Confluence INVESTMENT MANAGEMENT

Bi-Weekly Geopolitical Report
By Thomas Wash

June 26, 2023

Distinguishing My Wife From a Hat, an AI Story

In his book, *The Man Who Mistook His Wife for a Hat and Other Clinical Tales,* Oliver Sacks, a neurologist, details his experience with patients suffering from varying neurological disorders. In one such case, he dealt with a man who had lost the ability to recognize faces. The patient was a university music teacher who had always been known for his calm and collected demeanor but had suddenly began behaving strangely. He would sometimes fail to recognize his students and was often seen patting the top of water fountains and parking meters as if they were small children. His antics were widely regarded to be jokes since he didn't have trouble communicating, and his musical ability was as good as it had ever been.

It wasn't until the patient was diagnosed with diabetes that he sought professional help. Aware of the disease's impact on his eyesight, he visited an ophthalmologist, who reassured him that his vision was fine but referred him to see Dr. Oliver Sacks for a neurological exam.

During the visit, Dr. Sacks noticed something was off about the patient. The man seemed to be having trouble perceiving Dr. Sacks fully. Instead of looking directly at Dr. Sacks' face, the patient fixated on certain parts. He gazed at Dr. Sacks' nose, chin, right ear, and right eye, but never his face as a whole. After telling the patient the exam was over, the man attempted to find his hat but instead reached for his wife's head and tried to lift it as if he were about to put it on. To Dr. Sacks' surprise, the patient's wife treated this as if it were an everyday event.

The case of the man who mistook his wife for a hat is a great illustration of how artificial intelligence (AI) neural networks process information. Like the patient, neural networks do not have the ability to look at an entire image to judge what it is. Instead, they break down images into parts, specifically into an array of numbers called pixels. AI models (neural networks) see images by recognizing patterns in the numbers that represent the image. They can make distinctions between different objects through training which then teaches them to associate certain patterns with specific objects. Just like the patient who needed to see an eye, nose, and mouth to know that he was looking at a face, AI models need numbers to achieve the same task.

This report provides a beginner-friendly introduction on how AI learns and processes information. We will begin by discussing the similarities between AI and our brains. Next, we will explain how AI works and explore some of its most important applications. We will then discuss some of the challenges and limitations of AI. We end the report with market ramifications. While this is not intended to be an exhaustive summary of AI, readers should come away with a stronger understanding of the technology and why it is such a big deal.

The Nuts and Bolts

To fully understand artificial intelligence, one needs to understand the human brain. The brain is made up of over 80 billion neurons which are nerve cells that transmit messages throughout the body. These neurons are constantly learning and adapting, and they help us perform simple tasks such as walking, eating, and talking to more complex routines such as playing sports, dancing, and adopting new skills. As we learn new things, the connections between the neurons in our brain strengthen, and our understanding of the world deepens. AI works similarly, as it learns to perform tasks by being given examples and feedback.

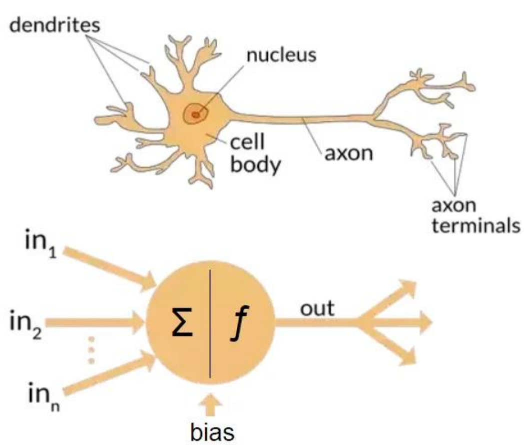

(Source: TowardsDataScience.com)

The makeup of an artificial neural network (ANN) was constructed to imitate its biological counterpart. A neuron is composed of a dendrite designed to receive information, a cell body to process it, and an axon to transfer it to other neurons. Similarly, an artificial neuron, also called a perceptron or a node, has an input channel that gathers data, a processing stage that transforms the data, and an output channel that sends information to other nodes.

An ANN has multiple layers, each with its own set of nodes. Information is transferred between layers through the connections between nodes. A simple ANN has three layers: an input layer, a hidden layer, and an output layer. The inputs, sometimes referred to as the visible layer, contain all the information the neural network has at its disposal to make a prediction. The hidden layer is responsible for transforming the data to fit the desired result. Lastly, the output layer displays the model's conclusion.

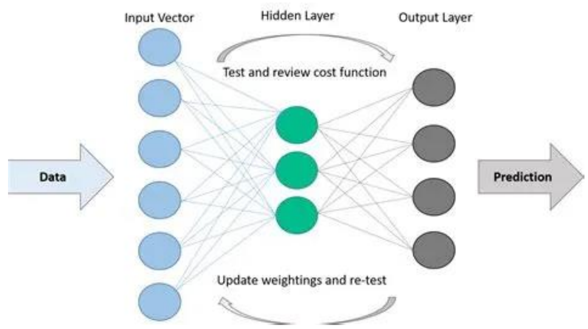

(Source: Medium)

The difference between a simple and complex ANN is the number of hidden layers. Simple ANNs have one hidden layer and are proficient at solving linear problems like time series and cross-sectional analysis. However, they cannot solve complex nonlinear problems, such as classification, because the lack of layers limits the model's ability to fit the data. To solve non-linear problems, additional layers are added which improves the model’s predictive power on tested data but hinders its ability to make accurate predictions on new data.

