

Artificial Intelligence Doesn't Exist

By: Steve Valentor

Artificial Intelligence (AI) is at the top of the news cycle. Is it smarter than we are? Can we control it? Should we fear it?

No one really knows how human thoughts are created. Similarly, we don't really know how human memories are formed or stored.

As of late 2024, here's what we do know:




A human brain contains about 100 billion nerve cells (called neurons). These are connected by 7 quadrillion or so synapses using 100,000 miles of myelin-sheathed fibers. Each communicates ~1000 signals per second. This amounts to more than 7 quintillion “bits” of information, almost an exabyte, per second. But remember, we don't really know if this information is in any way similar to binary bits.

Some information transfers are really slow, like the ones we use to see and hear. Others, like the ones I use to decide to procrastinate, are really fast.

Fuel for the brain is adenosine triphosphate (ATP). Our bodies make this from sugars and oxygen. By measuring the amounts that our brains use, we can estimate power consumption. Clue: It's more than a cellphone, but less than a super-computer.

The brain consumes 20% of the body's total energy. That's a lot considering that the brain comprises only about 2% of body mass. This equates to about 300 Watts per day. Since we sleep for about 1/3 of a day, we use about 200 watts during our waking hours. This is equivalent to about 6 smartphones. It's also about enough energy to drive a Tesla ¼ of a mile.



The El Capitan computer at Lawrence Livermore National Labs, currently the most powerful computer in the world continuously draws approximately 40 Megawatts (MW). So, to run it for the same 2/3 of a day that our brains are being actively used consumes 640 MW hours. That is equivalent to about 3.2 million human brains – or all of the brains of the entire population of Los Angeles – excluding most celebrities. El Capitan can perform about 2 exaFLOPS which is 2 quintillion floating point operations per second.

The fourth most powerful computer in the world today is the Eagle operated by Microsoft. But that's not the supercomputer being built for AI. Microsoft and OpenAI are reportedly working on a new system (and facility) called Stargate, scheduled to launch in 2028.

Stargate will cover many acres and continuously consume 5 gigawatts of power, making it 50 times more power hungry than El Capitan. Comparing that to human brains running for 16 hours, it is equivalent to the power consumption of the brains of all of the followers of either political party in the United States. That's enough to drive that same Tesla 37 million miles.

A University of California report estimates that a single query in ChatGPT consumes about 2.9 WH of power. That would project the query-capacity of Stargate at 17,000 per hour. If each one of those is equivalent to a human thought, it would quickly approach 5 "thoughts" per second.

That falls well short of the human capacity of processing 7 quintillion bits of information per second, which can be fueled by a small fraction of a Royale with cheese. But remember, we don't have any idea what a human thought is or how human memories are formed.



The term "Artificial Intelligence" was coined by mathematics professor John McCarthy of Dartmouth University in outlining the purpose of the summer-long workshop that he organized in 1956 to investigate the possibility of "thinking machines." This was far from the first time that humans looked to magical machines as a substitute for human thinking.

In Greek mythology, Hephaestus, assisted only by a cyclops, built Talos, a giant winged mechanized bronze robot capable of throwing giant boulders at any ships that dared to invade the Greek coast. In a similar way that we might simply unplug Stargate if it threatens us, Jason and the Argonauts defeated Talos by removing a plug in its foot which drained it of its other-worldly *ichor*, thought to be derived from the blood of the Gods.



There were a number of fictional attempts in the middle-ages to build thinking machines. Some used a little biology and chemistry. Paracelsus wrote that by mixing human sperm with horse dung and fertilizing it with human blood, a subservient but thinking human-like creature could be grown. Eleazar Rokeach described a process by which a clay golem could be animated by placing a piece of paper with a God's name written on it into the mouth of the clay statue. Onward from Goethe to Edgar Allen Poe, many thinking machines were created in fiction.

