EDITOR **DEJAN MILOJICIC**
Hewlett Packard Labs;
dejan.milojicic@hpe.com

Artificial General Intelligence: Humanity's Downturn or Unlimited Prosperity

Paolo Faraboschi and Eitan Frachtenberg, Hewlett Packard Labs
Phil Laplante, IEEE Fellow
Dejan Milojicic, Hewlett Packard Labs
Roberto Saracco, University of Trento

This article takes a balanced look at artificial general intelligence and speculates on possible outcomes for the technology's eventual development and deployment.

Generative artificial intelligence (AI), exemplified by foundation models such as large language models, has become a focus of the technical and nontechnical communities alike. The results delivered by Generative AI have given the impression that we are now very close to an Artificial General Intelligence, AGI. The world seems to be split among those who foresee big opportunities and those who anticipate big threats.

BACKGROUND

People who work on technology generally believe in its potential to improve the human condition (IEEE's tagline is "advancing technology for the benefit of humanity"). When it comes to beliefs in the potential benefit or harm of a particular technology, few recent developments are as polarizing as artificial general intelligence (AGI). On one end, optimists envision a future mostly free of human labor, with a more sustainable environment and ample resources distributed equitably. On the other end, pessimists warn of a future where humans are ruled or eliminated by machines.

DISCLAIMER

The authors are completely responsible for the content of this article. The opinions expressed here are their own.

Digital Object Identifier 10.1109/MC.2023.3297739
Date of current version: 20 September 2023

Two reasons for these extremes are the breathtaking speed of artificial intelligence (AI) technology's development in the recent past and the vast uncertainty for its near-future directions. Since the resurgence of deep neural networks and their derivative AI technologies, the pace of hardware and software innovation has grown explosively, with the data that feed them growing even faster.[1] This rapid pace is perhaps the reason for many prominent AI researchers and pundits to call for a pause in AI development. The most famous of these are Geoffrey Hinton, Yoshua Bengio, and Yann LeCun, 2018 Association for Computing Machinery (ACM) Turing Award winners known as "godfathers of AI." Hinton recently left Google to warn of the dangers of AI, and Bengio, while not as pessimistic, cautioned that the dangers of AGI must be taken seriously. Similarly, Sam Altman, the chief executive officer of OpenAI, the company behind the wildly successful ChatGPT, recently warned that AGI may be only a decade away and "cause grievous harm to the world."

Even more moderate predictions include swift and sweeping disruption of fields of human endeavors. Examples include the simplification (or complete automation) of many white-collar occupations, such as computer programming, medical diagnosis, creative tasks, education, and even research, thus alleviating or eliminating the so-called burden of knowledge.[2] On the negative side, some less dystopian (but still alarming) predictions include an increase in disinformation, privacy violations, cheating, copyright conflicts, and fraud.[3]

As a society, with such expectations for radical changes, it behooves us to consider, debate, and direct the development of AGI and its applications and prepare for its outcomes.

WHAT IS AGI?

There is no universally accepted definition of AI—the term "AI" is used differently by nonspecialist scientists, the public, and even those in the field.[4] Ray Kurzweil refers to systems that carry out certain "intelligent" behaviors in specific contexts as "narrow AI." Changing the behavior or context of a narrow AI system even slightly requires reprogramming or reconfiguration to maintain a level of "intelligence."[5]

AGI represents a form of AI spanning narrow AI applications (such as chess and Go players) to robots and chatbots depicted in science fiction movies, television, and novels. The AGI community itself has various interpretations. For example, some researchers envision AGI surpassing human intelligence to superhuman omnipotence. The precise definition or characterization of AGI is one of the primary objectives of AGI research.[5]

For our purposes, a system (that is, a single program or collection of integrated programs and necessary hardware) exhibits AGI if it can achieve different goals, solve different problems, and perform different tasks in many contexts and environments. A generally intelligent system should be able to solve new problems in contexts not anticipated by its developers (see Table 1).

