Welcome to the World of AI
(whatever that is)

Ray Willson
OLLI Investment Forum
July 5, 2023

Gartner predicts the market for AI software will reach almost USD 134.8 billion by 2025.[1]

I Facts and Figures

According to Statista, revenue from the Artificial Intelligence (AI) software market worldwide is expected to reach 126 billion dollars by 2025.

As per Gartner, 37% of organizations have implemented AI in some form. The percentage of enterprises employing AI grew 270% over the past four years.

According to Servion Global Solutions, by 2025, 95% of customer interactions will be powered by AI.

A recent 2020 report from Statista reveals that the global AI software market is expected to grow approximately 54% year-on-year and is expected to reach a forecast size of USD $22.6 billion.

What Companies are Doing with AI

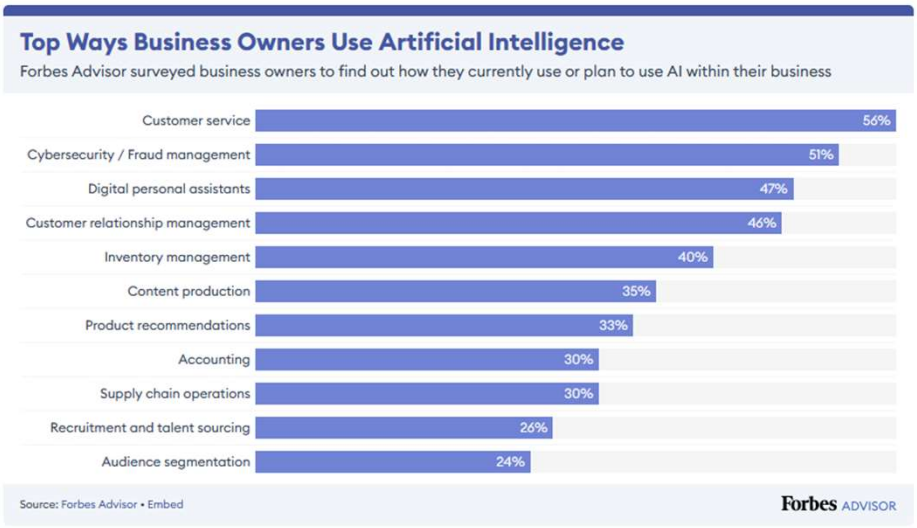

What People are Doing with AI

Plant Gardens
Plan Workouts
Plan Meals
Make a Gift
Design Parts for a Spaceship
Organize Computer Desktop
Write a Wedding Speech
Write an email
Get a first read
Play Devil's Advocate
Organize research for a thesis
Skim dozens of academic articles
Cope with ADHD
Sort through an archive of pictures
Transcribe a doctor's notes into clinical notes
Appeal an insurance denial
Build entirely new games
Create new proteins
Identify diseases in banana plants
Describe entire D&D worlds
Make a Spotify playlist
Play with language

NY Times, April 2023

What Is Artificial Intelligence (AI)?

Artificial intelligence (AI) refers to the simulation of human intelligence by software-coded heuristics. Nowadays this code is prevalent in everything from cloud-based, enterprise applications to consumer apps and even embedded firmware.

Source: Investopedia, By Jake Frankenfield Updated April 24, 2023 Reviewed by Gordon Scott

https://www.investopedia.com/terms/a/artificial-intelligence-ai.asp

Artificial Intelligence

While a number of definitions of artificial intelligence (AI) have surfaced over the last few decades, John McCarthy offers the following definition in this 2004 paper (PDF, 127 KB) (link resides outside IBM), " It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.“

Source: IBM at https://www.ibm.com/topics/artificial-intelligence

Artificial Intelligence (AI)

Artificial intelligence (AI) applies advanced analysis and logic-based techniques, including machine learning, to interpret events, support and automate decisions, and take actions.

Source: Gartner, IT Glossary

NATIONAL SCIENCE AND TECHNOLOGY COUNCIL

Artificial intelligence (AI)1 is one of the most powerful technologies of our time.

