.

Think of Attention Mechanisms like an expert conductor leading a grand orchestra. In the midst of all the instruments playing together, the conductor knows when to highlight the violins, when to focus on the drums, or when the flutes need to take the lead.

Similarly, Attention Mechanisms in AI models decide which parts of the data to “highlight” and focus on at a given time. These could be certain words in a sentence or certain parts of an image. By deciding what’s important and what’s not, Attention Mechanisms help the model make better predictions.

They are especially useful in tasks like language translation, where the meaning of a sentence often depends on understanding the context and relationships between different words.

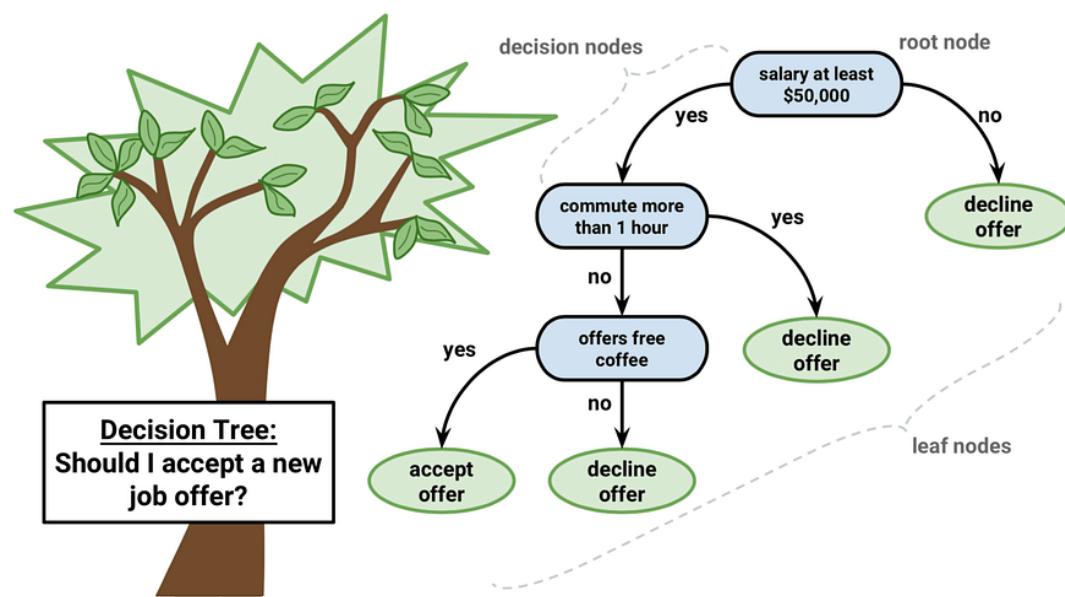
With Attention Mechanisms, models can “focus” on the important words that influence the meaning of others.

## 14. Decision Tree

You’re planning an outdoor birthday party. To decide whether to hold it in your backyard or move it indoors, you might consider a few factors:

- Is it likely to rain?
- How hot will it be?
- Is there too much wind?

This step-by-step decision-making process resembles how a Decision Tree works in machine learning.



### Decision Tree

A Decision Tree is like your personal party planner in the world of AI. It’s a flowchart-like structure where each internal “node” represents a question or “decision”, each branch represents the outcome of that decision, and each leaf node represents a final “verdict” or prediction.

What makes Decision Trees popular is their intuitive and visual nature. They can handle both categorical and numerical data, and they mimic human decision-making. But, like planning a party, there can be hiccups. Sometimes they overcomplicate the decisions (**overfitting**) or oversimplify them (**underfitting**).

It’s this branching path of decisions and outcomes that helps AI navigate through data and make accurate predictions!

## 15. Bias and Variance

Imagine you're playing a game of darts. Your goal is to hit the bullseye, but let's say your darts are consistently landing to the left. That's similar to what we call "bias" in machine learning — when your model consistently misses the mark.

Now, let's imagine your darts are landing all over the dartboard. That's similar to "variance," when your model's predictions are all over the place.