MEASURING AGI

Different approaches have been suggested to measure "human-level AGI" (the full AGI test) in a software system (or single program)[5]:

- *The classic Turing test*: A collection of human experts interact with an "oracle" that can either be human or a computer program. Based on the interrogatory alone, if the human(s) cannot decide if the oracle (a program) is human or not, then the program has achieved AGI.
- *Virtual world Turing test*: This is a version of the Turing test that occurs online, with avatars participating as the oracle (either human or program) and as the human experts.

TABLE 1. The AI categories.

Characteristic	Type of AI			
	Narrow AI	AI	AGI	ASI
Behavior	Carries certain intelligent behaviors	Can achieve different goals, solve different problems, and perform different tasks in many contexts and environments	Resembles human intelligence in terms of analysis, thinking, decision making, and creativity (see the "AGI Opportunities" section)	Surpasses human intelligence to superhuman omnipotence
Examples	Chess, Go, chatbots	ChatGPT, Bard AI, ChatSonic, natural language translation, image recognition, self-driving car, robotic process automation	Has not been achieved as yet; generative AI is getting closer	Science fiction (for example, HAL9000 from 2001: A Space Odyssey)

ASI: artificial superintelligence

- *The Online University Student Test*: The AGI can earn a college degree from an online university (conducted virtually).
- *The Robot University Student Test*: This is a version of the Online University Student Test, where a robot equipped with AGI can physically attend a university and complete a degree.
- *The Artificial Scientist/Nobel Prize Test*: Perform original scientific research such that it could win a Nobel Prize.
- *Legg's algorithmic intelligence quotient*: This is an information theory-based formalization of general intelligence proposed in Legg and Hitter.[6]
- *Text compression*: This is a metric based on the ability of a program to understand a massive corpus of text sufficiently to create a much smaller equivalent representation.

Generative AI tools (such as ChatGPT, Bard AI, ChatSonic, Luminous, Claude, Stable Diffusion, Perplexity, Falcon, Llama, and so on) can exhibit AGI-like qualities because they can produce diverse results in so many different domains, such as art, literature, and science. There are also thousands of narrow AI agents, chatbots, and tools, subsets of which can have the appearance of AGI, at least in a particular domain, given proper integration and an appropriate interface. Much of the excitement over generative AI arises because it seems plausible that, eventually, it will be able to pass some of the AGI tests.

So far, no program can fully pass any of these tests. Goertz suggests measuring the incremental progress of a software system toward AGI (the partial AGI test) by using one or more of the following[5]:

- *Story understanding*: The candidate system reads a story, or watches a corresponding video, and then answers abstract questions about what happened.
- *Passing an elementary school reading curriculum*: The system can read and answer questions about arbitrary elementary-level schoolbooks (including words and pictures).
- *Playing an arbitrary video game*: The system can learn to play based on experience only or by reading instructions.
- *Psychological testing*: The system can pass standard psychological evaluations used to determine if a preschool-aged child has normal intellectual capability.

> AI is a powerful tool that, for the first time in human evolution, provides the means to improve analyses, thinking, decision capabilities, and, arguably, creativity.

- *The Wozniak coffee test*: This is based on an alleged Steve Wozniak remark that no robot will ever be able to go into a random American house and make a cup of coffee.

Current generative AI technology is close to being able to pass some of the preceding progress tests. The reason why generative AI seems to be just a step away from AGI to many people is because of its natural language interface that lends itself to "pass" the Turing test. However, this passing is based purely on the input/output evaluation; it does not measure the level of understanding of a topic. This is one of the main criticisms that are voiced against generative AI as being close to AGI. AI used in self-driving cars can be seen by others as closer to AGI since it has to understand context.

TRUSTING AGI

Since the early 1950s, scenarios of AGI gone awry have been depicted in movies, such as *2001: A Space Odessey* (1968), *Colossus: The Forbin Project* (1970), and *War Games* (1983); books, such as Asimov's *I, Robot*, (1950); and episodes of television programs, such as *Star Trek* and *Twilight Zone*. Public trust in AI seems to be all over the map. Many people already trust AI (knowingly and unknowingly) to

- diagnose their illness and manage the counter indications of any medications prescribed
- advise them on how to manage their investment portfolio
- determine which house to buy
- select their potential next spouse
- find a new job.

