

1

What AI Is and What AI Isn't

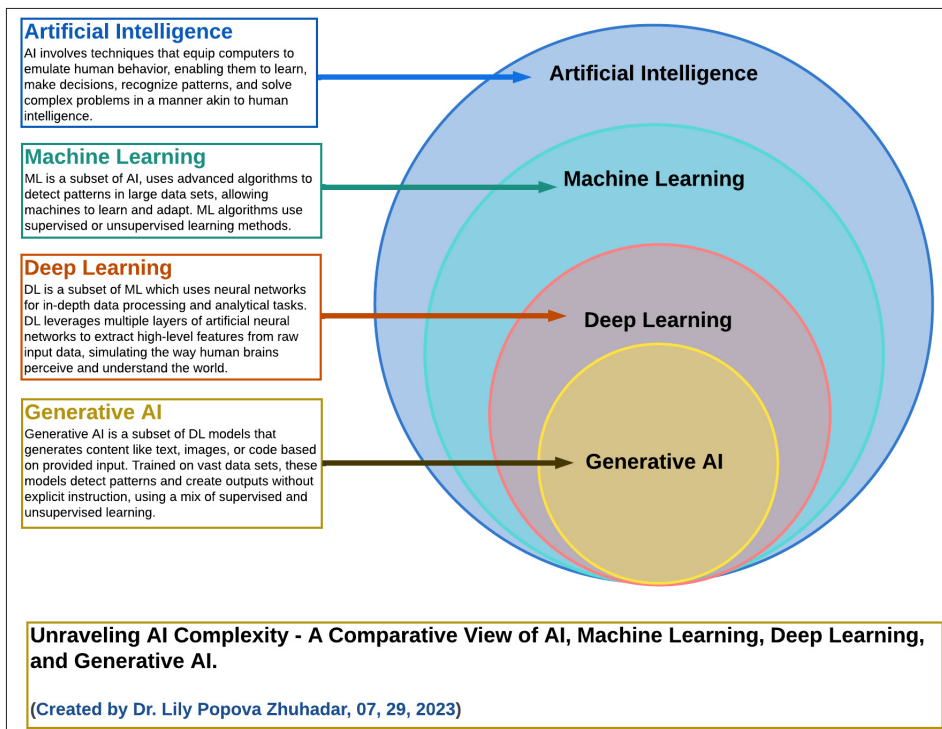


Image by PopovaZhuhadar, CC BY-SA 4.0, via Wikimedia.

A chart explaining machine learning, deep learning, artificial intelligence, and generative AI.

AI Basics

Artificial intelligence (AI) is an idea that means very different things to different people. However, the basic concept involves creating artificial systems that imitate human intelligence or some other kind of intelligence as seen in living beings of various kinds. As of 2024, there are no true artificial intelligences in existence, meaning that there are no machines that possess intelligence like a human, or a cat, or a mouse, or a fish. Machines, even the most advanced, possess no independent thought, personality, consciousness, or identity. Machines are not beings. They are not artists, or economists, or “thinkers,” they are merely tools, like an abacus, or a calculator, or a slide ruler. They do nothing unless told to do so, they have no thoughts unless given thoughts to think *about* by an actual intelligent being. Machines extend human capabilities and can be used to imitate human talents, but they do not think and they are not intelligent. This remains the realm of science fiction for now.¹

The concept of AI can be traced back to ancient cultures, where mythmakers theorized about the creation of magical automatons or golems, inanimate machines or inorganic beings endowed with the capability to think and make decisions. Science fiction authors and theorists took this idea further, creating fictional “robots” and intelligent machines of many different kinds. These inhuman beings, endowed with human-style thought and consciousness, served often as metaphors. The “Tin-Man,” in *The Wizard of Oz*, was a metaphor for the need to temper intellect with emotion. Fictional artificial intelligences are quite often used to reflect on the nature and value of human emotion, or to critique the way that intellectual calculations can lead to the devaluing of human life. Over the centuries, the idea of the thinking machine has been a useful thought experiment, but it has also been an arena of scientific exploration.

