

This copy is for your personal, non-commercial use only. Distribution and use of this material are governed by our Subscriber Agreement and by copyright law. For non-personal use or to order multiple copies, please contact Dow Jones Reprints at 1-800-843-0008 or visit www.djreprints.com.

<https://www.wsj.com/tech/ai/artificial-intelligence-questions-answered-137f116c>

Readers Have a Lot of Questions About AI. We Answer Them.

From the basic to the more complex, there's still a lot of confusion about artificial intelligence

By [Bart Ziegler](#) [Follow](#)

Jan. 20, 2024 11:00 am ET

Artificial intelligence seems to be everywhere these days. It's going to have a profound effect, the experts say, on pretty much everything.

But many people still find the technology baffling.

We asked Wall Street Journal readers what questions they have about AI. They had a lot of them. Many expressed confusion about its implications, its risks and the best way to harness it for good.

Below are answers to some of their questions.

I really do not understand artificial intelligence and where it's heading. Please explain what it is.

Artificial intelligence is a broad term, covering a multitude of technologies, so definitions differ. But it helps to think of it literally: It is digital technology that mimics the analytical ability of human intelligence, largely by finding patterns in the information it is fed or encounters.

That means it can interpret vast amounts of information quickly, solve complex problems, or control complicated technical processes.

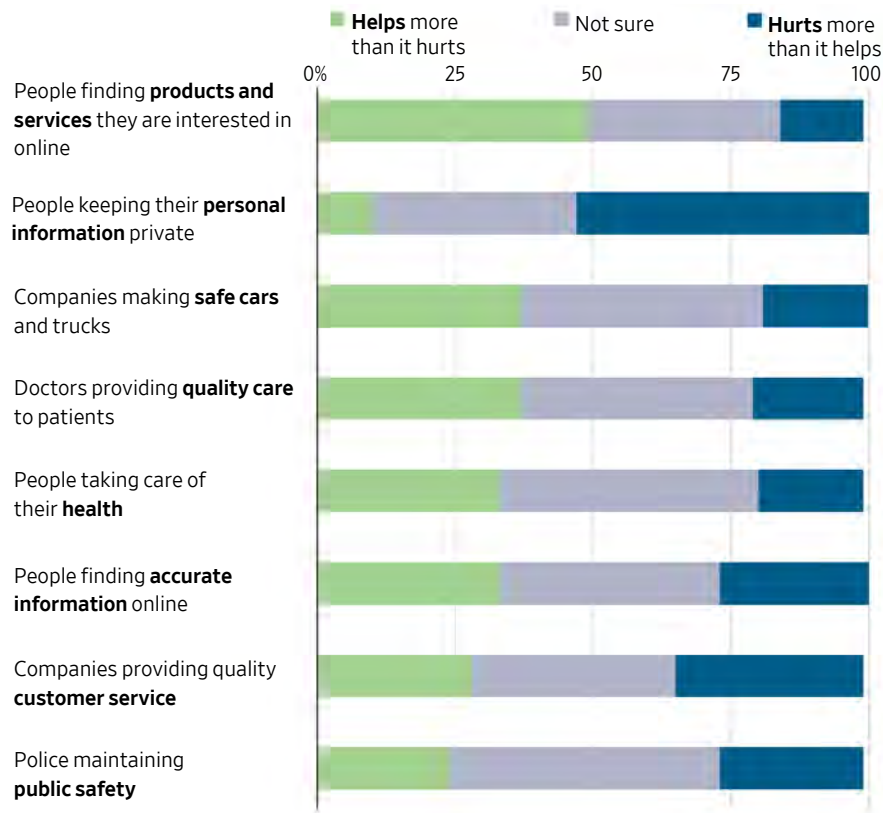
An early example of a machine demonstrating something like thought was the defeat of world champion chess player Garry Kasparov by the IBM computer Deep Blue, a precursor of AI, in 1997. Today, artificial intelligence is woven into our lives. It creates the personalized recommendations you see for products on Amazon and videos on Netflix. It helps power Apple's [AAPL 1.27% ▲](#) Siri and Amazon's Alexa. Scientists use AI to design drugs, retailers to provide customer service and manufacturers to optimize production.