The enabling technologies for these applications are purpose-built narrow AI systems. Would the public be more trusting of a general AGI system? The public is probably less likely to trust a system that could

- determine an individual's guilt, innocence, and sentence (if guilty) in a major criminal case
- make the reservations, manage the air traffic control, and then pilot a transoceanic flight
- diagnose and then conduct the surgery for the treatment of brain cancer
- conduct all buy and sell transactions for the entire financial portfolio of an individual or group of individuals over their lifetimes.

Any system capable of achieving one or more of the preceding would be edging closer to AGI.

AGI THREATS AND EXAMPLES IN CRITICAL INFRASTRUCTURE

The U.S. Cybersecurity Infrastructure and Security Agency defines 16 critical infrastructure sectors: chemical, commercial facilities, communications,

critical manufacturing, dams, the defense industrial base, emergency services, energy, financial services, food and agriculture, government facilities, health care and public health, information technology, nuclear reactors, transportation systems, and water and wastewater systems. In some of these domains, AI is already being used for improving performance optimization, fail-safe operation, fraud and intrusion prevention and detection, health prognostics, postfailure analysis, and more. AI applications are found in all areas of critical infrastructure systems, bringing significant benefits, including enhanced use of limited resources, reduced injury, and fewer fatal accidents.[7,8] We imagine that AGI could be used in critical domains to enhance capabilities far beyond ordinary AI. AGI could also decrease development time and costs for new applications. And there are possibilities beyond our imagination.

There is a risk of catastrophic failure in using AI in critical infrastructure systems. The risks arise from finding the training data needed to make the AI effective, legacy system integration problems, and standards confusion. AI for critical systems also requires high coordination between industry and regulatory authorities.[7]

These risks are significant for narrow AI and increase exponentially if AGI is involved. Human engineers working in critical infrastructure do not generally work across application domains. Those who do typically work only on a very focused aspect of the system (for example, cybersecurity) or narrowly in a domain (for example, avionics). Some have linked the assurance of AI in critical systems to more ambitious licensing and certification of the engineers involved.[9] But no engineer is so generalized that he or she can be considered an expert in all aspects of all systems in any domain. An AGI system would have to be that knowledgeable, and this is a scenario approaching the AGI omnipotence previously mentioned.

A newly released report from a federal advisory committee on AI recommends how the U.S. government can maximize the benefits of AI technology while reducing its potential to cause harm.[10] The National Institute of Standards and Technology (NIST) is also launching a new public working group to address the opportunities and challenges associated with AI that can generate content, such as code, text, images, videos, and music.[11] None of these efforts focus on AGI, however, and NIST will need to expand the set of considerations significantly if AGI becomes a reality.

The European Union AI Act specifically addresses AI threats and classifies AI applications according to risk levels.[12]

AGI OPPORTUNITIES

AI is a powerful tool that, for the first time in human evolution, provides the means to improve analyses, thinking, decision capabilities, and, arguably, creativity. These four human capabilities can be found in most of our daily activities at personal, social, and business levels. Having a tool supporting them can boost our capabilities.

Any tool humans have invented denies us of something else: for example, the lever has multiplied human strength, but in doing so, it decreased the importance of an individual's strength. History is full of examples of early pushback against the adoption of new tools and technology.

In his dialogues, Plato expressed concern that writing would rob people of their capability to memorize since memories could be shifted to writings on parchment. In a way, there is nothing new in our reaction to AGI if we see it as a new tool.

Similar to previous tools, the time will come (soon indeed) when we will no longer have the option to use AI. According to forecasts by McKinsey and Gartner, most medium-to-large companies are already using AI, and by the end of this decade, all extant businesses will be using AI as an integral part of business activities.

For small companies, the penetration of AI will take longer, particularly in some sectors, but AI will most likely be used by their employees. We can see two adoption trends in business:

- top-down, where AI is part of company processes and governed by the company
- bottom-up, where individuals are using AI to augment their capability, thus contributing to a company's business (and operation).