In the world of machine learning, we're often playing a game of balance between bias and variance, just like trying to perfect your aim in darts:

- High bias can cause an algorithm to miss relevant relations between features and target outputs (**underfitting**), akin to consistently missing the bullseye.
- On the other hand, high variance can cause an algorithm to model the random noise in the training data (**overfitting**), similar to your darts landing inconsistently.

The sweet spot is finding the right balance, a bit like training yourself to throw the darts so they land as close to the bullseye as consistently as possible.

It's this fine-balancing act that's key to improving the accuracy of our AI predictions!

## 16. Generative Adversarial Networks (GANs)

Ever seen a pair of painters engage in a friendly duel, each one trying to outdo the other, and in the process, creating fantastic pieces of art? That's a fitting metaphor for Generative Adversarial Networks, or GANs, in the dynamic arena of artificial intelligence.

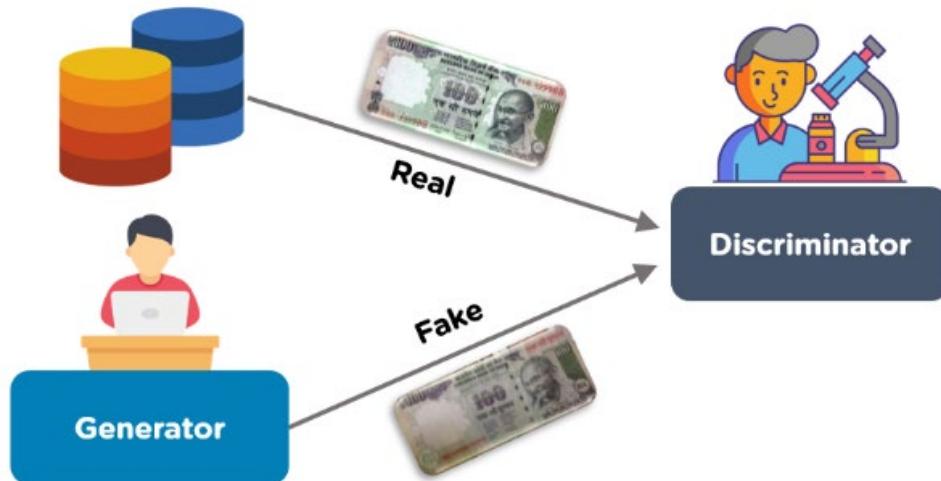
Here's how it works: GANs involve two AI models:

- The first model, the "**generative**" part, is like the imaginative artist, creating fake images from random noise. It's trying its best to create artwork (or any other form of data) so convincing that it looks like the real deal.
- Meanwhile, the second model, the "**adversarial**" part, plays the critic. It looks at these creations alongside real images and tries to tell which is which.

The catch?

The generator is always learning, always improving its craft based on the feedback from the discriminator.

The result is an AI system that can create astonishingly realistic data — from AI-generated art that could fool even the most discerning critic, to deepfake videos that are strikingly lifelike.



### Generative Adversarial Networks (GANs)

## 17. Convolutional Neural Networks (CNNs)

Do you remember playing the game “I Spy” as a child, trying to pick out a single object from a busy scene?

Now, what if we could train a computer to play “I Spy” at superhuman levels?

That’s where Convolutional Neural Networks, or CNNs, come into play in the world of artificial intelligence.

CNNs are a specialized kind of Neural Network that’s really good at handling image data. Picture CNN as an eagle-eyed detective, meticulously scanning every inch of a picture, one small patch at a time. This detective doesn’t just see the whole image, but can recognize intricate patterns, textures, and shapes that make up the bigger picture.

For example, it can break down a picture of a cat into features such as pointy ears, round eyes, and whiskers, then piece these features together to recognize the image as a cat.

It’s this ability to discern intricate patterns that make CNNs incredibly useful in areas like facial recognition, medical imaging, and even in self-driving cars.

## 18. Recurrent Neural Networks (RNNs)

Have you ever watched a gripping TV series where the plot builds up from episode to episode, and each storyline is intertwined with the previous ones?

Now, what if a computer could understand such sequences, recognizing patterns over time just like we do?

Meet the Recurrent Neural Networks, or RNNs, the time lords of the artificial intelligence world.