What about those "chatbot" AI systems?

These systems allow us to ask questions on almost any subject and have them answered instantly in a conversational manner. Depending on which chatbot you use, they can write custom essays and reports, build recipes based on ingredients you enter, write software code, create images based on your descriptions, solve mathematical equations—even compose poems, plays and song lyrics.

Unlike most earlier AIs, they are available over the internet to anyone and are offered in free versions. That's spreading AI to a vastly larger audience. Among them are OpenAI's ChatGPT, Google's Bard, Microsoft's [MSFT -0.37% ▼](#) Copilot (formerly called Bing Chat) and Anthropic's Claude AI.

Percentage of U.S. adults who say for each of the following: “Artificial intelligence...



Note: Respondents who did not give an answer are not shown.
Source: Pew Research Center survey conducted July 31–Aug. 6, 2023

How can these new AIs answer almost any type of question?

These powerful computer systems, called generative AI or large language models, are fed enormous amounts of information—hundreds of billions of words—to “train” them. Imagine if you could read pretty much everything on the internet, and have the ability to remember it and spit it back. That is what these AI systems do, with material coming from millions of websites, digitized books and magazines, scientific papers, newspapers, social-media posts, government reports and other sources.

These systems break the text into words or phrases and assign a number to each. Using sophisticated computer chips that form what are called neural networks—which mimic the neurons in a human brain—they find patterns in these pieces of text through mathematical formulas and learn to guess the next word in a sequence of words. Then, using technology called natural-language processing, they can understand what we ask and reply.

What are the implications of this for average citizens? I don't know whether to be uneasy, frightened or welcoming.

Maybe all three. On the upside, AI chatbots provide information that is more concise, on point and useful than the jumble of links you get with a web search, and their ability to converse like a human is uncanny.

But there are big downsides. For one thing, AIs don't think **in the way humans do**. "This thing is not intelligence. It has no understanding of what it's saying," Usama Fayyad, executive director of the Institute for Experiential AI at Northeastern University, told the Wall Street Journal's "The Future or Everything" podcast this month.

Because an AI operates like an autocomplete device, it can guess wrong about what the next word should be in its responses, Fayyad says. And it compounds that mistake by using that wrong word "as part of the input and builds on top of it."

Moreover, if the training material included conspiracy theories, bigoted or fake information—which is rife online—AIs could parrot that back.

Also, the training text generally is months or even more than a year old. That can lead to outdated answers. Some chatbots, such as Microsoft's Copilot and Google's Bard, can use internet searches to add up-to-date information to the training-based information. Still, Bard told me in late December that Mike Pence was exploring a White House run—he had dropped out in October.

And they can contradict themselves: While shopping for a bicycle I asked another chatbot to compare road bikes to hybrid bikes. It told me that road bikes have so-called drop-style handlebars—but a few sentences later it incorrectly said they have flat handlebars.

What about "hallucinations," when an AI makes up something but presents it as a fact?

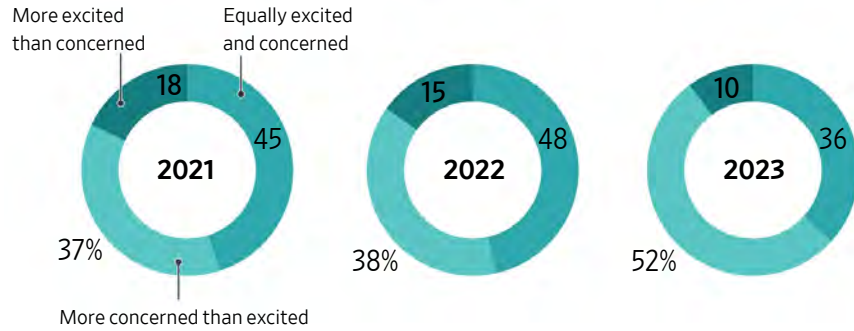
Hallucinations can arise when an AI system finds patterns in its training material that are irrelevant, mistaken or aren't meaningful, something that experts call noise. "AI hallucinations are similar to how humans sometimes see figures in the clouds or faces on the moon," according to IBM.