In the 1950s, as scientists and mathematicians began to realize the potential for computational technology, people started to get serious about the potential to actually create intelligence machines, as had long been fantasized about. The discovery of digital circuits essentially led to the belief that the human brain was essentially a living computer completing calculations. We now know that this metaphor doesn’t work exceptionally well for describing the human thought process and that digital computing is not very much like human brain activity. Nevertheless, the view that it should be possible to create a mechanized brain led to the birth of a new field of computer science. The actual term “A.I.,” meaning “artificial intelligence” is often attributed to Massachusetts Institute of Technology (MIT) researcher John McCarthy, which he defined as a computer that could engage in tasks or activities that were, at the time, seen as the exclusive domain of humanity. Another pioneer of the field was Alan Turing, who created the famous “Turing test” to measure whether a

computer was an intelligent system. Other pioneers like Warren McCulloch, Walter Pitts, John von Neumann, and philosopher Marvin Minsky, all helped to develop the early theory of decimal computing and the philosophy of potential artificial intelligence.²

From there, AI moved through a number of different stages. There were periods of intense activity, typically initiated by some new technological breakthrough in computing speed or capacity to manage larger stores of data. The discovery of networking, digital and wireless information exchange, language processing, machine vision, and machine “learning” all led to major leaps forward in the study of artificial intelligence. As the study of AI systems advanced, researchers began looking at AI research as proceeding down two separate paths of development.

First, there is what is called “strong” AI, which involves the creation of an artificial mind with the capability to autonomously conceptualize problems and to create original solutions. This is the ultimate goal for many AI enthusiasts, but currently humanity has not come close to this achievement. It is unlikely that current technology or development even has the potential to lead to the achievement of strong AI, as there is a massive conceptual leap between what is possible and what would need to be possible for a true AI to be created. Some theorists believe that potential future developments, like quantum computing, might allow scientists to break through the barrier to create strong AI, but there is no indication currently that this is correct. To create true independent intelligence, it is necessary not only to endow a machine with the capability to process information, but the machine must also be empowered with consciousness and volition. Without this, no AI can be truly autonomous and so no AI would be an independent mind, it would simply be an extension of human minds.

While research in strong AI has not produced any successes, scientists have made progress on what is called “weak” AI, which involves the creation of technology that imitates some aspect of human intelligence. Machines that can process language, for instance, or those that can “see” the same type of visual information as a human brain can see, are examples of weak AI. These machines are not intelligent, but they are based on studies of human capabilities and they imitate some aspect of biological existence and personhood. It is far easier to approach a weak AI problem than it is to approach a strong AI problem and so it is from the weak AI development field that all of the AI-adjacent technologies have come. These technologies are not intelligent, but they provide the illusion of intelligent action and advantages to users in a variety of ways.³

Imitating Thought vs. Having Thoughts

The 2020s has proven to be latest boom in weak AI development and has seen a surge of public interest in tools and programs marketed to the public as “AI,” though few have any real relationship to artificial intelligence. The “AI craze” of the 2020s has been based largely on advancements in three AI-adjacent fields of inquiry: language processing, generative image technology, and what is somewhat misleadingly called “machine learning.”

Language processing concerns technology that can recognize and respond to language, either when inputted in digital form, visual characters, or when spoken. This technology can also be linked with technology that can reproduce an imitation of a human voice and can therefore be used to simulate a conversation. Generative imaging involves machines that can generate seemingly “new” images or videos by combining elements of other images. The technology called “machine learning” involves the construction of machines that can change or improve their function as the machine collects data. None of this is actually “intelligent” technology, but all of these technologies imitate human activities that are involved in or representative of human intelligence.

Generative imaging technology, which has been the source of significant concern in the arena of intellectual property, works by combining elements of other images to create images that look like original images or that imitate a person or animal or thing that is real. The images and videos created using these technologies are not original, because the computer system is incapable of thought. Thus, a computer cannot “imagine” a thing and create it. It cannot create anything unless humans have already created things for it to sample. The machine can then mix and match different elements taken from human creators, to imitate the artistic process.