---

1 There are multiple definitions of AI and AI systems used by the federal government, including from in the National Defense Authorization Act for Fiscal Year 2019 [Public Law 115-232, sec. 238(g)], the National Defense Authorization Act for Fiscal Year 2020 [Public Law 116-617, sec. 5002(3)], and the NIST AI Risk Management Framework, as well as a more encompassing view of automated systems articulated in the Blueprint for An AI Bill of Rights. The R&D priorities defined in this document are applicable and important to the full breadth of technologies covered by these definitions.

Source: NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN 2023 UPDATE

What Does Artificial Intelligence (AI) Mean?

Artificial intelligence (AI), also known as machine intelligence, is a branch of computer science that focuses on building and managing technology that can learn to autonomously make decisions and carry out actions on behalf of a human being.

AI is not a single technology. Instead, it is an umbrella term that includes any type of software or hardware component that supports machine learning, computer vision, natural language understanding, natural language generation, natural language processing and robotics.

Source: Techopedia, Margaret Rouse editor

IBM's View of AI
(Matryoshka Concept)

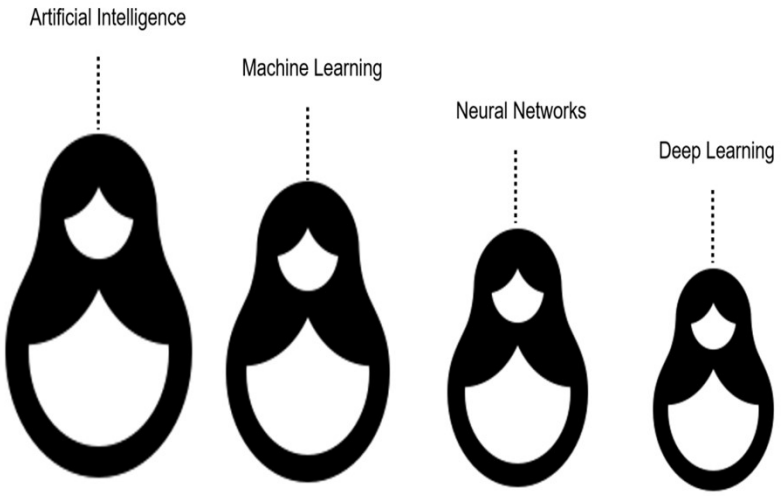

Machine Learning

- Machine learning is a branch of artificial intelligence (AI) and computer science which focuses on the use of data and algorithms to imitate the way that humans learn, gradually improving its accuracy. *IBM.com*
- Machine Learning is said as a subset of **artificial intelligence** that is mainly concerned with the development of algorithms which allow a computer to learn from the data and past experiences on their own. The term machine learning was first introduced by **Arthur Samuel** in **1959**. *Javapoint.com*
- Machine learning (ML) is a discipline of artificial intelligence (AI) that provides machines with the ability to automatically learn from data and past experiences while identifying patterns to make predictions with minimal human intervention. *Spiceworks.com*

Neural Network

- An artificial neuron network (neural network) is a computational model that mimics the way nerve cells work in the human brain. Artificial neural networks use learning algorithms that can independently make adjustments – or learn, in a sense – as they receive new input. *Techopedia.com*
- A neural network is a series of algorithms that endeavors to recognize underlying relationships in a set of data through a process that mimics the way the human brain operates. In this sense, neural networks refer to systems of neurons, either organic or artificial in nature. *Investopedia.com*
- An artificial intelligence system used for learning and classifying data and applied in research on pattern recognition, speech recognition, machine translation of languages, and financial prediction, among other areas. Neural networks are usually abstract structures modeled on a computer and consist of a number of interconnected processing elements (**nodes**), each with a finite number of inputs and outputs. *APA Dictionary of Psychology*

The first neural network was conceived of by Warren McCulloch and Walter Pitts in 1943. They wrote a seminal paper on how neurons may work and modeled their ideas by creating a simple neural network using electrical circuits.
At a basic level, a neural network is comprised of four main components: inputs, weights, a bias or threshold, and an output.
The application of neural networks to artificial intelligence (AI).
AI research quickly accelerated, with Kunihiko Fukushima developing the first true, multilayered neural network in 1975.