Our Brain vs. AI

Our brains are constantly bombarded with sensory information, which can sometimes lead to perception errors. One example of this is the McGurk effect. In this illusion, a subject is shown a video of someone saying "ba" while hearing audio of the person saying "fa." The result is that the subject hears the sound "da." AI models, on the other hand, do not have eyes to see the video nor ears to hear the audio. As a result, they are not easily fooled by illusions.

Instead of mammalian senses, ANNs perceive the world through numbers. This allows them to process information more efficiently and disregard irrelevant information more effectively. To see an image, an ANN will transform it into a grid of numbers called pixels. For plain text, it changes letters and characters into a numerical sequence called ASCII. Even sound is converted into binary numbers, each of which represents a sound wave. Through these numbers, an ANN can distinguish a person from a hat, an "A" from a "B," and a cough from a purr. It can also use numbers for more complex tasks such as recognizing speech patterns, editing pictures, and making predictions.

AI models are able to see the world by analyzing large volumes of data and looking for patterns within that numerical data. There are two approaches by which ANNs learn data: ***supervised learning*** and ***unsupervised learning***.

Supervised learning attempts to find patterns through labeled input and output data. The purpose of this style of learning is for either classification or regression analysis.

- **A Classification Model** is when data is separated into groups based on predefined parameters. A real-world example of a classification problem is when banks decide to deny or approve a loan. The bank will take information such as a person's income, overall debt burden, job status, etc., and then decide whether applicants are creditworthy.
- **A Regression Model** is designed to make inferences about a future event. This is seen when an analyst tries to forecast future monthly sales revenue using the previous month's sales receipts, market size, and cost structures.

Unsupervised learning tries to find patterns by looking for similarities without the help of labels. This type of learning is often used for tasks such as association, clustering, and dimensionality reduction.

- **Association Learning** is designed to find relationships between different features in a dataset. This type of algorithm is used when making recommendations for movies and songs based on data from consumers with similar tastes.
- **Clustering** is used to group data points together based on their similarities. This is typical in marketing when trying to find target groups for advertisements.
- **Dimensionality Reduction** is used to reduce the size of a dataset while preserving the most critical information. This is utilized when aiming to improve sound quality.

I Think, Therefore I Am

Despite the sophistication of AI models, they are limited in their ability to interpret information. ANNs are adept at analyzing data and drawing inferences based on inductive reasoning. This means they can take specific information and make general conclusions. For example, an ANN could look at a particular arrangement of numbers and determine that it is likely a picture of a person's wife and not a hat.

(Source: Nvidia)

Unlike ANNs, humans have the capacity for deductive reasoning. This means we can use general information to draw specific conclusions. For example, we can look at a person's face and determine that they are our spouse. We can then describe those features to someone else who will also be able to see a picture of a person with those features and come to the same conclusion. A person's ability to draw conclusions through deductive reasoning allows them to be able to think deeper than AI models and create a better understanding of the world.

The lack of deductive reasoning means that ANNs struggle to work backward to find an answer. Hence, it is difficult to start with an assumption that is correct and prove it is so. To overcome this flaw, AI models generate fictional answers that sound right but are ultimately wrong. This was demonstrated in a recent case where a lawyer cited six nonexistent cases developed by ChatGPT to help with his defense. The lawyer had mistakenly believed that the chatbot had access to real-world legal cases. As a result, he filed a court brief full of phony legal precedents, and he now faces possible punishments.

Market Ramifications

Artificial intelligence is a computational program designed to mimic the functions of the mammalian brain. It accomplishes this through a mathematical algorithm, often referred to as an Artificial Neural Network, that functions similarly to a biological neural network. This framework allows the AI model to learn through supervised and unsupervised processes. Even though these models mirror how the brain functions, they are primarily useful for solving problems that can be pieced together to form more extensive theories, but not the other way around. Therefore, these AI models are likely still far away from replacing human workers.

However, the AI craze is unlikely to end anytime soon as businesses will look for ways to incorporate this technology into their practices. The technology's ability to analyze large volumes of data allows businesses to become more efficient and improve their product offerings. While AI is already being used in phones and cars, we believe it will become increasingly integrated into consumers' everyday lives. Possible applications range from web searches to driving cars to creating internet pages. The recent rally in the tech sector reflects investors' expectations for the widespread adoption of AI.

While the potential of AI is vast, investors should not be overly optimistic about the pace of adoption and the technology's overall capabilities. AI is still in its infancy, and it will likely be some time before firms understand the best ways to integrate it. As the previous example of the lawyer shows, the technology's flaws are still being worked out. Additionally, countries are still determining how to create regulatory guardrails that will allow the technology to thrive while maintaining consumers' sovereignty over their data. As a result, investing in AI-related stocks does have a lot of upside potential, but also comes with many risks.

Thomas Wash
June 26, 2023

Confluence Investment Management LLC

Confluence Investment Management LLC is an independent Registered Investment Advisor located in St. Louis, Missouri. The firm provides professional portfolio management and advisory services to institutional and individual clients. Confluence's investment philosophy is based upon independent, fundamental research that integrates the firm's evaluation of market cycles, macroeconomics and geopolitical analysis with a value-driven, company-specific approach. The firm's portfolio management philosophy begins by assessing risk and follows through by positioning client portfolios to achieve stated income and growth objectives. The Confluence team is comprised of experienced investment professionals who are dedicated to an exceptional level of client service and communication.