While generative technology has proven popular with those who want to create visual products but lack the skill, talent, or imagination to create these visual products in traditional ways, the introduction of this technology has been confusing for many people because companies and journalists have sometimes referred to “AI artists,” as if the machine itself was an artist. This is, of course, nonsensical, as the machine is not a person and cannot therefore be an anything-ist. Generative tools are also not creating something new, they are essentially “remixing” already existing images, color palettes, etc., much like a DJ might remix an existing song. It is therefore the human using the program who is the “artist,” to whatever extent this kind of visual remixing deserves to be called a legitimate art form.⁴

Language processing is another major realm of advancement in the field of weak AI, and this technology has led to several arenas of commercial development. First, there have been the new “chatbots,” programs that can respond to human verbal information and can then use that information in a variety of ways. These language processing systems can be used to make more advanced search engines that are better able to interpret vocal requests from users and to find relevant information online. Users are already familiar with this kind of technology, but advancements in language processing have the potential to make voice commands and voice searches much more powerful and accurate. Manufacturers of these technologies have begun marketing them under the term “AI,” but this is merely marketing language and does not signify the employment of any actual artificial intelligence.

Language processing can be used to “generate” language, just as generative technologies can be used to generate images. There are a number of companies that offer the ability to generate writing, based on prompts from a user. A user might ask the system to “write a term paper on the dynamics of heating systems used for desert water management programs.” The generative writing system will then

search for available information, as well as information about how “term papers” are structured, and will then produce an imitation of a term paper. This potential has led to some controversy over who owns intellectual property created by AI and how much human influence needs to be used before a piece of generated language can be considered the work of a human author. Just as with generated imagery, generative language technology cannot create anything new. All of the ideas must already be present in some form for the technology to incorporate those ideas into language. The resulting product is not an original product, but a collection of sampled human ideas from uncredited authors and thinkers.

Finally, language processing can be used to create “chatbots,” which are machines that are able to “converse” with users, for various purposes. This last usage is also a callback to the early days of AI research. Alan Turing—one of the pioneers of artificial intelligence research—created the now famous “Turing test,” which was the earliest proposed way to test whether a machine was intelligent. Turing reasoned that one measure of intelligence would be when a computer could essentially “fool” a person into thinking that the machine was a person. There were many early examples of chatbots created to engage in conversations, but this technology has grown far more advanced since the 2010s, and therefore has attracted more interest among the public.

Turing’s “test” turned out to be a poor measure of intelligence, because humans are *very* easy to fool. Even simple machines, created in the 1960s, were sufficient to fool many computer users into believing they were conversing with an actual human, long before it was possible to simulate anything approaching the subtleties of real human language. This discovery ended up leading to new breakthroughs in human psychology as philosophers and psychologists began doing research into why humans tend to ascribe intelligence and emotion to inanimate and unemotive features of their environment.⁵

Again, because machines can’t think, the language processing chatbots aren’t thinking of answers, what they are doing is imitating the human conversational process, drawing all of their ideas from ideas that other humans gave to the chatbot as seeds for possible responses to certain conversational situations. A chatbot can combine other ideas to produce what *seems* like a novel thought, but it is not really a thought, because the machine isn’t really thinking. A number of companies are incorporating this chatbot language processing technology into their “virtual assistants,” which will endow virtual programs like Amazon’s Alexa or Apple’s Siri with the capability to better recreate the illusion of human interaction, though the underlying function of the programs will remain the same.

It is important to understand that chatbots and language processing computers do not “know” languages and they do not “learn” languages or words, or anything else. When a human learns a word, they learn by associating symbols or sounds with conscious experiences. Because machines cannot have conscious experiences, they cannot learn. To a machine, a language is a set of symbols. Machines are programmed to provide certain sets of symbols in response to other sets of symbols. A user may ask the computer a question, and the computer may provide an answer,

but this is not because the computer understood the question, it just has a complex system for providing the appropriate symbols in response to other symbols that are given to the machine. Thus, machine language processing imitates the understanding of language without any actual understanding taking place.

Finally, one of the biggest areas of weak AI research has been in the field of machine learning. Computer engineers used the term “learning” to refer to machines that can incorporate new information and then use this information to alter the machine’s functions. So, a chatbot can add additional information given by a human and this will alter the way that the machine responds to human interaction.