The original goal of the neural network approach was to create a computational system that could solve problems like a human brain. However, over time, researchers shifted their focus to using neural networks to match specific tasks, leading to deviations from a strictly biological approach. Since then, neural networks have supported diverse tasks, including computer vision, speech recognition, machine translation, social network filtering, playing board and video games, and medical diagnosis.

The main difference between regression and a neural network is the impact of change on a single weight. In regression, you can change a weight without affecting the other inputs in a function. However, this isn't the case with neural networks. Since the output of one layer is passed into the next layer of the network, a single change can have a cascading effect on the other neurons in the network.

Deep Learning

- Deep learning is a subset of machine learning, which is essentially a neural network with three or more layers. These neural networks attempt to simulate the behavior of the human brain—albeit far from matching its ability—allowing it to "learn" from large amounts of data. While a neural network with a single layer can still make approximate predictions, additional hidden layers can help to optimize and refine for accuracy. *IBM.com*

Deep learning requires a tremendous amount of computing power. High performance graphical processing units (GPUs) are ideal because they can handle a large volume of calculations in multiple cores with copious memory available. However, managing multiple GPUs on-premises can create a large demand on internal resources and be incredibly costly to scale.

Large Language Model

LLM is a supervised learning algorithm that uses ensemble learning method for regression. Ensemble learning method is a technique that combines predictions from multiple machine learning algorithms to make a more accurate prediction than a single model. *expert.ai*

A **large language model (LLM)** is a language model consisting of a neural network with many parameters (typically billions of weights or more), trained on large quantities of unlabeled text using self-supervised learning or semi-supervised learning. LLMs emerged around 2018 and perform well at a wide variety of tasks. *Wikipedia*

Investing in AI

I don't know what the hell this "logistics" is that Marshall is always talking about, but I want some of it. - Admiral E. J. King

Runyon's Law: The race is not always to the swift, nor the battle to the strong, but that's the way to bet. *Damon Runyon*

Gartner predicts the market for AI software will reach almost USD 134.8 billion by 2025.[1]

I Facts and Figures

According to Statista, revenue from the Artificial Intelligence (AI) software market worldwide is expected to reach 126 billion dollars by 2025.

As per Gartner, 37% of organizations have implemented AI in some form. The percentage of enterprises employing AI grew 270% over the past four years.

According to Servion Global Solutions, by 2025, 95% of customer interactions will be powered by AI.

A recent 2020 report from Statista reveals that the global AI software market is expected to grow approximately 54% year-on-year and is expected to reach a forecast size of USD $22.6 billion.

Investment Options

- Companies supplying an AI Product or components needed for AI
  - AI Product: Google, Microsoft, Amazon, C3.ai
  - AI Elements: Steppers, GPU chips, LLMs, Servers, Databases, Data Warehouses, Back-up generators, Cyber security
- Funds investing in companies providing the above capabilities or services

AI-Related Companies

Symbol Name	Forbes	US News	Motley Fool	Nerdwallet
MSFT Microsoft	X	X	X	
AMZN Amazon	X		X	
NVDA Nvidia Corp	X	X	X	X
AI C3.ai Inc	X		X	X
GOOGL Alphabet Inc	X	X		
MU Micron Technology Inc	X			
TSLA Tesla Inc	X			
IBM IBM Corp			X	
TSM Taiwan Semiconductor Mfr		X		
META Meta Platforms Inc		X		
ASML ASML Holding Co		X		
SAP SAP SE		X		
RELX RELX Plc		X		
BIDU Baidu Inc		X		

Sample of AI ETFs

Symbol Name	$AUM	Holdings	YTD%	Expense
AIQ Global X AI & Technology	391M	91	39.43	0.68
BOTZ Global X Robotics & AI	2.5B	48	39.81	0.69
IRBO iShares Robotics and AI Multisector	435M	123	29.27	0.47
ROBT First Trust Nasdaq AI and Robotics	386M	100	27.2	0.65
THNQ Robo Global AI	48.6M	66	32.78	0.68
As of Jul 3, 2023				

All of these are index-based ETFs. ROBO, IRBO and ROBT are equal-weighted, while BOTZ is market-cap-weighted.