RNNs are like AI detectives with a knack for remembering past events. Unlike traditional neural networks, they save the “memories” or outputs from their previous steps and use this information to influence their future decisions.

Picture an RNN as a bookworm, engrossed in a gripping novel, remembering past plotlines and characters to understand the unfolding story.

This ability to recall past information makes RNNs ideal for handling sequential data, like sentences in a text or time-series data. Whether it’s predicting the next word in a sentence, understanding the sentiment in a movie review, or analyzing stock market trends, RNNs are at the heart of it all.

Now, at this stage, you may ask:

Is GitHub Copilot leveraging RNNs in order to give developers suggestions for what they want or have to write next in their source code?

That’s a great question! Behind GitHub Copilot is powered by a language model called **Codex**, which is a sibling to **GPT-3**, developed by OpenAI.

While GPT-3 and Codex use Transformer architectures, the details of their training processes are proprietary information and not publicly available.

**Transformers**, the architecture used by these models, have proven to be more effective for many tasks compared to Recurrent Neural Networks (RNNs).

Transformers overcome some limitations of RNNs, like handling long-range dependencies in sequences, by using a mechanism called **attention** that weighs the importance of different elements in a sequence.

### Examples of RNNs application

There are many popular AI tools and applications that have been developed based on Recurrent Neural Networks (RNNs) because of their unique ability to handle sequential data. Here are some examples:

1. **Google Translate:** For several years, Google Translate leveraged an architecture called Long Short-Term Memory (LSTM), a type of RNN, for its translation services. While they’ve moved to Transformer models more recently, RNNs played a crucial role in the initial development of this tool.
2. **Amazon Alexa and Google Assistant:** These voice-activated virtual assistants often employ RNNs for their natural language processing capabilities, including understanding user commands or generating human-like responses. They’re particularly useful in processing spoken language, which is inherently sequential.

3. **Siri:** Apple's Siri is another virtual assistant that uses RNNs for voice recognition and natural language processing tasks.
4. **Facebook:** Facebook has used RNNs for a variety of tasks, including text translation, sentiment analysis, and even for generating automatic captions for videos.
5. **Grammarly:** This writing assistance tool uses deep learning technologies, including RNNs, to detect grammar and spelling errors, suggest better phrasing, and even assess the tone of the text.

Remember, while RNNs have been fundamental in the growth of these tools, they often work in concert with other AI technologies and models. And as AI continues to evolve, many applications are moving towards using more advanced models like Transformers, which can often provide better performance on many tasks.

## 19. Long Short-Term Memory (LSTM)

Do you remember studying for exams, where you had to remember key facts for the test but then forgot them soon after?

Interestingly, our brains have a mechanism to remember critical information for short periods, while also retaining crucial long-term memories. In the world of artificial intelligence, there's a model that mimics this process called Long Short-Term Memory (LSTM).

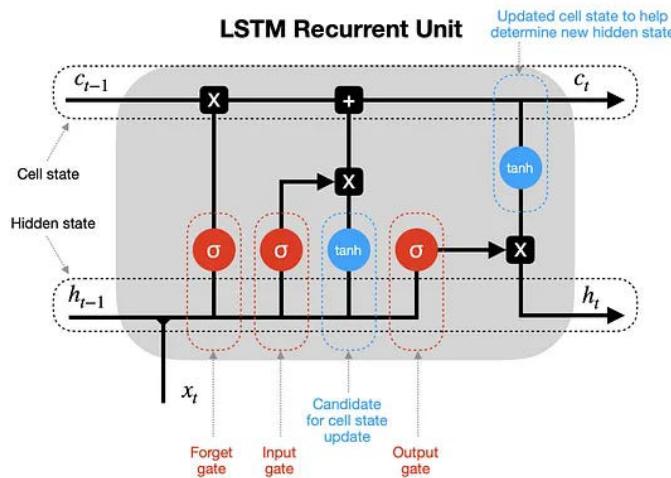
Picture LSTM as a selective librarian in the vast library of AI memory. This librarian not only knows which books (or pieces of information) to fetch for immediate use but also which ones are worth keeping for the long haul.