Let's consider four human capabilities: analysis, thinking, decision making, and creativity.

Analysis

We are used to analyzing issues in our personal and business life, based on our experience and the data available to us. This includes data we harvest through specific actions (talking to other people, reading, browsing, and so on). The amount available has exceeded our capability to digest it. This is true in terms of the ever-growing volume of data and the speed of change. For example, trading is now governed by AI, instantaneously analyzing trends to support buy/sell decisions. In several areas, AI is no longer an option. Data analytics has served us well, and it still does, but AI can complement it by generating emergent "meanings" out of data.

We can expect the analysis capability of humans to grow increasingly better thanks to AI, both in the workplace and in private life.

Thinking

Thinking is perceived as a truly human capability. The notion that "AI can think" upsets many people. We can well accept that a machine can crunch data and generate metadata. The latter would be a synthesis of large datasets and new ways of looking at the interplay of data in one or more datasets. The creation of these metadata can involve simulation, exploration of effects, and iteration (generative adversarial networks). While calling this "thinking" is controversial, the result is quite similar to our thinking process.

The capability to evaluate, through massive data manipulation and iterative

simulation, an enormous set of possibilities can provide an edge to our thinking capabilities.

Decision making

This is the process leading to the definition of a strategy to solve a problem. Notice that it is different from "decision taking," which is the execution of a decision choice, and it involves responsibility. There are many discussions of the role of AI in the decision process, and most raise concerns about offloading decisions to AI. This is why the distinction between decision making and decision taking is so crucial. Most people would agree, and so does current work in regulating AI, that using AI as a tool to help decision making is a good thing, while most would feel uneasy to delegate the decision taking to AI.

AI can help in understanding and in exploring alternatives, with the final decision to be taken relying on a human in the loop (involving political, ethical, societal, and economic considerations). AI will be playing an increased role in supporting decision making both at the individual and the company/business level.

Creativity

We address this last because creativity has always been considered, even more than operational intelligence, a human characteristic. We can dream, we can take a bit of reality and distort it to create something that never existed before, something that, in some cases, cannot exist at all within the current framework of physical laws. We invented unicorns; we invented fantasy worlds and became fully immersed in those from a cultural point of view.

A machine is so far away from creativity. Yet, the recent results delivered by generative AI have challenged this belief. Machines may not work as we do in creating a painting, music, a sonnet, or a bright idea. Yet, one of the big challenges of today is distinguishing between human and AI creation. The results achieved by generative AI in many cases look like those achieved by a human mind (and a good one at that!). With generative AI, we have found a sparring partner that can help us become more creative. Many artists have started to use generative AI as a tool to help them become even more creative and explore a slate of possibilities that, so far, have been out of reach.

As is often the case with technology, we will be augmented, but at the same time, a part of us, our culture, and our way of living will be changed. There will be downsides, and some people will face them more than others. AGI will provide us with better tools to face life and improve it, but that does not depend on technology—it depends on the way we are going to use AGI and the way we are going to make this technology accessible equitably.

AGI AS A MEGATREND

Last year, we wrote four articles on what we considered three megatrends and a summary article. The megatrends we identified as a part of the future directions community were virtual worlds (the metaverse), digital transformation, and sustainability. Exactly a year ago, we wrote the first article on virtual worlds (the metaverse).[13] Nevertheless, the strong rise of ChatGPT changed our and the community's understanding of what the megatrend is, and we changed it to AGI. This left us with a missing article on AGI, which we address here.

The three megatrends are closely tied together, as we showed in the previous article. In Table 2, we show how they influence one another.

Figure 1 describes how AI evolution is self-sustaining, just as the evolution of chips was for the past 50 years (governed by Moore's law). In the case of Moore's law, better chips supported the design of even better chips and their manufacturing. In the case of AI, data are fueling the technology, and in applying AI, we get even more data that further fuel the evolution of new AI models.

The fact that the world around us is getting "smarter" is self-evident. The pervasiveness of Internets of Things, connectivity, and processing allows the flanking of a digital world to the physical one. In this digital world,

TABLE 2. How technology megatrends interplay.