The belief that artificial intelligence is possible is based on the assumption that human thought is a process that can not only be understood completely, but also can somehow be mechanized. As medical research advanced between the 1930's and the 1950's, scientists and physicians began to hypothesize that human brains were a network of neurons transmitting all-or-nothing electrical signals. Nobel laureate Sir Roger Penrose posits the theory that human consciousness is not computational. He suggests that it is much more mysterious and requires quantum mechanics to explain. According to Penrose, thoughts may be fleeting coherencies of entangled pairs of quantum particles in superposition in microtubules located deep within neurons. These superpositions collapse into thoughts when we need them. His theory is probably as good as any other. Perhaps when the promise of quantum computing is realized, true artificial intelligence may indeed be possible. Since the first quantum qubit was demonstrated in 1998, scientists have predicted that useful quantum computers are no more than 5 years away. That's about the state of the technology today. But I digress.

Engineers including Alan Turing devised ways to employ digital signals to simulate bits of information. These were then organized into bytes, words, double-words, kilobytes, megabytes, terabytes, and soon exabytes in the Stargate system.

As computer hardware and software evolved, the general population became more astounded with every breakthrough. Once a computer was able to solve a maze by exhaustively testing every possible path, the news media of the day proclaimed that human thinking would soon be eclipsed by thinking machines. The need for humans to perform thinking jobs would vanish and we could begin to enjoy lives in total leisure.

The prophecy continued to project that this future would become reality within – you guessed it - the next five years. Intelligent computers could be miniaturized and built into robots that would take over all tedious, monotonous tasks. Human life could be reduced to pure decadent pleasure.

Such power would unlikely be used only to serve good and noble causes. Imaginations ran wild with the thought of how AI might revolutionize warfare. Governments began to make significant investments. DARPA provided MIT with \$2.2 million in 1963 followed by \$3 million per year through to 1970, with similar grants at Carnegie Mellon and Stanford.

It was critical that we humans get this AI thing right. As we learned in the 1968 film "2001: A Space Odyssey," the AI running on the computer HAL could easily kill humans if it perceived them as a threat.

With little progress and no commercial success, funding became limited and AI entered its first extended winter phase. This continued until 1980 with the dawn of "Expert Systems" and "Inference Engines." This was my first professional exposure to AI.

These systems attempted to capture all of the known knowledge about a particular subject and organize it into a database. It could then be used to follow logical steps similar to how an expert

consultant might assist a client to solve a particular problem. This is exactly what DynaQuest, the startup that I worked for did.

During the 1960s and 1970's, universities, government agencies, and large businesses bought multi-million-dollar mainframe computers from companies like IBM, Honeywell, and Digital Equipment Corporation (DEC). These systems required giant teams to operate them. Before the personal computer was introduced in 1981, smaller businesses typed letters, invoices, checks, and receipts on typewriters and mailed them through the postal service. Now that these companies could afford to buy small computer systems for as little as \$8000 (the price of a fairly expensive car in 1981), they had to make decisions about how to buy and configure those systems. They typically hired consultants who could easily charge multiples of the system cost for advice.

DynaQuest sought to automate that process by using an expert system. Here's how it worked.

A sales person would interview the client and capture several pages of information. This was called a "profile." It included information such as how many letters they might mail each month, how many invoices they sent, how many receipts they created, whether they might need a data connection, and so on.



The profile would then be processed by the "expert system" which generated a 75-page report overnight. The report, which was delivered for only \$175, instructed the business owner how to purchase and configure the perfect computer system.



At this time, the internet existed, but was not available to the public. Computer communications was done with dial-up modems which communicated at 300 bits-per-second. My job as a mathematician was to compress the "profile," upload it to the mainframe computer, run it as a batch process on the "expert system," then compress it, download it, and finally print out the finished report.

There was a story in the Wall Street Journal about how DynaQuest would render the business consulting market obsolete. The article continued that there were no limits to the power of this new technology. It could be applied to robots, and in short order, most human jobs would likely be eliminated.

In April of 1983, Omni magazine declared that intelligent robots would change the way we work and live.

DynaQuest failed in less than a year.