In an LSTM network, data flows through components called cells, each equipped with an input, output, and forget gate, much like doors that control the flow of information:

- The **input gate** decides what new information to let in,
- the **forget gate** chooses what to forget or discard,
- and **the output gate** determines what information should be output to the next layer of the network.

This ability to remember crucial long-term dependencies while forgetting irrelevant details makes LSTMs a powerful tool for handling sequential data. They've been used for everything from speech recognition and language translation to predicting stock market trends.

# LONG SHORT-TERM MEMORY NEURAL NETWORKS



[Long Short-Term Memory \(LSTM\) Neural Networks](#) by

[Saul Dobillas](#)

## 20. Feature Engineering

Let's imagine you're a detective, trying to solve a mystery. Your success depends largely on picking up the right clues, interpreting them correctly, and connecting them to form the bigger picture.

In the world of machine learning, this crucial process of selecting and crafting the right “clues” from raw data is known as Feature Engineering.

Feature Engineering is like the discerning detective of the AI world, identifying important pieces of information (or ‘features’) in the data that can help solve the mystery — in this case, make accurate predictions. It involves transforming raw data into a format that is more suitable for models to work with.

For example, let's say you have data on houses and you want to predict their prices. The number of rooms, location, size of the house, and age could all be useful features that influence the price. But sometimes, the clues aren't so obvious. Maybe the combination of features, like the ratio of bathrooms to bedrooms or the proximity to schools, might be more important.

Feature Engineering is all about finding or creating these key pieces of information.

It's a creative and crucial process that can significantly improve a model's performance.

## 21. Hyperparameter Tuning

We all know the key ingredients to bake a cake: flour, sugar, butter, and eggs.

But baking a perfect cake is not just about mixing these ingredients together; it's about getting the quantities just right. A little too much butter or a pinch less sugar could alter the taste dramatically.

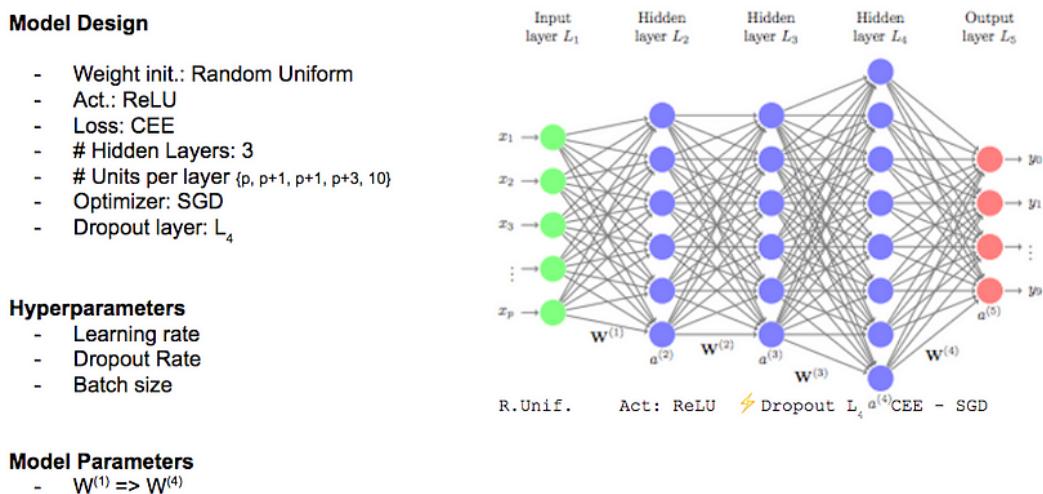
Similarly, in the world of machine learning, there's a process called Hyperparameter Tuning, which echoes the same meticulous tweaking and adjusting. It's the art of perfecting a recipe for machine learning. However, instead of adjusting the quantity of flour or sugar, we adjust parameters like:

- learning rate,
- the number of layers in a neural network,
- or the depth of a decision tree.

Just as too much or too little sugar can spoil a cake, incorrect hyperparameters can lead to underperforming machine learning models.

Interestingly, these hyperparameters aren't learned from the data during the training process; they need to be set beforehand. Finding the best values often involves a blend of experience, trial-and-error, or a systematic search, much like baking the perfect cake.