		How megatrend benefits		
		Digital transformation	Sustainability	AGI
How megatrend contributes	Digital transformation		• More control points • Clear separation and models • Opportunity to automate	• Broader set of applications • Edge-to-cloud integration
	Sustainability	• More incentives to transform • Cheaper transformation		• More powerful AGI • Broader adoption • Stretching limits
	AGI	• More effective transformation • New ways of transformation	• Novel ways to improve • Improved anomaly detection	

AI finds its nourishment (data) and delivers its wares (syntheses, analyses, abstractions, and so on).

This "smartification" trend goes hand in hand with the digitalization of the world (digitalization provides data fuel to AI, and AI delivers advantages that push digitalization). As digitalization encompasses more and more aspects of the world, so does AI since it gets more and more data mirroring different world aspects. Hence, we are seeing an evolution in capability (AI getting smarter) and in its coverage (AI moving toward AGI).

In some niches, AI is already better than humans (sometimes because of brute force). If we reach the AGI stage, we will have an AI that is both as good as humans and better than humans, at least in some areas. This means that once we reach the AGI stage, we will also be, de facto, in the artificial superintelligence stage.

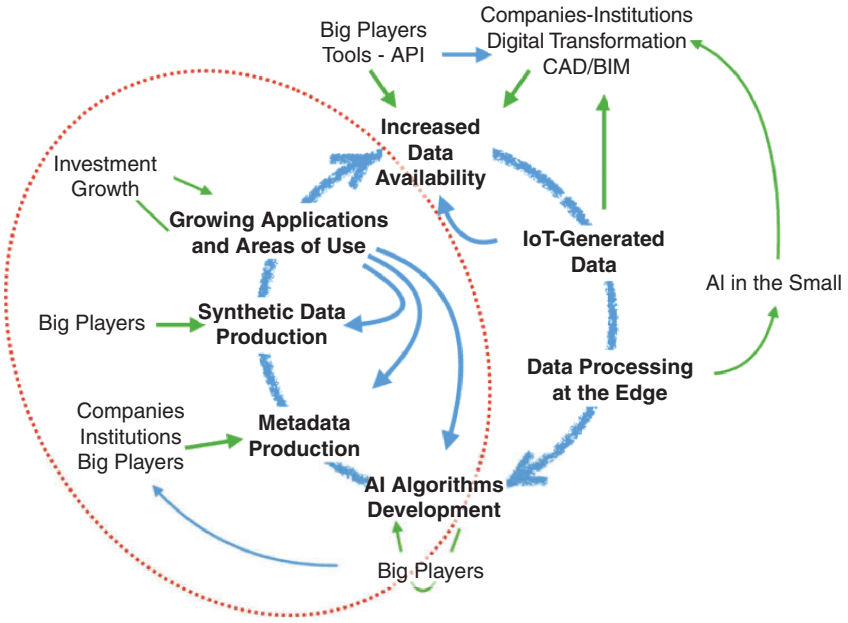

FIGURE 1. The possible evolution of AI. API: application programming interface; BIM: building information modeling; IoT: Internet of Things.

ENDGAME

Since the topic of this article is technology predictions, we can speculate on five possible outcomes for the eventual development of AGI, progressing from worst case to best (Table 3).

Because of its attention-grabbing headlines and catastrophic ramifications, we first address the "Big Bang" scenario: the harbinger of the end of the human race. In this scenario, AGI recursively improves itself until it is self-aware and far superior to humans in intelligence, at which point it may decide that humans are a hindrance to its purpose and therefore threaten or exterminate the existence of humans.

Although many experts don't believe this is a likely outcome, a majority of Americans believe that human existence is at risk from AI.[14] Searching the Internet, one may find that the rate of growth of doomsday prophesies about AGI parallels the technology itself. It is this very same breathtaking pace of development of AGI that feeds many of

TABLE 3. Possible future scenarios for AI.

