General

Can machines really think, learn, and act intelligently?

In this post, we're going to define what machine learning is and how computers think and learn. We're also going to look at some history relevant to the development of the intelligent machine.

jeff
Sep 18, 2023 • 11 min read



Can machines really think, learn, and act intelligently?

😎 There are so many introductory posts about AI and ML, and yet I decided to write this one. Do you know why? Because they're all boring. Not this one though, this one is cool. It also includes an awesome *optional* quiz to test your AI superpowers. But here's the deal, you promise to subscribe if you do the quiz. Okay?

Introduction

Before we start talking about all the complicated stuff about Artificial Intelligence and Machine Learning, let's take a big-picture look at how it began and where we are now.

It can be hard to focus and understand what everything means with all the buzzwords floating around, but fear not! If you're new to this, we're going to break down the most important things you should know if you're just getting started.

(Man, I love it when "getting started" fits into paragraphs nicely.)

History

So, humans have brains, we know that. In more recent news however, research in neurology has shown that the human brain is composed of a network of things called "neurons". Neurons are connected to each other and communicate using electrical pulses.

So what do neurons in the human brain have to do with Artificial Intelligence? Oh, they're just the inspiration behind the design of neural networks, a digital version of our biological network of neurons. They're not the same, but they kind of work similarly.

> 💡 I am going to expand on neural networks later in this post, so keep on reading.

Buffalo, New York 1950s

Meet the Perceptron. The first neural network, a computing system that would have a shot at mimicking the neurons in our brains. Frank Rosenblatt, the American psychologist and computer scientist who implemented the perceptron, demonstrated the machine by showing it an image or pattern, and then it would make a decision about what the image represents.

For instance, it could distinguish between shapes such as triangles, squares, and circles. The demo was a success because the machine was able to learn and classify visual patterns and it was a major step toward the development of machine learning systems.

However, the perceptron wasn't perfect, and interest in it steadily declined due to the fact that it was a **single-layer neural network**, meaning it could only solve problems where a straight line could be drawn to separate the data into two distinct cases.

> 💡 The perceptron is like a ruler, it could only draw straight lines. If things are all mixed up and can't be split neatly with a line, it'll get confused and fail.

Toronto, Ontario, 1980s

Meet Geoffrey Hinton, a computer scientist and a big name in the field of artificial intelligence and deep learning. Why? He had a huge impact on the development of deep neural networks (aka multi-layered neural network).

Unlike a single-layer neural network, Hinton developed a multi-layered network and contributed significantly to the backpropagation algorithm which allows the network to “learn”.

To train a multi-layer neural network (or deep neural network), you'll need to strengthen the relationship between the digital neurons. Meaning, that during the training process, and using iterations, a specific input will fine-tune the connections between individual neurons so that the output becomes predictable.

Makes sense? Great. If not, have no fear as we're going to dive deeper into this topic below, so keep on reading! 📚

The "Thinking Machine"

Why did Hinton and Rosenblatt spend time and effort to develop artificial intelligence in the first place? Well, for the same reasons as to why we invented cars, planes, or houses. To improve life! And by having a machine do the boring work for us, oh boy how much time we'll save to focus on the real stuff that matters, like browsing through Reddit, and X (Formerly Twitter). *(just kidding)*

Okay, so we, as humans are natural innovators, and scientists believed that they could one day teach machines to think for themselves so we can delegate all the boring repetitive tasks.

Still, while there was interest in such wild ideas, others wondered what would happen to humanity if machines could think. Will machines take over the world? Will computers become our masters? Will everyone become unemployed?

Well, to understand the possibilities of what could happen, we must understand if machines are actually able to think, and if they are, what are their capabilities and limitations.

What does "think" mean?

For humans, thinking is generally a process that involves (at least) the important steps below:

1. **Perception**: Recognize things from the world around us.
2. **Memory**: Remember things, by accessing previous knowledge through storing them somewhere in our heads.
3. **Attention**: Understanding intentions from given information.
4. **Logic**: Analyzing information based on perception, understanding, memory and, determining intent.
5. **Creativity**: Creating new things, like art, music, stories, etc..

Today, many argue that computers can think just like humans do. To illustrate this, I've listed the cognitive processes in the table below and provided an example of how modern computers perform the same functions of thinking as humans do:

😎 **Personal opinion**: I do somewhat agree with the idea that computers nowadays can think, similarly to how we can. Actually, the recent progress is mind-blowing.

To attain intelligence, just thinking, however, is not enough. We'll need to teach computers so they can figure things out on their own. What this means is that we'll want computers to correctly interpret things that they've never been exposed to before.

This all brings me to my final artificial intelligence formula:

$$\textbf{thinking} + \textbf{learning} = \textbf{artificial intelligence}$$

💡 By implementing a neural network we'll have thinking, and by training the machine we'll fine-tune the network and achieve learning.

Okay, now we're going to take a look at the two parts that make up artificial intelligence. First, we'll just define what neural networks are, and then briefly discuss machine learning.