While expert systems, inference engines, and other incarnations of AI during the 1980s continued to solve very limited problems, these solutions became referred to as “data processing” and would no longer warrant the sizable investments that had been made in AI. By 1985, the second AI Winter had set in. The US government “Strategic Computing Initiative” didn’t see AI as viable and dramatically reduced funding. Even DARPA had abandoned AI in favor of what many referred to as science fiction – Ronald Reagan’s Star Wars Defense Initiative which was known as SDI.



Through the 90’s and into the new millennium, a number of technologies marketed as AI came and went. There was Fuzzy Logic, Natural Language Processing, Symbolic LISP, neural networks, biological circuits...

And in actual 1984, James Cameron showed us that a time-travelling AI cyborg from a future Orwellian dystopia would look just like Arnold Schwarzenegger and would kill humans indiscriminately to preserve its own future. In 1987, we learned that a vestige of a human brain could be integrated into a cyborg and that a reliable RoboCop could be fabricated and trusted.



But then in 1997, IBM’s Deep Blue defeated chess master Garry Kasparov and the R&D spigot cracked open again. By 2004, Will Smith found a way to collaborate with a cyborg, though the film “I, Robot” was projecting life in 2035. In the 1999 film, “The Matrix,” humans were needed as a cheap, renewable energy source for the computers that ruled the majority of human existence.

Meanwhile, back in the real world, Apple introduced SIRI in 2011 and Amazon launched Alexa in 2013. They were called “digital assistants” and you could start a conversation with them by calling their name. Google went a step further in 2016 when they launched their version which was touted to respond to natural language. Even after more than 13 years of technological advances and extensive training, they still cannot transcribe a voicemail reliably. And I’m confident that most people continue to be embarrassed by apparently context-ignorant “auto correct.”

In 2015, a group of current and future billionaires formed the non-profit OpenAI with a combined pledge of \$1 billion. Elon Musk resigned from the board in 2018 citing a conflict of interest with the development of AI for the self-driving modes of Tesla cars. By 2019, the

billionaires had only contributed \$130 million, and it became clear that OpenAI could not continue without a profit-seeking component. When they made the transition, Microsoft invested \$1 billion in OpenAI then followed in 2023 with an additional \$10 billion. With that much money, and no immediate pressure for revenue or profitability, a tremendous amount of development continued into 2024.

At that point, Elon Musk sued OpenAI for breach of contract, alleging that OpenAI had strayed from its not-for-profit mission and closed its source code to outside developers. Much like the communities around Linux, Python, and PHP, OpenAI intended to provide open source code and free licenses which




would ultimately be managed by the community that formed around it. Musk later agreed to drop the suit if OpenAI was willing to change its name to ClosedAI.

Considerable progress has been made in text-to-image functionality, music composition, and generally fooling the public again into fearing sentient machines. In a rapidly approaching eventuality, these machines will determine that the human race is a threat to be first made irrelevant, then completely eliminated.

Before we get too carried away, we need to look at a few of the limitations of the current systems. AI scored in the 93rd percentile on the SAT reading test and in the 89th percentile on the math test. It scored in the 90th percentile on the bar exam. It hasn't yet applied to college or law school.

With successes like these, it is easy to believe that AI has advanced reasoning skills. However, that's not the way AI operates. It is trained on enormous volumes of data. It looks for patterns in the data, then predicts other patterns of data that are most likely to be associated with that original data. If it is trained (which it is) by seeing many SAT and bar exam questions, it will perform well on the exams simply because it has memorized the answers.



On the GSM8K, which is a word problem test designed for grammar school students, AI performs very well. A group of Apple researchers published a paper in late 2024 where they tested the reasoning skills of the state of the art AI systems. They tested them with revised

versions of the GSM8K exam. They found that when they changed the proper names and specific numbers, AI continued to do very well.

However, when they added information that seemed as if it might be important to the problems, but was totally irrelevant, the AI scores dropped significantly. The normal version of the test might contain this question:

Q: If John drives a car at 60 MPH for three hours, how much distance does he travel?
A: John travels 180 miles

The modified question might be:

Q: If Janelle drives a car that gets 40 MPG at 80 MPH for two hours, how much distance does she travel?
A: Confusing or ambiguous information prevents a clear result

The accuracy of the results drop by as much as 65%!