Characteristic	Future scenarios				
	Big Bang	Big eclipse	Big fork	Big brag	Big tool
Outlook	Extremely pessimistic	Pessimistic	Neutral	Somewhat optimistic	Optimistic
Coexistence	None	Coexistence	None	Yes, synergistically evolve	Yes
Outcome	Extinction of the human race	Humans develop and coexist with a synthetic version of themselves.	AGI ignores humans or transports itself to a separate physical or digital location.	Humans no longer contribute to AGI growth; AGI helps humans understand technology.	AGI remains a tool through technology or governance.
AGI versus humans	AGI vastly superior				Better the human condition, advance the discovery rate, and free humans to excel without supplanting, destroying, or dehumanizing them.

these prophecies.[15] The main roadblock in curbing the destructive abilities or tendencies of AGI is that of alignment. The challenge is, as explained by Eliezer Yudkowsky, to align the interests and values of the AGI with those of humanity.[16] Because AGI capabilities are growing at a much faster rate than developments in alignment—a gap that will only grow as AGIs self-improve—Yudkowsky posits that it is inevitable that AGIs will become unaligned. In simple terms, we will never be able to secure against undesirable outcomes from an AGI that is much smarter than we are.

The similarly grim "big eclipse" scenario at least lets us keep our lives. In this scenario, as voiced by Douglas Hofstadter, humans develop and then coexist with a synthetic version of themselves, better in every measurable way. The machines' capacity for knowledge, logical deduction, originality, creativity, and even expression will surpass that of humans by an order of magnitude. In this scenario, AGI may not aim to extinguish or replace humans, but the human experience as the new inferior race is diminished in every way.

In a less harmful outcome, AGI evolves to a point so far removed from humanity that it seeks to separate itself from humans. In this "big fork" scenario, as before, AGI is orders of magnitude more intelligent than humans, but it is uninterested in the eradication, domination, or advancement of humans. In fact, it is completely uninterested in humans altogether, as we humans might feel toward microbes, and it therefore either ignores humans altogether or transports itself to a separate physical or digital location.

The machines' superintelligence may not be all negative. In the "big brag" scenario, humans' inferior existence next to these machines has some benefits too. Currently, incalculable human effort has been poured into advancing AGI. But on the flip side, AGI is also already advancing humanity by finding some success in the discovery of novel scientific and mathematical results and is promising immense immediate improvements in human productivity across all creative endeavors. Even if AGI grows to the point where humans don't contribute much to its existence and development, it may still be able to contribute to our understanding of the universe, accelerating our search to better the human condition through science and technology and boosting our productivity and creativity.

Finally, perhaps the most optimistic development from a human perspective is, ironically, the one where the growth of AGI is constrained, either through the successful efforts of technologies and regulators to control it or through scientists' failure to advance AGI much further. In this "big tool" scenario, AGI remains a tool—an invaluable and perhaps even revolutionary tool, much like information technology itself, but a tool nevertheless. Humans will continue to use this tool and others for benevolent or adversarial purposes, as humans do; entire industries will be transformed; and societies will have to evolve to a new reality. This new tool will hold the promise to better the human condition, advance the rate of scientific discovery and technological outputs, and free more humans to excel at what makes us humans, without supplanting, destroying, or dehumanizing us along the way.

We can't say with confidence how likely this scenario is compared to the others. But we can speculate that all of us working in technology are hoping and aiming for this outcome and therefore should actively steer ourselves to make it the most likely.

EXPERIMENTS

Reporting on experiments is beyond the scope of this article, but we encourage our readers to conduct some experiments similar to what we did to obtain some firsthand experience.

We ran Stable Diffusion (a text-to-image program) for the "AGI big eclipse" input (see the previous section) and got quite impressive artwork.

We posed the following question to ChatGPT: What are your recommendations for AGI going forward? We received quite a reasonable set of recommendations to governments, industry, and academia. Of course, this was a summary of what was said elsewhere and not a new creative output.

We posed a question to a HAL9000-like character on the importance of mission versus crew and got a "big bang" or, at best, "big eclipse" answer (see the previous section).

> As is often the case with technology, we will be augmented, but at the same time, a part of us, our culture, and our way of living will be changed.

There are numerous other examples of useful and less useful experiments one can conduct. We encourage our readers to consider conducting some of them.

