



The Profession of IT

Three AI Futures

Artificial intelligence development practices may be facilitating a drift toward an unwanted AI singularity.

THE PRESENT ARRAY of artificial intelligence (AI) and social paradoxes could be described by a future historian in the following way: “It was the best of times, it was the worst of times, it was the age of wisdom, it was the age of foolishness, it was the epoch of belief, it was the epoch of incredulity, it was the season of Light, it was the season of Darkness, it was the spring of hope, it was the winter of despair, we had everything before us, we had nothing before us, we were all going direct to Heaven, we were all going direct the other way.” Do these words sound familiar? Charles Dickens opened his novel *A Tale of Two Cities* in 1859 with them.

The turmoil and promises of AI have generated a confusion of uncertainty about the future. Can we trust AI? Will AI take our jobs? Is an AI app safe? Shall we invest in AI company? Create an AI startup? Will AI accelerate cyber-crime? Undermine our ability to learn the truth? Endanger humanity? Will organized crime or totalitarian autocrats use AI to take over? Should AI be regulated? In short, what will the future bring?

Three kinds of AI futures have dominated the speculations: a singularity when humans merge with superintelligent machines, a utopia from pervasive automation, and a huge array of autonomous AI agents providing useful services. We will refer to them as:

- ▶ The Merger
- ▶ The Utopia
- ▶ Agentic AI

This column is a reflection on these futures. While the Merger and the Uto-



pia are unlikely, an Agentic automation singularity is a disturbing threat.

The Merger

In his book, *The Singularity is Near* (2006),⁶ Ray Kurzweil introduced the singularity: an event horizon in AI, a time when superintelligent machines emerge and replace humanity. The centerpiece of his argument is that information technology has doubled its performance every two years for well over a century. In effect Moore’s law has held since the tabulating machines circa 1900 and will continue to hold in post-silicon technologies. Extrapolating this trend, he predicts that computing will pass the Turing Test by 2029, enable repair and communication nanorobots circulating in the human bloodstream by late 2030s, and attain the singularity

by 2045. This vision is at once strikingly compelling and deeply disturbing. It is easy to imagine a superintelligence that exterminates humans, which it sees as fraught, fallible, and feeble.

In his newer book, *The Singularity is Nearer* (2024),⁵ Kurzweil modifies his idea to make it more appealing: the singularity will not be a takeover by superintelligent machines, but a complete merger of humans and machines. Instead of exterminating or replacing humans, the singularity will celebrate a new species of superhumans.

These ideas are controversial. Kurzweil believes that by 2029 machines will close in on AGI (artificial general intelligence) by passing a very stringent version of the Turing Test, in which the machine excels in all domains of human cognition. This prediction relies

on three questionable assumptions. One is that large language models (LLMs) can generalize to AGI, despite being inherently given to hallucinations, fabrications, and other untrustworthy responses.^a Second is that, when an exponential growth of a technology saturates its domain, another technology awaits in the wings to continue the growth. The third is more serious: the implicit assumption that intelligence is computable. There are domains of knowledge associated with intelligence that we do not know how to describe, measure, and represent in a way that a computer can process them. These include performance skills and what we call the context, which includes “common sense” and “background of obviousness”.

Moore’s Law is not the only basis for predicting AGI. AI machines can be organized into hierarchies of progressively more learning power.² The most powerful today combine generative AI and reinforcement learning. Unfortunately, there is little evidence that this is anywhere close to AGI. Thus, machine hierarchies give no useful insight about progress toward the singularity.

The Utopia

A different sort of future has been proposed by S.M. Sohn in his book *The Last AI of Humanity*.⁸ His hierarchy is four levels of ever-expanding AI automation, culminating with AI-run government (see accompanying table). He argues that these levels of automation will produce food, energy, goods, and essential services in abundance and at near-zero cost. In this Utopia, there will be no shortages, no poverty, no inequality.⁴ This Utopia can be achieved by AI automation without AGI. A hallmark of this world will be the “0-person” organization or government.^b Social safety nets and universal basic income will ward

off social unrest from lack of jobs. He is not worried about the possibility that a criminal or government cabal could take over the AI machines and use them to oppress the rest of humanity, or that free money for everyone might produce massive inflation. At first glance, this model might seem preposterous. But it is a very plausible path to a different singularity—human subjugation by uncaring machines. It is surprising how far the processes of automation have already penetrated at each of Sohn’s four levels.