## E.g. DNN



## Hyperparameter Tuning: Variables classification example

## 22. Model Evaluation Metrics

Think of a sports coach analyzing a player's performance. They don't just consider one factor, like the number of goals scored. They take into account assists, defensive contributions, accuracy, and more.

These varied factors provide a comprehensive evaluation of the player's performance.

Similarly, in machine learning, we use Model Evaluation Metrics to assess the effectiveness of our models.

Model Evaluation Metrics are like the coach's performance stats for machine learning models. They provide a way to measure the quality and predictive power of models based on different aspects that vary depending on the type of problem we're trying to solve. For instance:

- In a **classification problem**, we might look at **accuracy, precision, recall, or the F1 score**.
- In a **regression problem**, we might consider **mean absolute error or mean squared error**.

Just as a coach wouldn't rely on a single stat to evaluate a player's overall contribution, we shouldn't rely on a single metric to evaluate a model. It's a multi-faceted evaluation that helps ensure our AI models are ready for the big leagues!

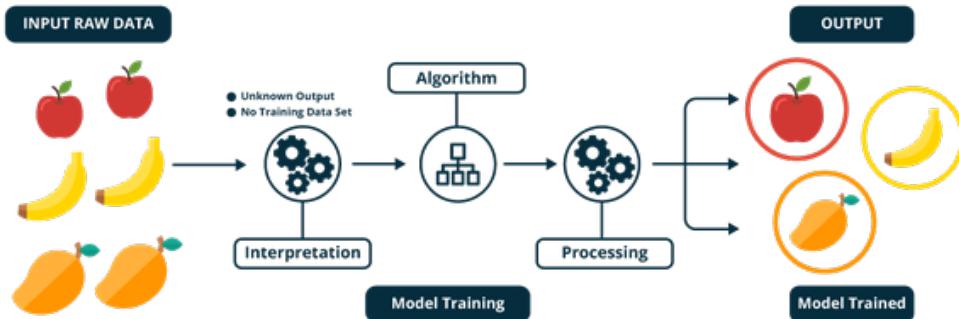
## 23. Trained Data

Think back to the time when you were learning to drive. You didn't just jump in the car and hit the freeway, did you? Instead, you began with understanding the controls, practicing in a safe, open space, and gradually taking on more complex driving situations with an instructor guiding you. In the realm of AI, 'Trained Data' is akin to this practical driving manual.

Trained Data is the AI equivalent of a detailed guidebook. It's a compilation of data that's been carefully pre-processed and labeled, serving as a blueprint for AI systems to learn from. This trained data helps the AI to understand the patterns it should look for and the decisions it should make.

### Examples

- Consider a **voice recognition** AI like the one powering your smartphone's voice assistant. The Trained Data for this system would include countless sound clips of spoken words and phrases, tagged with the correct text that they represent. By learning from these examples, the AI becomes proficient at transcribing your spoken commands into text.
- Similarly, in **weather forecasting**, Trained Data composed of historical weather patterns and outcomes helps predictive AI models generate accurate weather forecasts for future dates.



Unsupervised learning, explained by some tasty fruit

## References

- Stanford University's course CS231n: [Lecture on neural networks](#)
- Google: [Federated Learning: Collaborative Machine Learning without Centralized Training Data](#)
- Trained Data: [Google AI blogs](#)
- OpenAI: [Generative Models](#)
- [Explainable Artificial Intelligence \(XAI\)](#) by DARPA.
- [Applied Deep Learning — Part 3: Autoencoders](#)
- [Reinforcement Learning: An Introduction](#) by Richard S. Sutton and Andrew G. Barto.
- [Attention is All You Need](#)
- [The Illustrated Transformer](#) by Jay Alammar
- [Decision Trees](#)
- [Gentle Introduction to the Bias-Variance Trade-Off in Machine Learning](#)
- [Generative Adversarial Networks](#)
- [Understanding LSTM Networks](#)
- [Fundamental Techniques of Feature Engineering for Machine Learning](#)
- [How to Tune Algorithm Parameters with Scikit-Learn](#)
- [Metrics and scoring: quantifying the quality of predictions](#)