Introduction to Neural Networks

A neural network is a subset of machine learning which itself is a subset of artificial intelligence. The structure and function of neural networks are inspired by the human brain (as we've seen earlier), in a way where connections between neurons (a network comprised of multiple neurons) are strengthened and reinforced through training to improve its overall prediction accuracy.

Every single neuron (or node) in a neural network has the following properties:

1. **Inputs:** A neuron can receive one or more inputs from other neurons in the network. Inputs can be continuous or binary values.
2. **Weights:** Each input to the neuron has a weight that determines the strength of the connection between this neuron and the one sending the input. Weight for each connection between neurons is then strengthened during training.
3. **Activation Function:** A neuron will perform the necessary computation to determine the output based on the inputs and weights. Examples of this function include the sigmoid function that outputs values between 0 and 1.
4. **Bias:** A bias is a value that is used to shift the activation function in a particular direction. Think of the bias as a knob that tweaks sensitivity to certain information, the higher the sensitivity the more likely this neuron will fire.
5. **Output:** Each neuron then has an output, which is essentially the result of the activation function.

Here is a visual representation of a single neuron:

Single Neuron in a Network

The output is the result of the activation function which will include, the inputs, weights, and bias represented as such:

$$Y = f(w1 * X1 + w2 * X2 + b)$$

So, we've seen what every single neuron is, and as we know the collection of neurons makes up the neural network. Awesome, now let's quickly go over what Machine Learning is.

Introduction to Machine Learning

Ok, so now we know what a neural network is composed of and we also know that a neural network gives the machine the power of thinking. But how does it learn?

We'll need to train it. Just like teaching a child the difference between an orange and an apple through trial and error, we train the machine by providing information and correcting its guesses until it makes the right prediction.

This process is iterative, and won't stop until the machine is able to predict the desired output correctly.

I will not go into the details here to keep this simple, but this is done mathematically using algorithms that determine the best values for the weights and biases which themselves represent the strength of a connection between two neurons.

> 💡 One popular example is **backpropagation**, which is used in conjunction with an optimization algorithm such as stochastic gradient descent.

Basically machines learn through training iteratively, until the output matches what we want. It’s boiled down to finding the right values for the weights and biases.

Let’s take this for example:

In a puzzle, you have the pieces but you don’t know how to put them together so it looks like the picture on the box. You try to put the pieces randomly and compare them with the picture, but hey, you notice that it doesn't look right! Now you keep repeating the process until you have the right picture. The relationship between pieces is re-adjusted and re-evaluated with every try. This is somewhat an oversimplified example of backpropagation but kind of gives you an idea of what's happening behind the scenes.

Now it’s important to point out, that a Neural Network is a type of machine learning process, also known as deep learning that is used to train and teach machines. There are other techniques to teach machines, which we're not going to touch on in this article. Maybe in future ones? Let me know!

> 👉 If you're interested in more details, check out this post some other time.

It's about a type of learning called **Supervised Machine Learning**.

Okay, let's now define what Artificial Intelligence is since we've covered Neural Networks and Machine Learning.

What is Artificial Intelligence?

Put simply, it is a broad term that covers everything that is a non-human form of intelligence. I've prepared this little awesome quiz to help you determine if you can properly classify buzzwords and common terms.

I call it "Can You Identif-AI":

QUIZ: Can You Identif-AI?

Is "Neural Network" Artificial Intelligence or Not?

- AI
- Not AI

(This quiz runs in your browser only and does not store any info whatsoever)

Putting it all Together

We have AI on top and then all other stuff branch out. Neural Networks for example would go under Machine Learning. So the next time you're having a discussion with someone who sounds confused about the topic or mixing things together, you know where to tell them to go.

Right? 👀

Here's a visual representation of how things are connected:

Most Popular Branches of Artificial Intelligence

Final Thoughts

So can machines learn? Yes, they can and they're pretty good at it and the accuracy rate is only getting better. There's even a lot of talk nowadays about technological unemployment, a term that refers to mass unemployment in the future due to machines taking over our jobs.

And that's exactly why you're doing great by sharpening your skills and learning about Artificial Intelligence and its different branches so that when the time comes you'll stay relevant in the future economy.

And hey! Thanks for reading through. This stuff takes time to write, and I don't use ChatGPT or other Generative AI to write my posts, so show some love by subscribing for free before you go, I'll make sure to send you only awesome content.

I'm also active on X (Twitter) if you'd like to connect.

Cheers! ✌️

Member discussion

0 comments

Start the conversation

Add comment

LangChain vs Semantic Kernel: Which framework is better for your RAG app?

Confused between LangChain and Semantic Kernel? This short post will help you choose the most suita…

Feb 5, 2024 · 2 min read

A must-read if you're using or planning to use AI to generate blog posts

Using AI to generate blog posts? It's important to understand how to do it effectively to maximize…

Feb 2, 2024 · 6 min read

How to use the Gemini Pro API, LlamaIndex, and Pinecone to build your next RAG app

Let's talk about building a simple RAG app using LlamaIndex, Pinecone, and Google's Gemini Pro…

Jan 30, 2024 · 9 min read · 4 comments