This is exactly the sort of complication that arises in driving a car. I recently tested Tesla's latest Full Self Driving (FSD). Much of the time it is remarkably capable. At other times, it drives about as well as an angry, possibly inebriated teenager. With no warning, it will appear to lose its self-awareness it will do something like make a right turn from the left turn lane – while switching from its left to right turn signal. In my opinion, it is not ready for unsupervised operation. I can envision a



day when dynamic communication between FSD-equipped vehicles communicate with each other, intelligent infrastructure, and external safety systems. But I feel that the dawn of this nirvana it is not consistent with the aggressive schedules projected by proponents of the technology. When it works, I am all-in.

Then there is sarcasm. It is unlikely that Large Language Model (LLM) systems, even with the most sophisticated voice-to-text synthesis will be able to detect sarcasm. For example, consider the statement “Yeah, right.” It might mean that the speaker is agreeing with something that was said. “Yeah Right.” But even with the exact tone and inflection, it might mean that the speaker is sarcastically disagreeing. “Yeah Right.” I concede that this situation may also be difficult for humans to navigate.



There is no doubt that a Roomba from iRobot can safely clean your floor without injuring pets or children. These systems are equipped with sensors and control software to avoid dangerous situations. Possibly more dangerous



are robotic lawn mowers with high-speed spinning blades like the Husqvarna 415X that uses GPS data to navigate your lawn. Interestingly, neither are marketed as having AI capabilities.

Here is what I think it all may boil down to. For millennia, we humans have been seeking thinking machines to make our lives easier. We have created all sorts of them, and they are truly amazing. With each advance, our lives get a bit easier. When it comes to AI, I don't think it really exists in the sense that it is an artificial recreation of human intelligence. I do think that it is the state of the art of data processing.

Part of the reason I feel this way is that I have many decades of experience developing processor based computers, chips and software. Every computer computes binary bits of information the same way. Deep in the core of the processor, below the application software, the compiler, the linker, the application program interface (API), the operating system, peripheral devices, device drivers, the Basic Input Output System (BIOS), interfaces to memory, and even the registers where bits are stored for immediate computing, lies the Arithmetic Logic Unit (ALU).

It is in the ALU that arithmetic and logical operations are performed. All arithmetic can be reduced to addition. Subtraction is just addition with a minus sign, multiplication is just repeated addition, and all other math functions can be built from these. Graphics Processing Units operate exactly the same way, except that they have hardware which is optimized to perform many math operations simultaneously. For example, in less than a single clock cycle, they can perform a multiplication and an accumulation of the result, known as a multiply-accumulate. These functions dramatically speed up calculations in data searches (used in AI), encryption, cryptocurrencies, block chain and graphics.

Paradoxically, this has driven the Nvidia stock price to a point where it will have to return three times its current **revenue** as **profit** to justify its current valuation of more than 331,465 times its IPO value.

NVIDIA Corporation (NVDA)

☆ Follow


↔ Compare

🕒 Time to buy NVDA?

145.06 -0.07 (-0.05%) **144.76** -0.29 (-0.20%)

At close: 4:00:00 PM EST

After hours: 5:17:38 PM EST



Back to the technical detail. The other part of the ALU performs logical functions on bits. These include AND, OR, NOR, and XOR. There are plenty of videos online that explain this quite clearly.

In order to have an approximation of human intelligence, we would first have to find a way to recreate human thoughts and memories. You may recall that we don't really have any idea what these are or how they are formed.

No matter what the AI system appears to be doing at the application level, at its lowest level, all it can do is add, or perform logic on individual bits. Let's say that you have two bits. Each one can be either 0 or 1. Logical AND and OR operators generate the following truth tables.