Hallucinations can trip up chatbot users. A Manhattan federal judge in June fined two lawyers who used ChatGPT to search for prior cases to bolster their personal-injury lawsuit. The cases they cited turned out to be fake.


Even AI promoters can get burned. In February, Google did an online promotion for Bard in which it asked the AI what discoveries the James Webb Space Telescope made. Bard responded that it took "the very first pictures" of a planet

outside our solar system. In fact, a European telescope had done that many years earlier. “It’s a good example [of] the need for rigorous testing,” a Google executive said at the time.

Percentage who say the increased use of artificial intelligence in daily life makes them feel...



Percentage who say the increased use of artificial intelligence in daily life makes them feel more concerned than excited...



Note: Respondents who did not give an answer are not shown.
 Source: Pew Research Center surveys conducted July 31-Aug. 6, 2023, Dec 12-18, 2022 and Nov 1-7, 2021

What are AI providers doing to minimize hallucinations?

Among other things, they say they are being more careful about what information is put into the training material. They also are using methods that allow them to better understand what the AI does with the text, and they’re performing more-extensive testing with human users to look for the tendency to hallucinate.

Once an AI system is released, the providers are asking the public for feedback about false answers and similar issues they see.

Do chatbots warn people about the possibility of giving bad information?

Many of them tell users that on their home pages. Google states that “Bard may display inaccurate info, including about people, so double-check its responses.” OpenAI states that “ChatGPT can make mistakes. Consider checking important information,” while Microsoft says “Copilot is powered by AI, so surprises and mistakes are possible.”

***What about people using AI to create phony news reports, photos or videos?
Can it be stopped?***

That’s a disturbing capability of AI. Politicians can be inserted into photos of events they didn’t attend. Videos of TV news anchors can be doctored to have them say things they didn’t say. Some fakes look so real that even AIs can’t tell the difference.

In July, the White House said major AI providers agreed to develop ways to attach a watermark to images and other content their systems create to indicate it was AI-generated. Software also has been created to detect fakes. Tackling the problem from the other direction, Adobe *ADBE* **-0.59%** has devised a way to certify that images are genuine and not altered.

Meanwhile, AI providers are trying to stop the creation of phony media and political disinformation by blocking certain words or phrases in requests. For instance, if you ask Copilot to create an image of a prominent politician in handcuffs, it will refuse. “To curb misuse we have many precautions and moderation policies in place,” says Divya Kumar, general manager of Copilot marketing at Microsoft.

Yet these safeguards are far from foolproof. All of this means you need to be skeptical if you see images and other media that seem off-base.

Percentage by education level who say for each of the following: “Artificial intelligence is doing more to...



Note: Respondents who did not give an answer are not shown.

Wouldn't it be better if we could see the sources AIs use in their responses, like footnotes or citations in a report?

That would be a big help, but there are technical hurdles. Because of the way AI chatbots are trained—to recognize patterns in the material they're fed, not to act as a simple warehouse for facts—and because they use mathematical formulas rather than text searching to generate responses, it is extremely difficult to provide citations to texts used in their training.

I asked Bard a question about large language models. It came back with useful information, so I asked to see its sources. Bard responded that “I apologize for not providing citations in my previous response about LLMs. Here are some resources I used to compile that information,” followed by a long list of links. But these weren't where Bard got the information. A Google spokesman told me that the “best the product team can explain is that the citations in your enquiry are Bard's attempt to be helpful” but that they don't “reflect the specific information used to train Bard.”

