

What Is AI?

Thanks to science fiction, many people probably feel like they have a good grasp on artificial intelligence. The image that comes to mind is some kind of computer, maybe in the form of a robot, that can communicate and make decisions just like a human. Thanks to modern advances, certain technologies on the market today bear some resemblance to this imagined ideal, making it seem as though AI will soon be able to perfectly replicate human learning and thought processes.

The reality is a little different. For starters, artificial intelligence can be thought of as a very broad spectrum of capability. If the basic definition is “computing that mimics human intelligence,” then most of computing history is actually about the evolution of artificial intelligence, and many of today’s standard IT practices were yesterday’s cutting-edge AI. Secondly, there is a large difference between human-like behavior in a specific use case and the ability to have that behavior across a wide variety of situations. By fundamentally understanding what AI is and how today’s AI is different from previous models, people can be more informed about how this trend will fit into business and society.

The Different Types of Artificial Intelligence

At the highest level, artificial intelligence is software programming. There are some ways in which AI is no different from any other software application and some ways in which it is a completely new discipline. Framing AI as software programming helps demonstrate how AI has been an evolving concept since the invention of the first computer.

If AI is computing that mimics human intelligence, then there are many aspects of human intelligence that can be considered. One aspect is the ability to quickly calculate mathematical operations. The first commercial computers replicated (and exceeded) this particular ability. Another aspect of human intelligence is the ability to reach conclusions based on specific inputs. For a long time, software has been evolving to navigate decision trees of increasing complexity. Today, we do not think of these functions as artificial intelligence, but they serve as examples of how computing has steadily assumed tasks once performed only by people.

Modern AI represents a step into a previously unexplored aspect of human intelligence: the ability to use large amounts of disparate data to arrive at a conclusion that has a high chance of being correct. In other words, the ability to make a guess.

This shift brings computers much closer to the way humans learn and make decisions, but it also introduces a new way of thinking about computing results. In the past, software programs have always been deterministic—certain inputs always led to certain outputs, and those outputs could be trusted (assuming the inputs were correct). With AI, the results are probabilistic. There are many different inputs, and the computer makes decisions based on the data it collects and the algorithms that provide guidance. The result may be a never-before-seen insight, or it could be way off the mark based on some faulty logic.

Think of it as the difference between tic-tac-toe and chess. Tic-tac-toe has a finite number of moves, and each move can be coded based on the current setup of the board. Chess essentially has an infinite number of moves and board configurations. Coding every single move would be impossible. Instead, computers are told the basic rules of chess and allowed to make their own decisions. Sometimes these

might be bad decisions, but often they are decisions that have never been played before which might open up a new avenue for a win.

The main takeaway for businesses using AI, whether they are developing their own AI tools or procuring them from a vendor, is that there must be new processes for integrating the results of AI. The results must be reviewed to ensure they make sense, and there must be some ability to troubleshoot or reverse-engineer the mechanics of the AI system. For the foreseeable future, this reality shows why human-machine hybrids will be the dominant model. AI can produce new and innovative results, but it requires checks and balances to avoid bias and reverse bad decisions.

Pieces of Modern AI

If AI is a blanket term that shifts over time as computing becomes more powerful, then this most recent incarnation has specific characteristics that are worth understanding. Three concepts in particular play a large role in defining the unique properties of modern AI.

- **Machine Learning.** To a large extent, machine learning is actually the end goal for most of the cutting-edge work in AI. Machine learning generally refers to techniques that allow computers to gather data and build their own pathways for analyzing that data. There is no single approach or algorithm that defines machine learning; instead, there are many different methods being used to produce machine learning capability. This means that an end user needs to know the details behind the specific machine learning process they are implementing.
- **Deep Learning.** Deep learning is one of the most common approaches for machine learning. With deep learning, computers analyze problems at multiple layers in an attempt to simulate how the human brain analyzes problems. Visual images, natural language, or other inputs can be parsed into various components in order to extract meaning and build context, improving the probability of the computer arriving at the correct conclusion.
- **Neural Networks.** These software models consist of five parts—a decision-making unit called a neuron, inputs, weights that can be assigned to the inputs, biases that shape the analysis, and a single output (also called an activation function). Multiple stages of neurons can be stacked together to result in the multi-layered analysis found in deep learning.

While the underlying neural networks have some similarities to models used in traditional software programming (such as inputs and outputs), the way that the model is built and the use of weights and biases to make predictions and learn on the fly makes AI programming a very different task.

There is one final piece of modern AI that is a huge requirement but may get overlooked in the tendency to focus on the software operations. Data sets are critical for AI in two main areas: first as the base for training the program and second as the ongoing fuel for operations. Many companies have struggled with data management as they try to absorb new streams of data with varying structure, and a lack of proficiency around building the right data sets can prevent AI initiatives from moving out of a prototype stage.

Most AI activity today is happening within companies with the resources to perform their own AI development, but over time AI will become more and more productized. For businesses to truly reap the benefits of AI beyond moderate improvements to existing processes, they will have to embrace the changes necessary for handling probabilistic results and build the capabilities needed for managing massive quantities of data.