# Artificial Intelligence: Trusting AI?

This report is the fourth I have written on Artificial Intelligence (AI). The reports are online at Blog.UylessBlack.com. To download them, enter **Blog.UylessBlack.com** in your browser. My blog homepage will appear. Scroll down to **21. Computers and Networks**, where you will find links to the articles, including this one.

The first three articles on artificial intelligence brought forth a lot of responses. They varied greatly. Some were (are) baffled by AI. Some said they were reading or watching news about AI finding its way into practically everything. One person said, “Every time I check the news, AI is usually one of the topics.”

This comment substantiates that artificial intelligence is a hot topic in the media. The wide coverage reflects the fact that many governments and commercial enterprises are using or planning to use AI.

Many of the major news magazines cover stories have been about artificial intelligence. Recently, I received a copy of *Time*, which has the subject of AI as a cover story, as shown in Figure 1.

**Figure 1. AI on *Time* cover.**

Some responders expressed concern about AI, how it might have a negative impact on their lives. I am sure this last point is at least partially because my three articles on AI were written with the intent to raise the level of awareness about the potential gains, but the risks as well, of using AI.

This article deals with the responses I have received thus far, as well as some additional thoughts about AI. It is divided into general, non-technical discussions and somewhat technical discussions of AI technology. The latter is identified as sidebars with the text inside a box. Those who do not wish to go into this level of detail can skip around the sidebars and read the general comments. First, some observations, based on your responses. We begin with an example of AI use by Shin Bet, Israel’s internal security agency.

## Israel’s use of AI in the Gaza War[1]

Shin Bet is considered by many nations and security experts to be one of the best security/intelligence services in the world. Thus, most of the world was surprised that Israel was surprised by the Hamas attack on 7 October 2023. Israel and Shin Bet did not anticipate the attack and was not prepared for it.

[1] The sources for the information in this section and other parts of this essay are from:
(a) https://www.bing.com/search?PC=YF73&q=gaza+war&FORM=YF73DF,
(b) https://en.wikipedia.org/wiki/Gaza%E2%80%93Israel_conflict,
(c) “The War Lab,” *Time*, 30-35,
(d) “The Pentagon’s Silicon Valley Problem,” *Harper’s Magazine*, March 2024, 25-30.
Unless otherwise noted, direct quotes are taken from source (d).

Shin Bet had extensive information pointing to the attack. Hamas openly carried out practice assaults on "…mock-ups of the border fence and Israeli settlements." Videos were posted on social media about these activities.

Months before the attack, Shin Bet announced it had developed its own AI system, one based on the widely used ChatGPT, developed by OpenAI. (More information on ChatGPT and similar AI systems is provided in the sidebars.) ChatGPT, and its newer off-shoot, GPT-4, rely on having access to a huge amount of data it can process in order to make its decisions known to humans. Fittingly, this kind of AI is called a *large language model* (LLM). If these data are not accurate, an LLM will most likely produce unacceptable, perhaps damaging output. Just like the pre-AI days: Garbage in, garbage out.

Ronen Bar, the Shin Bet director, gave a speech about this system. A preview of his speech was made available on social media through the news site, Tech12. Here is an excerpt from this speech:

> The system knows everything about [the terrorist]: where he went, who his friends are, who his family is, what keeps him busy, what he said and what he published. Using artificial intelligence, the system analyzes behavior, predicts risks, and raises alerts.

These confident assertions were made based on the reliance of a large language AI model. As stated earlier, like all automated systems, going back to the times we humans first used them: Garbage in, garbage out.

Through informants, Hamas knew about some of the Israeli intelligence practices. Enough, that for this attack, Hamas created erroneous information about what they were up to.

> They [Hamas] signaled that the ruling group inside Gaza was concentrating on improving the local economy by gaining access to the Israeli job market, and that Hamas had been deterred from action by Israel's overwhelming military might. …AI, as it turned out, knew everything about the terrorist except what he was thinking.

Regarding that last sentence: It was not the fault of AI. It was the fault of humans mis-using AI. "…Hamas would short-circuit Shin Bet [AI] algorithms by feeding the system false information (AI algorithms are explained in the sidebars).

**Sidebar One: Beyond Conventional Data Processing**

Most AI systems provide their results based on being fed massive amounts of data (*data mining*), a process called *training*. The more data, the better the AI software can perform; the more it can *learn* from this data. Learn, in the sense of using the data to create more information beyond what the data itself might yield in providing information to a user. In non-AI systems, programmers' write software code that determines the software's output to users. In an AI system, the software may very well make some of its own decisions in determining the output.

Of course, as stated in the first AI article, most any software application could be considered AI-capable if it produces output an ordinary human (or groups of humans) cannot produce. And if, as stated in this first article, the AI-based computer system is able to perform tasks that normally require human intelligence.

AI goes further than ordinary computer systems. Like many computer programs, it takes in data, but it analyzes the data to correlate millions of pieces of the data into *patterns*. Assembling and making sense of this vast amount of information is usually beyond our cognitive powers…not to mention our physical stamina and persistence.