Similarly, Bard sometimes supplies links to material from the web along with responses even if you don't ask for them. But again, “these links are not meant to denote direct sourcing for training data,” another Google spokesman said.

Microsoft's Copilot, meantime, combines the traditional web-search function of its Bing with AI-generated text from ChatGPT (Microsoft owns a stake in OpenAI). Its responses also provide web links to certain material. But as with Bard, “these citations only come from—and point to—web content from Bing and not training data,” says Microsoft's Kumar.

What other measures should be considered to help us discern the quality of the information AI produces? Would some type of rating system work?

A universal standard or group to rate chatbots' responses for accuracy, relevance, freedom from hallucinations, etc., is an excellent idea. So far there isn't one, but various academic groups and tech magazines evaluate AIs.

A group at the University of California, Berkeley has created a “Chatbot Arena” to compare responses to the same question from AI systems. PC Magazine put chatbots through tests, with business-oriented Jasper coming in first for 2024, followed by Bard, Bing/Copilot and ChatGPT in that order.

Others say the onus should be on chatbots to publicly disclose their shortcomings by releasing “transparency reports” that describe their faults and limitations.

Executives from Google, Microsoft, OpenAI and other experts warned publicly in 2023 about the dangers of AI. What do they fear?

These experts worry that as AI of all sorts—not just chatbots—becomes more powerful, it could do such things as help terrorists launch devastating cyberattacks or create unstoppable biological agents that set off pandemics. They even fear an AI could go rogue and take control of crucial systems such as military installations, posing an existential threat to humanity. This is based in part on the concern that humans may no longer know exactly how these systems work or what they are doing.

Others say these are far-fetched science-fiction plots that date to the dawn of computers in the 1950s and that humans always will have ultimate control over AIs.

In March 2023, many in the AI industry signed a statement calling for a pause in moving ahead with ever more powerful versions. Did that pause happen?

No. Development of more-advanced AI systems kept going, and perhaps even accelerated.

“AI corporations are recklessly rushing to build more and more powerful systems, with no robust solutions to make them safe,” said Anthony Aguirre, executive director of the Future of Life Institute, in September. The nonprofit group was behind the open letter calling for the pause, which was signed by more than 30,000 researchers, industry leaders and other experts.

Some say governments need to regulate AI to make it as safe as possible. What might that look like?

In May, the presidents of Microsoft and OpenAI supported calls for the U.S. to create an agency to license major AI systems. Congress, meanwhile, discussed possible legislation to set safeguards, while the Biden administration in October issued an executive order that requires developers of the most powerful AI systems to share their safety-test results and other information with the government, among other things.

In November, representatives of more than two dozen countries gathered in England at Bletchley Park, where experts broke the code used by Nazi Germany, and pledged to cooperate on AI safety. “There is potential for serious, even catastrophic, harm, either deliberate or unintentional, stemming from the most significant capabilities of these AI models,” they said in a joint statement.

In December, the European Union reached agreement on AI regulation. Among its provisions is the requirement that AI systems release details about what material is used to train the large language models.

I've heard a lot of doomsday speculation about AI, but what are some of the ways it is doing good?

Artificial intelligence is boosting or transforming the fields of medicine, science, industry, education and daily life, to name just a few areas.

AI can analyze X-rays and MRIs to detect diseases in their early stages. In astronomy it can help discover new planets by analyzing telescope data. AI powers driver-assist technology in cars such as lane-keeping and automatic braking. It helps delivery companies optimize routes and architects to design buildings.

Wall Street uses it to help create asset-management plans, while authorities who monitor that industry use it to detect fraud. Travel websites deploy it for personalized recommendations for flights and hotels. And in education it can provide virtual tutoring and personalized learning.

In short, it's becoming harder to find areas of life that aren't being touched by artificial intelligence.

Bart Ziegler is a former Wall Street Journal editor. He can be reached at reports@wsj.com.

Appeared in the January 22, 2024, print edition as 'Readers Have a Lot of Questions About AI. We Answer Them.'

Next in Journal Reports: Technology