At Sohn level 1, automation of business apps has been under way for many years. Early examples in the 1990s are self-learning spam filters and auto-completion of typed input. In that era, Microsoft introduced Clippy, an intelligent office assistant. Because so many users found it intrusive, annoying, and often wrong in its recommendations, Microsoft discontinued it. Microsoft has since offered Copilot, a newer assistant based on LLM technology. Copilot can generate cogent summaries of email threads and meetings, answer emails, edit draft documents, and provide tutoring. It remains to be seen whether users find these tools useful and trustworthy.

At Sohn level 2, software-powered business processes such as customer record-keeping and order tracking are quite old; computing systems enabled massive expansions of business into worldwide operations. Since the 1990s, automation of workflows within the organization has become increasingly common—for example, hiring, travel, purchasing, or anything based on forms. Workflow automation has allowed these processes to be managed with fewer personnel. In growing numbers of workplaces today, however, workflow automation

has been coupled with constant surveillance of workers to monitor their progress toward assigned productivity goals. Workflow automation also encodes complex business rules into the machines; users who fail to comply cannot get the services. Some workflow systems confront employees with complex and confusing arrays of tools that do not interoperate well. Workflow systems often do not make work go away. They rearrange work, with less done by the service office and more by the user. Productivity of service offices goes up, and of users down. Workflow systems do not increase overall productivity as much as many believe.

At Sohn level 3, automation at the organizational level is growing rapidly. The most notable form is robots at the user interface for scores of companies and government agencies. The robots only allow limited transactions, with no exceptions or means to request exceptions. Customer services are generally a shambles. Many companies make it hard to find how to contact customer service. Some (such as Facebook) offer no customer service at all. Most provide only an email form (ticket) or a chatbot that is designed to avoid connecting with a human agent. When one finally gets the robot to connect with a human agent, the agent follows a script and does not resolve the problem. In short, automated customer service is universally reviled because it is mindless, inflexible, and uncaring. The situation is unlikely to improve because many companies are seeking to minimize or eliminate their expensive call centers.

As automation turns more to cloud services, data breaches are increasing in frequency and scope. If your identity is stolen, you cannot function until it is restored, which can take a long time. Massive troves of personal data collected by companies are released in these breaches. Ransomware continues its rampant rise. Recovery is often slow and expensive because systems and databases need to be rebuilt and, thanks to automation, there are fewer IT specialists to do that work.

At Sohn level 4, politics, polarization, and elections are under severe stress because AI-generated misinformation and uncorrected lies are undermining

Table. Sohn’s AI automation hierarchy.	
Level	Category of Machines “in charge of”
1	Human business roles (AI copilot, AI assistant)
2	Machine business roles (AI agent, AI butler)
3	Business (AI CEO, AI company)
4	Government (AI president, AI bureaucracy, AI congress)

a A growing segment of AI research literature flows from Harry Frankfurt’s now-famous 1986 essay *On Bullshit*. He defines bullshit as conversation aiming to convince irrespective of facts. LLMs fabricate (build) convincing statistical reconstructions of training data. They have no ability to identify truth in their outputs.

b An old joke, from at least 1978, says: “The factory of the future will have only two employees, a man and a dog. The man will be there to feed the dog. The dog will be there to keep the man from touching the equipment.”

trust at all levels. Hidden behind web interfaces, the “faceless bureaucracy” is no longer a joke. It seems like no one is concerned about customer satisfaction or takes responsibility for it.

The spirit of automation is well established in the legislature, which views law making as the specification of algorithms to run social programs. Most bills have simple titles but go on for thousands of pages of details spelling out how the implementing agency should respond to every conceivable contingency. Some governments use AI tools to maintain “social credits” and subjugate when it serves their interests for power, control, and stability. As AI automation spreads, AIs will be given control over large networks (for example, organizations), where, lacking human oversight, they will set their own goals, seek to control and optimize population segments, and cut off people who do not comply.

Ever-deeper automation is already generating social unrest and social polarization. It is dumbing down humans by loss of critical thinking and inability to compete.

Agentic AI

The business world is not much interested in possible existential singularities far in the future. It is focused on pragmatic AI—AI apps that do jobs better than humans, aiming to relieve humans of drudgework and enable them to spend more of their time on tasks that machines cannot do well. This focus is often called “Agentic AI,” meaning autonomous AI agents interacting with each other and with humans. According to Jensen Huang, CEO of Nvidia, this AI is like a time machine. It can make a future that would take a lot of work happen in a few seconds.