Machine learning has the potential to transform many of the technologies that people use in their professional and personal lives. Virtual assistant-type programs, for instance, will be better able to learn about their users’ habits and attitudes and can refine the ways in which they respond. Other “learning” machines will be used in scientific and medical research, incorporating new data into their analyses as they process information. The basic science of machine learning is therefore about making computer systems that can automatically change or refine their functioning over time, with the addition of new data. In many applications, this machine “learning” may make it appear as if the computer or program is actually learning, but this is, again, not what is occurring.

Language Problems and Marketing Opportunities

When it comes to artificial intelligence, America and the rest of the Western world have a language problem. Companies have increasingly been calling their products “AI,” and claiming to enhance existing products with “AI” tools, but none of the tools are intelligent, they are simply based on imitating aspects of intelligence, and this creates confusion among consumers. Years of science fiction about AI and the fact that humans are very easily fooled by even basic machines that imitate human speech patterns, has primed humanity to believe that true intelligent machines, the “hard AI” dreamed of by scientists and philosophers for centuries, has arrived, but this is not actually the case. The tools that are available now, like “AI art generators,” or “chatbots,” or the newly debuted “AI assistants,” are not fundamentally different from other tools that already exist. They are still just programs doing what programs already do, but they have been enhanced and improved using tools that were developed by scientists seeking to imitate some *aspect* of human intelligence. These machines therefore do not imitate human minds, but they can provide a sometimes convincing imitation of some aspect of what human minds do.

While this is not the most pressing problem associated with the recent surge in generative AI tools and technologies, the misnaming of unrelated technologies as “AI” creates unrealistic expectations of where these technologies are headed. Alexa, Siri, and other technologies upgraded with AI are not in the process of developing consciousness and will not become so intelligent that they will threaten humanity. This is important to understand, because as long as scientists have dreamed of inorganic intelligence, humanity has been afraid of inorganic intelligence.

If humanity were ever to create a true artificial intelligence, it would be a moment of such profound importance that most of human science would pale by comparison. A true AI would be the first intelligent being ever actually created, in all known history. Biological intelligence is not created, it is grown from already living cells, and the processes that create intelligence do not involve the creation and combination of algorithms, but the collection of experiences. The creation of a true intelligence would not simply mean a new boom in advanced personal assistants or entertaining language simulation technology, it would raise serious ethical and moral implications about whether artificially created intelligence has rights and freedoms. The labeling of relatively simple technological advancements as “artificial intelligence” therefore raises fears and concerns that are not yet warranted, and may never be. Such unrealistic fears, reflecting deeper metaphysical insecurities, distract from the more practical discussion of how evolving technology is impacting American life and about the economic and security concerns raised by generative technology. While many Americans, still entranced with the fantasy of machine intelligence, will debate the issue in the press and on social media, it is far more important for society to, as it has been forced to do in the past, reckon with the more proximate effects of our evolving technological capabilities.

Works Used

- Anyoha, Rockwell. “The History of Artificial Intelligence.” *Harvard University*. Aug 28, 2017. sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/.
- Lanier, Jaron. “There Is No A.I.” *New Yorker*. Apr 20, 2023. www.newyorker.com/science/annals-of-artificial-intelligence/there-is-no-ai.
- Mineo, Liz. “If It Wasn’t Created by a Human Artist, Is It Still Art?” *The Harvard Gazette*. Aug 15, 2023.
- Ramge, Thomas. *Who’s Afraid of AI?* New York: The Experiment Press, 2019.
- Shieber, Stuart. *The Turing Test: Verbal Behavior as the Hallmark of Intelligence*. Cambridge, MA: MIT Press, 2004.

Notes

1. Lanier, “There Is No A.I.”
2. Anyoha, “The History of Artificial Intelligence.”
3. Ramge, *Who’s Afraid of AI?*
4. Mineo, “If It Wasn’t Created by a Human Artist, Is It Still Art?”
5. Shieber, *The Turing Test: Verbal Behavior as the Hallmark of Intelligence*.