Most AI experts explain this process by stating that a developed pattern is an *AI algorithm*, which is used to analyze and correlate similar data, say thousands of human faces fed into a face recognition application. These computations go on and on, consuming immense computational resources and electricity.

To summarize, the difference between the conventional data processing method of simply grinding away at possibly massive amounts of data and AI, is that AI software learns from its processing and, as one AI expert puts it, "improves itself." Maybe so, but any such improvement is dependent on the data given to the software.

Nonetheless, given enough data, say, human faces, an AI facial recognition application can tell the difference between one face and another with amazing accuracy.

Conventional computer data processing systems are not designed to learn and possibly improve on the data they process.  AI systems are.

## AI Vendors are *Feasting* on *Your* Data

As discussed, large language models rely on being able to process massive amounts of data, often called *tokens*. Tokens vary, depending on the particular AI system that is processing them. For this discussion, GPT-4

> …was trained on as many as 12 trillion tokens, or words or parts of words. The next iteration of the model, GPT-5, could require up to 100 trillion tokens. That's more than all of the useful language and images available on the web.[2]

[2] "Scarcity: The Great AI Data Gap," *The Week*, April 26, 2024, 20.

Let's pause for a moment to consider the implications of this last quote. Embedded within this reality is the fact that much of the web's content is data (information) about you and me. Much of it is personal information, which Internet vendors use extensively for advertising and tailoring social media interactions with users. As I wrote in *Fractured*, user data, created by people accessing the Internet's social media and websites, provide the silage on which many Internet vendors feed for their existence.

**Sidebar Two: AI's Learning Process**

AI uses a process called *machine learning*. One approach in AI-based machine learning, among others, is the use of mathematical methods (risk minimization, backpropagation as examples) that help the AI system learn from the data at its disposal, and to also use this data to *generalize: Taking the known data and formulating other data*.

This other data, often called *unseen data*, translates into the power of AI: It can "perform tasks without explicit instructions,"[3] without a software programmer's code. Thus, certain AI implementations can become independent and act on their own. As Arthur Samuel said during the early years of computers and software (1959), performing computations "without being explicitly programmed."[4]

Another key component of AI systems are *neural networks*. For machine learning in computers, and not humans' nervous system, a neural network uses mathematics to approximate nonlinear functions. (A function in which the change of the output is not in proportion to the change of its input.)

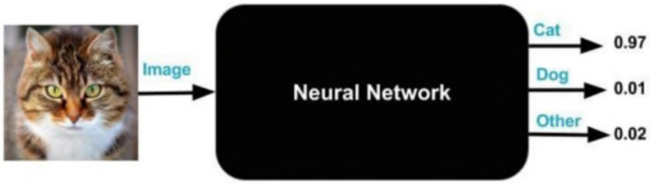

Nonlinear systems' output, and associated neural network-based AI systems, may not be predictable and may not contain completely accurate output. "They can appear to be counterintuitive, unpredictable or even chaotic."[5]

The picture above is an illustration of the (potential) problem.[6] In this example, the neural network is not 100 percent accurate. One might say, "Few things in life are." True, but that begs the point. (Besides, a conventional computer program, one properly debugged (checked for errors), is 100 percent accurate.) The point being that you likely do not want to be flying in an airplane designed on purpose to be 97 percent reliable.

However, that does not mean AI methods should not be used. Quite the opposite. In some applications, they are preferable to conventional data processing programs. Of key importance, an AI system is wholly application and data specific. As an example, an AI application is not intended to predict crop yields, and at the same time, discern the image of the cat, shown above.

I refer you to an excellent analysis of the pros and cons of these systems, available at the website cited in footnote 6.

[3] For more information on generalization and AI's performing these tasks, see (a) Richard Nordquist, "Definition and Examples of Hypernyms in English," https://www.thoughtco.com/hyponym-words-term-1690946.
(b) "Scale and Generalization," Axis Maps, October 14, 2019.

[4] A.L. Samuel, "Some studies in machine learning using the game of checkers," *IBM Journal of Research and Development*, 44, 206–226.

[5] See "Nonlinear Dynamics I: Chaos Archived," MIT OpenCourseWare: The Wayback Machine, 12 February 2008.

[6] https://builtin.com/data-science/disadvantages-neural-networks.

Consequently, user data, and other kinds of data (such as news broadcasts and advertisements) are vital for AI tokens, for AI large language models. Aside from the issues of impinging upon user privacy and rights to "owning" one's own information about oneself (addressed in other reports), the social media vendors who have captured this invaluable information resource are becoming more committed to claiming ownership of this data. Ownership in the sense of denying others access to it.

> The supply [to AI systems] is shrinking further because "social media platforms, news publications, and others have been curbing access to their data for AI training"[7]

"Their data." A substantial amount of their data is user data; that is, your data and mine. We should be aware of this fact and perhaps be a bit more discerning about how we use social media. I try to avoid using it.