Bit 1	Operation	Bit 2	Result
0	AND	0	0
0	AND	1	0
1	AND	0	0
1	AND	1	1

0	OR	0	0
0	OR	1	1
1	OR	0	1
1	OR	1	1

Despite being abstracted into high level languages like Python, which was the most prevalent language used to develop current AI systems, all instructions compile down to these basic ALU functions. How one might find intelligence to emulate human reasoning in this remains a complete mystery to me. And I have been working with this stuff for 40 year.


At that top level, we do have very application-specific expert systems and large language models. Both can generate information that is statistically likely to be appropriate in various situations. They may be wrong. But most of the time they will be reasonably correct. These systems are improving rapidly. In most situations, we can rely on them. But every so often, if we allow it, they will try to make the car turn left from the right lane, cite a legal precedent that does not exist, make you late for a meeting by provide clear directions to a baseball field (rather than a Raytheon facility), lock Dave out of the spacecraft, terminate Sarah Connor, create a Matrix to appease humans so they can be used as a power source, fire a missile, or determine that humans are no longer required.

At the point that we give these machines too much autonomy, we need to be sure that we can unplug them. In aviation, systems that use software are qualified in five levels of Design Assurance defined in Document Order 178 (DO-178).

Design Assurance Level	Condition	Objectives
A	Catastrophic	71
B	Hazardous	69
C	Major	62
D	Minor	26
E	No safety effect	0

In order to prevent the AI computer system HAL from locking the astronaut Dave out of the spacecraft in *2001: A Space Odyssey*, NASA would have rated the failure to open the spacecraft door as “Level A,” where a failure is deemed “Catastrophic” and required that the software running on HAL satisfy all 71 objectives. For Level A, 30 of the 71 objectives must be “With Independence.” This means that the person who reviews the code must be different from the person who wrote the code. The requirements go much further beyond the code author and reviewer. The test procedures and testbench apparatus must meet the same objectives criteria.

If we allow that HAL might have some form of machine learning (ML) and might reprogram itself to rewrite its own software in such a way that it no longer would meet the objectives, then redundancy would be required. It is likely that the system designers would require a manual override to allow Dave to open the door independent of HAL. I would certainly demand it if I were designing the spacecraft.



Here is a great use of AI where there is no risk of ML causing violence or death. I asked ChatGPT to generate a picture of German Shepherd that looked like Einstein so that I could amuse my family with the potential dog name: K-9-Stein for this exceptionally smart breed. No creatures were harmed in the creation of this image.

Far more annoying is the application of AI in education. At the university where I teach, administration has licensed “AI Detector” from Turnitin and made it available to faculty. According to the Turnitin website, more than 130 million papers have been processed, 3.5 million have been identified as containing more than 80% AI content, and 12.5 million have more than 20% AI-written content. And this data is as of October 2023!

Every semester, I regularly assign writing tasks to my students. As an experiment, I “assigned” one of them to an AI bot. It generated B-level work. I then refined the instruction to add formatting, footnotes, and references. It produced A-level work. When I read it carefully and checked the references that were cited, I found major discrepancies. In some cases, the references did not contain the material being referenced. In other cases, the references were not even relevant to the topic. In many cases, the references were written at a level far beyond the competency of my 200-level students and did not relate to the assigned topic. Most often, the references were just fabricated and didn’t exist.

One thing that I am certain of is that most students who use AI to generate their writings learn far less than students to complete their own original research. If a student uses AI to generate even a first draft, they must still complete the research, read the references, and ensure that they are relevant. Some might argue that this takes more effort than completing the research directly.

Extending this metaphor to the battlefield, we might find that one side has tired of sacrificing soldiers to enemy fire. So they build a swarm of Unpiloted Air Vehicles (UAV) and program them to seek out enemy uniforms and destroy the people wearing them. This is well within the capabilities of today's technologies. After a horrific battle, the other side is resolute not to let that happen again. So they build a force of UAV-killing robots and program them to seek out and destroy the enemy UAVs.



In arrogant retribution, the opposition increases the level of armaments and reprograms its next swarm of UAVs to seek and destroy the enemy robots. The volleys continue until both sides have run out of AI weapons which lay in a giant pile of broken technology on the battlefield (or in a landfill with all the broken renewable energy equipment used to power it all – but that is a different topic).