The portfolio composition of these ETFs can vary quite significantly, and by extension, so can their returns.

Over the past year, returns have ranged from a loss of 30% to a loss of 8.5%

Sample of AI ETFs

Symbol Name	$AUM	Holdings	YTD%	Expense
AIQ Global X AI & Technology	219M	96	34.26	0.68
BOTZ Global X Robotics & AI	2.12B	49	37.27	0.69
IRBO iShares Robotics and AI Multisector	338M	153	25.67	0.47
ROBT First Trust Nasdaq AI and Robotics	294M	113	25.98	0.65
THNQ Robo Global AI	31.3M	73	32.78	0.68
As of Jun 9, 2023				

All of these are index-based ETFs. ROBO, IRBO and ROBT are equal-weighted, while BOTZ is market-cap-weighted.

The portfolio composition of these ETFs can vary quite significantly, and by extension, so can their returns.

Over the past year, returns have ranged from a loss of 30% to a loss of 8.5%

Sources

Books

Robot-Proof, Higher Education in the Age of Artificial Intelligence by Joseph E. Aoun

Machine Learning by Ethem Alpaydin

Deep Learning by John D. Kelleher

The Myth of Artificial Intelligence by Erik J. Larson

I, Human by Tomas Chamorro-Premuzic

Four Battlegrounds, Power in the Age of Artificial Intelligence by Paul Scharre

Sources

Newspapers, Periodicals

The Washington Post
Wall Street Journal
US News
Barron's
Financial Times
Economist
Forbes

Websites

IBM
Bankrate
Investopedia
Wikipedia
Motley Fool
The Street
Money
Nerdwallet

Sources

Papers

Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond (Cite: arXiv:2304.13712v2)

Authors: Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Bing Yin, Xia Hu

Top 10 Silicon Wafer Mfrs

- Shin-Etsu Chemical
- Sumco
- Globalwafers
- Siltronic AG
- SK Siltron
- Soitec
- Wafer Works
- Okmetic
- Episil
- Xinsheng Semiconductor

Top 10 Foundries

- TSMC
- Samsung
- UMC
- Global Foundries
- SMIC
- HuaHong Group
- PSMC
- VIS
- Tower
- Nexchip

MIT Press Series

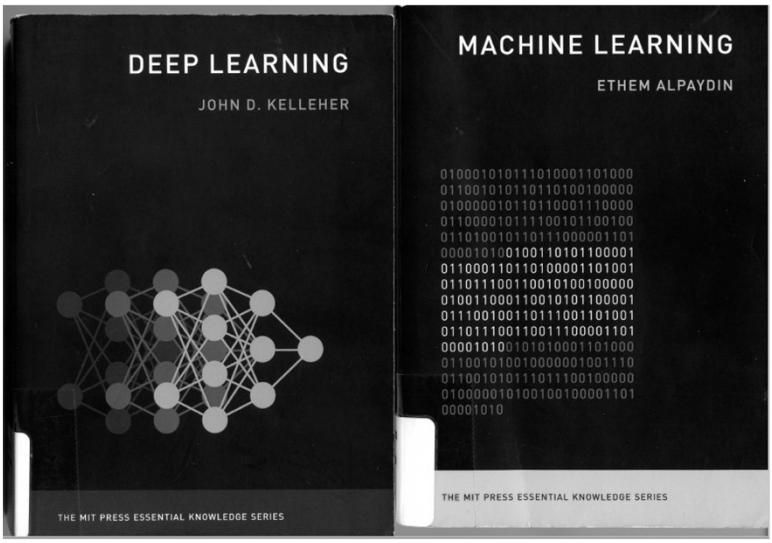

MIT has produced a series of over 40 pocket-sized books on a variety of topics of what they call "current interest" Topics range from scientific and technical to cultural and historical