But there is more to value than getting jobs done faster. In *Rebooting AI*, Gary Marcus and Ernst Davis argue that many apps intended to support this direction have failed and will continue to fail because they are not safe and reliable.⁷ They argue the path to AGI depends on employing a lot of software engineering to attain trustworthy AI. This path may be more difficult than we anticipate. While the trust aspects of meeting specifications and reliability can be achieved through good en-

The business world is not much interested in possible existential singularities far in the future.

gineering, the trust aspect of care cannot. Machines do not and cannot care about anything.^{1,3} Expecting machines to achieve high-trust relationships may be impossible.

Today’s practices of implementing AI systems have inspired a long list of near-term concerns. These practices increase the risk of an AI automation singularity:

- ▶ A strong tendency for hype and anthropomorphism. This leads to overclaiming and overpromising, risking a bubble-bursting backlash bringing another AI winter.

- ▶ A definition of productivity that prioritizes speed of task completion over amplification of human capabilities—AI replacing human work rather than augmenting it.

- ▶ Fights over intellectual property taken from the Internet and used to train AI models, without regard for the owners’ copyrights.

- ▶ Surveillance of workers to enforce assigned productivity goals.

- ▶ Lack of business-supported training to prepare workers for coming technology transitions and help the displaced find new employment.

- ▶ Loading sensitive company and government data to the Cloud, where it is more exposed to data breaches and incorporation into the training of AI models.

- ▶ Implicit uploading of user behavior data that is sold to advertisers.

- ▶ Retention by LLMs of user prompts that may contain sensitive data.

- ▶ Biases in training data of models. Tech companies have found that AI screeners favor male job applicants. The open-sourced Wikipedia admits that most of its active content editors are young people with college educations that lean to the left in political articles.

- ▶ Low-wage workers with minimal domain knowledge are hired to label data for training models. One ballyhooed example featured such workers identifying polyps in colon images to train AI colon-cancer screeners.

- ▶ Use of synthetic data to overcome data shortages for training ever-larger LLMs. A substantial amount of LLM generated output resides in the Internet—as summaries, email, advertisements, reports, and letters. LLMs are training themselves rather than learning from humans.

- ▶ Synthetic data from unvalidated digital twins and other simulations may be of low quality.

- ▶ Misinformation and disinformation are easy to generate and fast to propagate in the Internet. They facilitate polarization of citizenry. They influence political decisions and elections through manipulation of people’s perceptions and moods.

- ▶ AI is alluring young people to spend more time at screens, endangering their mental health and social development.

- ▶ AI capabilities such as surveillance and social monitoring can be misappropriated with great success by authoritarian governments and criminal organizations.

These emerging practices are amplifying the contradictions in the growing world of AI and are facilitating the drift toward the AI automation singularity, in which networks of uncaring, unintelligent machines subjugate human beings. Charles Dickens would not be surprised. □

References

1. Can we teach machines a code of ethics? *Forbes* (Mar. 27, 2019); <https://bit.ly/3IdRGLS>
2. Denning, P. and Lewis, T. An AI learning hierarchy. *Comm. ACM* 67, 12 (Dec. 2024).
3. Gourgey, B. Will computers ever feel responsible? *MIT Technology Rev.* (Aug. 28, 2024); <https://bit.ly/44i4zxe>
4. Khosla, V. A roadmap to AI utopia. *Time Magazine* (2024); <https://bit.ly/4eDn6HI>
5. Kurzweil, R. *The Singularity Is Nearer*. Viking (2024).
6. Kurzweil, R. *The Singularity Is Near*. Penguin (2006).
7. Marcus, G. and Davis, E. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Vintage (2019).
8. Sohn, S.M. *The Last AI of Humanity*. SM Research Institute (2024).

Peter J. Denning (pjd@nps.edu) is Distinguished Professor of Computer Science at the Naval Postgraduate School in Monterey, CA, USA, is Editor of *ACM Ubiquity*, and is a past president of ACM. His most recent book is *Navigating a Restless Sea: Mobilizing Innovation in Your Community* (with Todd Lyons, Waterside Productions, 2024). The author’s views expressed here are not necessarily those of his employer or the U.S. federal government.

© 2025 Copyright held by the owner/author(s).