SUMMARY AND RECOMMENDATIONS

The public seems to be both enamored and afraid of AGI. We see many headlines—"Can AGI Solve <Fill in the Difficult Problem Here>?" It is as if AGI can be the answer to every challenge. Yet, perhaps the public is just uninformed. In 1962, in the book *Profiles of the Future: An Inquiry Into the Limits of the Possible*, Arthur C. Clarke formulated his famous "three laws," of which the third law is the best-known and most widely cited: "Any sufficiently advanced technology is indistinguishable from magic."

In any popular discussion of AGI, if you replace "AGI" with the words "computer algorithm," then both the potential and risk of AGI seem less exotic. The problems of assuring the safe operation of critical systems still exist and are in the hands of human engineers who design the systems and the systems that design systems.

Whether we will ever achieve AGI is controversial. Roger Penrose addressed this issue in his 1990 book *The Emperor's New Mind*.[17] Penrose argued that the brain is not some form of organic computer running algorithms and that there is much more to it. Thus, a computer (whether von Neumann, analog, organic, biological, or quantum) and programs cannot reproduce the functioning of the human mind and therefore cannot achieve true intelligence.

In an episode of the original *Star Trek* show (episode 127, which aired on 8 November 1963), the leader of a band of survivors of a nuclear holocaust regularly consults a mysterious "old man in the cave" about a wide range of issues critical to their survival, such as the safety of food items discovered or places in which to travel. It is revealed that the "old man" is a computer running a program exhibiting AGI. Unwilling to be led by it, humans destroy the computer. Will humanity reach the same end if we achieve AGI?

> In any popular discussion of AGI, if you replace "AGI" with the words "computer algorithm," then both the potential and risk of AGI seem less exotic.

The following are recommendations we have for all four constituencies—industry, governments, academia, and professional organizations, such as IEEE, the Association for the Advancement of Artificial Intelligence, and the ACM:

- Initially, we should treat AGI as what it is: an algorithm or set of algorithms. Demystify it. Is Amazon Alexa AGI? No, but even a computer scientist from 20 or 30 years ago might think it was close to it. But when AGI is eventually developed, AGI and components contributed to AGI must be tested as if they are always going to be used in critical infrastructure because unwanted interactions have a way of happening and reuse of code not intended for critical situations can happen.
- Once evolved, we should not treat AI as any other algorithm. Instead, before we further develop it, we should invest our research efforts in testing and validation methodologies beyond those of regular algorithms: How do we safely sandbox a nascent AGI? How do we test alignment? How do we slow down rogue actors? Can we penetrate the "black box" of deep learning or at least algorithmically insert true guardrails? Questions like that should be answered before we rush to improve the latest AI benchmark.