**Synthetic Data**

To find more data, AI companies are searching far-and-wide. Meta is rumored to be considering purchasing publishing houses, such as Simon & Shuster, as part of their data mining endeavors.[8]

While they were in operation, older systems, such as MySpace and Friendster, captured a wealth of information from millions of their users. AI companies are negotiating with these (defunct) enterprises for their data. Billions of photos and videos, not to mention text, have been archived by these organizations.

Notwithstanding these efforts, what to do when a vital resource to an organization's well-being is in danger of becoming exhausted? One alternative: Make something up. Create a substitute, something *synthetic*. That is what occurred as a run up to the financial meltdown of 2008. The banking and financial investment industries invented instruments out of thin air, ones that had no solid assets behind them. They sold them as if they were backed by secure commodities, such as bonds and mortgages.

AI vendors, looking for training data for their models, are experimenting with *synthetic data*. This kind of information is not actual data, it is created by the AI system.

> …essentially [this means] AI is training itself. That's a problem, especially if the mentoring AI is biased or inaccurate, which is often the case. Researchers have compared the practice "to the deeply inbred Habsburg dynasty" and worry it will

[7] Ibid.

[8] As an aside, this writer reminds AI vendors that he owns the copyrights to his published books. His publishers do not. So, Meta and others, do not go to Pearson Publishing, McGraw Hill, or my other publishers to negotiate using the "data" in my books. I'm open to your offers, if you can convince me that you have found ways to mitigate your AI systems from hallucinating erroneous output, and we can jointly figure out royalty/Fair Use Law arrangements.

create an "inbred mutant" AI.[9] But such companies may have no alternative if the quality data pool runs dry.[10]

At this stage of the game in which artificial intelligence is in its formulative stages, AI users having any influence on their AI vendors' products should require the vendors to divulge their source and kind of data used to train their products.

**Sidebar Three: Human and Artificial Neural Networks**

Continuing the discussion of the role neural networks have in artificial intelligence, a short explanation of neural networks in humans will be helpful. The network in our brains is make up of neurons (nerve cells) connected to each other across synapses, billions of them. Working together, neurons and synapses are the fundamental tools for human thought.

Research going back to the 1940s concluded that "neural networks can *change and learn* over time by strengthening a synapse every time a signal travels along it."[11] This process leads to our brains becoming, as described by the experts, more elastic.

As discussed in this series on artificial intelligence, that is what computer-based artificial AI systems do: *change and learn*. So, artificial neural networks and artificial intelligence are aptly named.

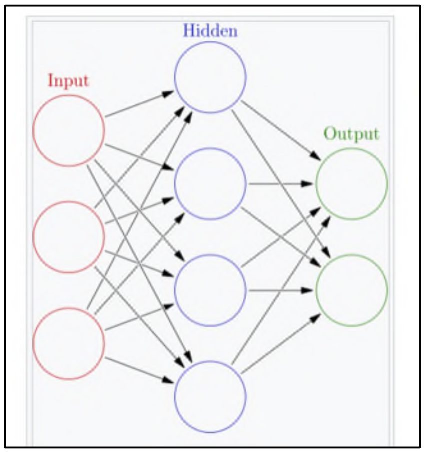

Like the brain, in which electrical and chemical signals are passed between neurons, information is passed in an artificial neural network through software generated storage areas called *artificial neurons*, or in a less esoteric description: nodes.

As shown in the figure on the left, an artificial neuron receives and sends *signals* from/to other neurons, based on mathematical calculations performed in the software.

This signal, the output of the neuron, is a number that is computed by a *nonlinear* function of the sum of all of the inputs into that neuron. (Don't forget: Neural networks based on nonlinear algorithms may not be predictable and may contain errors from faulty training data.)

The intermediate calculations in the nodes are aggregated into *hidden layers*. They are called hidden because the layers may perform different alterations on the input to the layers without human intervention. This figure shows one layer, but signals (data) might pass through more than one intermediate layer. If so, AI designers call it a *deep neural network*.

[9] Centuries ago in Europe, the Habsburg family increased their power across countries by frequent marriages among close relatives. This practice compromised their gene pool. The inbreeding resulted in epilepsy, insanity, facial distortions, and early deaths. The family learned, too late, that "kissing cousins" had its limits.

[10] Ibid, *The Week*.

[11] D.O. Herb, *The Organization of Behavior*, (New York: Wiley & Sons, 1949).

## An Intelligence Approach to Using Artificial Intelligence

In 1956, Lewis Mumford, a widely known and respected social commentator in his day, published a book titled *The Transformations of Man*. In this book, he offered this thought (paraphrased): *One of the functions of intelligence is to take account of the dangers that come from trusting solely on intelligence*.

For our examination of artificial intelligence, we can alter Mumford's observation: *One of the functions of intelligence is to take account of the dangers that come from trusting solely on artificial intelligence*.

Israel's AI system failed to warn about the attack from Hamas because the Shin Bet intelligence experts were too trusting of artificial intelligence.