At that point, if the humans that created the AI systems can remember what they were fighting about in the first place, they climb over the piles of AI carnage and begin hand-to-hand combat until their issue is resolved.

As casualties compound on both sides, they quickly realize that the effort is futile. They then resort to AI to wage a cyber-war against their opponent. They use AI to feed mind-altering material to their enemy's children, hoping to turn them into sympathizers as quickly as possible. In retaliation, the other side uses AI to wreak havoc on their enemy's banking system. Both sides then escalate, destroying each other's infrastructure, education, government, food supply, and medical care systems. In the end, neither side has sufficient resources to wage additional warfare.

Since both sides have converted each other's youth, they now continue to battle in hand-to-hand combat, but on opposite sides of the original cause. Scary. No robots were harmed when I used AI to safely generate this battle-scene image.

Back to fiction. I asked a free AI image generator to create an image of Jason and the Argonauts fighting Talos. It generated this.

Then, I asked it to summarize the end of Jason and the Argonauts. In the second paragraph of its four-paragraph summary, without any specific prompting from me, it pointed out that Jason and Medea defeated Talos, the bronze, boulder-throwing AI robot using one of the few resources they had left, magic.



Does it require magic to control AI? NO!

My feeling is that AI is just software performing data processing. Every time a computer does something that is difficult for humans to do, we hail it as a technological marvel and declare it to be "intelligent." Then we become accustomed to it, and it falls to the mediocrity of being just

data processing. Indeed, in 1986 when I toured the British Telecom research facility at RAF Martham Heath, I was shown their latest voice recognition system. It ran on a giant mainframe computer, and it was remarkable that the system could recognize a vocabulary of more than a dozen words, even when spoken with my American accent. The software on the iPhone that transcribes voice mails has a much larger vocabulary, but may not be as useful.

Three researchers at the Hebrew University of Jerusalem trained one of these artificial deep neural networks to perform the same calculations as a simulated neuron. Keep in mind that the software construct of a “neural network” is what scientists of the 1950’s surmised that human neural networks must contain. Yet - we still don’t know how they function.

Regardless, to build a neural network to simulate the function of a single simulated neuron required between five and eight layers of interconnected constructs, which are also called neurons.

We have learned in the last 70 years that actual neurons are unknowingly more complicated than the concept used for the neuron model in this exercise. Between five and eight layers of 256 artificial neurons in each layer were required to accurately simulate the function of a single neuron model with 99% accuracy at a one millisecond level. This equates to approximately 1000 artificial neurons required to simulate the function of a limited model of a single biological neuron.

We now know that biological neurons are dramatically more complex, and we do not have any way to model the input-output function of biological neurons. So, will AI replace human intelligence? I doubt it!

AI is just data processing. It’s fun to play with; can save time on some tasks; is as unreliable as the information used to train it; can be managed with redundancy; can be unplugged if it misbehaves; and will eventually become just software processing data.

Computers performing data processing are amazingly fast. But I do not see a path to computers developing the ability to “reason” until some new paradigm emerges. Perhaps it will be quantum computing – in five years.

If you are considering investments in these and other exciting new technologies, you should consult a Registered Financial Advisory like my firm, Polynomial Ventures. Visit www.polynomial-vc.com or schedule an introductory meeting through the QR Code below.



Steve Valentor is a 30+year technology industry veteran who has worked in computer engineering, semiconductor R&D, software development, and manufacturing for companies ranging from startups to the Fortune 200. He has held positions from entry level engineer to senior technical management, CEO and board chair. Currently the managing partner of Polynomial Ventures and an adjunct professor at DePaul University, Valentor holds an M.B.A. in finance and a bachelor's degree in math, both from the University of Illinois at Chicago.

Polynomial Ventures invests venture capital in early stage technology companies outside of Silicon Valley and Boston. The Chicago-based firm is a registered investment adviser (RIA) providing a range of wealth management services.