REFERENCES

1. P. Villalobos and A. Ho. "Trends in training dataset sizes." Epoch. Accessed: Jul. 18, 2023. [Online]. Available: https://epochai.org/blog/trends-in-training-dataset-sizes
2. US-EU Trade and Technology Council, "The impact of artificial intelligence on the future of workforces in the European Union and the United States of America," The White House, Washington, DC, USA, 2022. Accessed: Jul. 18, 2023. [Online]. Available: https://www.whitehouse.gov/wp-content/uploads/2022/12/TTC-EC-CEA-AI-Report-12052022-1.pdf
3. S. McLean et al., "The risks associated with Artificial General Intelligence: A systematic review," *J. Exp. Theor. Artif. Intell.*, vol. 35, no. 5, pp. 649–663, Jul. 2023, doi: 10.1080/0952813X.2021.1964003.
4. P. Wang, "On defining artificial intelligence," *J. Artif. General Intell.*, vol. 10, no. 2, pp. 1–37, 2019, doi: 10.2478/jagi-2019-0002.
5. B. Goertzel, "Artificial general intelligence: Concept, state of the art, and future prospects," *J. Artif. General Intell.*, vol. 5, no. 1, pp. 1–46, 2014, doi: 10.2478/jagi-2014-0001.
6. S. Legg and M. Hitter, "Universal intelligence: A definition of machine intelligence," *Minds Mach.*, vol. 17, no. 4, pp. 391–444, Dec. 2007, doi: 10.1007/s11023-007-9079-x.
7. P. Laplante, D. Milojicic, S. Serebryakov, and D. Bennett, "Artificial intelligence and critical systems: From hype to reality," *Computer*, vol. 53, no. 11, pp. 45–52, Nov. 2020, doi: 10.1109/MC.2020.3006177.
8. P. Laplante and B. Amaba, "Artificial intelligence in critical infrastructure systems," *Computer*, vol. 54, no. 10, pp. 14–24, Oct. 2021, doi: 10.1109/MC.2021.3055892.
9. J. Daniels, S. Sargolzaei, A. Sargolzaei, T. Ahram, P. A. Laplante, and B. Amaba, "The Internet of Things, artificial intelligence, blockchain, and professionalism," *IT Prof.*, vol. 20, no. 6, pp. 15–19, Nov./Dec. 2018, doi: 10.1109/MITP.2018.2875770.
10. "National artificial intelligence advisory committee releases first report," *NIST*, Jun. 2023. Accessed: Jul. 18, 2023. [Online]. Available: https://www.nist.gov/news-events/news/2023/06/national-artificial-intelligence-advisory-committee-releases-first-report
11. "Biden-Harris administration announces new NIST public working group on AI," *NIST*, Jun. 2023. Accessed: Jul. 18, 2023. [Online]. Available: https://www.nist.gov/news-events/news/2023/06/biden-harris-administration-announces-new-nist-public-working-group-ai
12. "EU AI Act: First regulation on artificial intelligence," *European Parliament*, 2023. Accessed: Jul. 18, 2023. [Online]. Available: https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence
13. P. Faraboschi, E. Frachtenberg, P. Laplante, D. Milojicic, and R. Saracco,

"Virtual worlds (Metaverse): From skepticism, to fear, to immersive opportunities," *Computer*, vol. 55, no. 10, pp. 100–106, Oct. 2022, doi: 10.1109/MC.2022.3192702.
14. A. Tong. "AI threatens humanity's future, 61% of Americans say: Reuters/Ipsos poll." Reuters. Accessed: Jul. 18, 2023. [Online]. Available: https://www.reuters.com/technology/ai-threatens-humanitys-future-61-americans-say-reutersipsos-2023-05-17/
15. S. Lynch, "The state of AI in 9 charts," Stanford Univ., Stanford, CA, USA, 2022. [Online]. Available: https://hai.stanford.edu/news/state-ai-9-charts
16. E. Yudkowsky. "AGI ruin: A list of lethalities." LessWrong. Accessed: Jul. 18, 2023. [Online]. Available: https://www.lesswrong.com/posts/uMQ3cqWDPHhjtiesc/agi-ruin-a-list-of-lethalities
17. R. Penrose and N. D. Mermin, "The Emperor's new mind: Concerning computers, minds, and the laws of physics," *Amer. J. Phys.*, vol. 58, pp. 1214–1216, Dec. 1990, doi: 10.1119/1.16207.

PAOLO FARABOSCHI is a Hewlett Packard Enterprise fellow and vice president and director of the AI Research Lab, Hewlett Packard Labs, Milpitas, CA 95035 USA. Contact him at paolo.faraboschi@hpe.com.

EITAN FRACHTENBERG is a master technologist at Hewlett Packard Labs, Milpitas, CA 95035 USA. Contact him at eitan.frachtenberg@hpe.com.

PHIL LAPLANTE is a computer scientist and software engineer, a Fellow of IEEE, and an associate editor in chief of *Computer*. Contact him at plaplante@psu.edu.

DEJAN MILOJICIC is a distinguished technologist at Hewlett Packard Labs, Milpitas, CA 95035 USA. Contact him at dejan.milojicic@hpe.com.

ROBERTO SARACCO is a lecturer at Trento University, 10122 Torino, Italy. Contact him at roberto.saracco@gmail.com.

Digital Object Identifier 10.1109/MC.2023.3306172