MIT Technology Review
Insights

In association with

intel

Artificial intelligence is a primary driver of possibilities and promise as the Fourth Industrial Revolution unfolds.

Getting smart about the future of AI

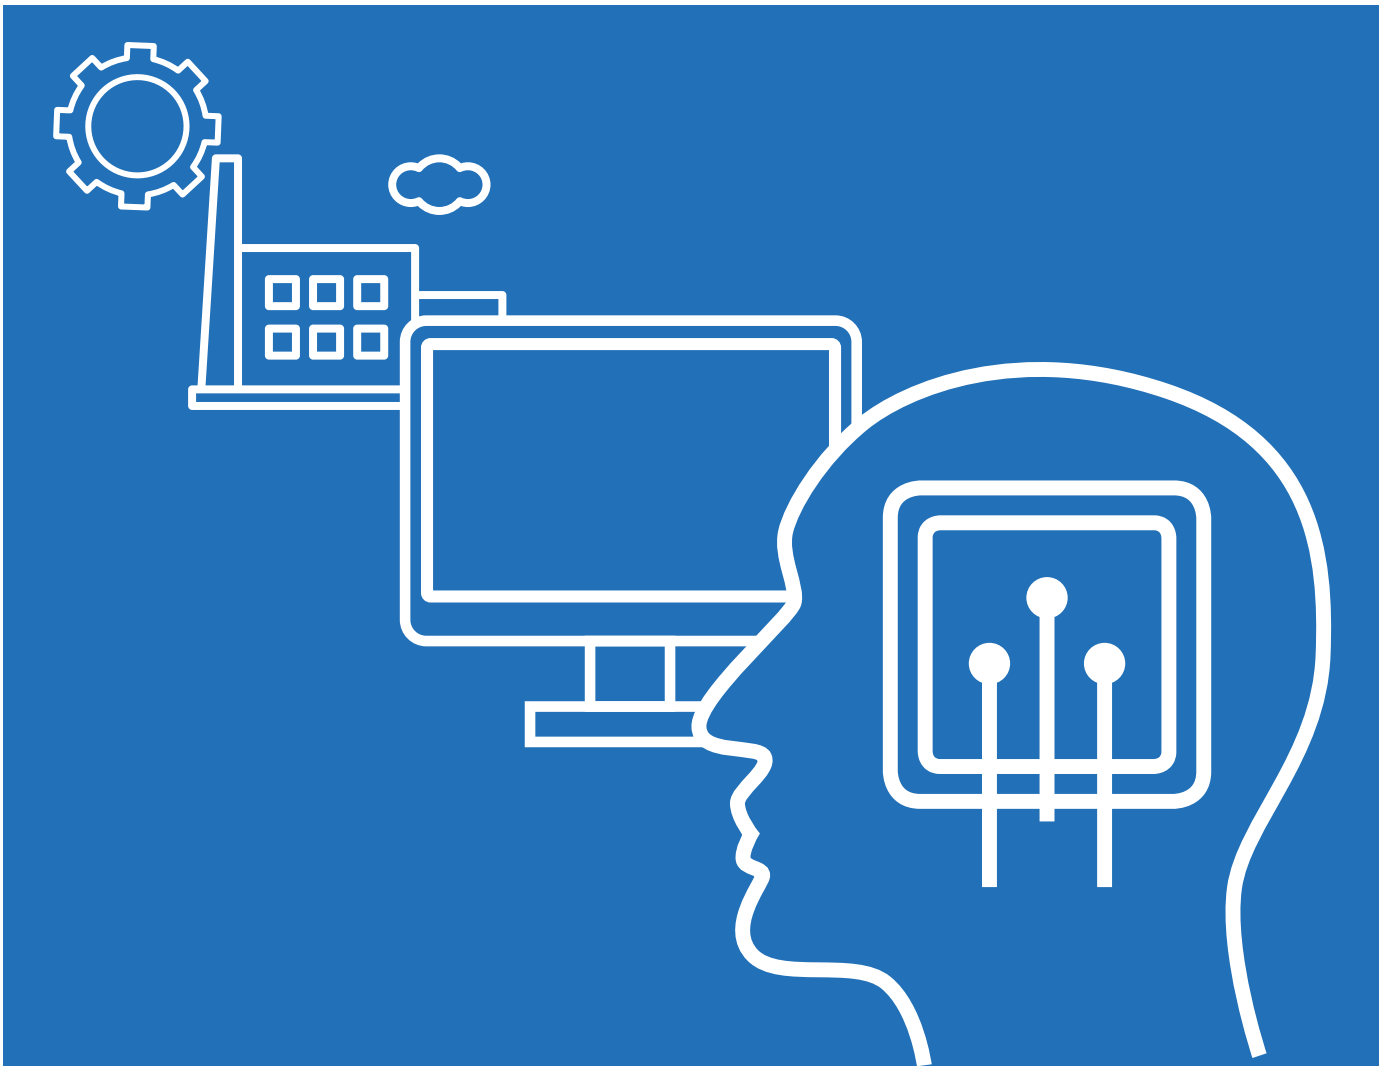

The Industrial Revolution conjures up images of steam engines, textile mills, and iron workers. This was a defining period during the late 18th and early 19th centuries, as society shifted from primarily agrarian to factory-based work. A second phase of rapid industrialization occurred just before World War I, driven by growth in steel and oil production, and the emergence of electricity.

Fast-forward to the 1980s, when digital electronics started having a deep impact on society—the dawning Digital Revolution. Building on that era is what's called the Fourth Industrial Revolution. Like its predecessors, it is centered on technological advancements—this time it's artificial intelligence (AI), autonomous machines, and the internet of things—but now the focus is on how technology will affect society and humanity's ability to communicate and remain connected.

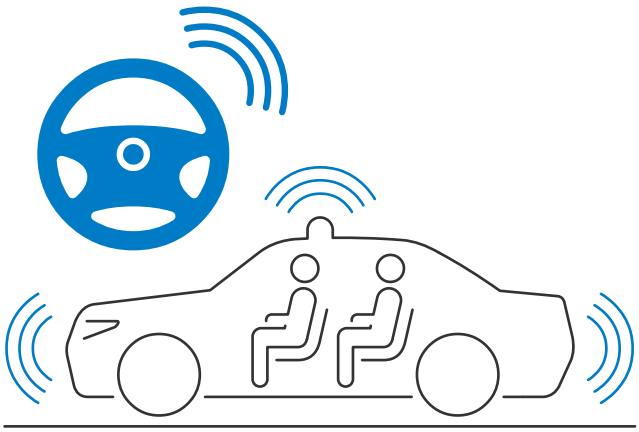

"What does it mean to have systems that will be autonomous? How will we feel about that? What will it mean?"

Genevieve Bell, Distinguished Professor, Australian National University

Key takeaways

1. Artificial intelligence (AI) is a focal point of the Fourth Industrial Revolution, the latest era of technological advancement.
2. As AI becomes widespread, it is essential to consider the potential effects of the technology on society.
3. Challenges include establishing guidelines on the use of AI, determining how to ensure AI systems are safe and ethical, measuring their effectiveness, and devising ways to interact with them.

"In the first Industrial Revolution, we replaced brawn with steam. In the second, we replaced steam with electricity, and in the third, we introduced computers," says Guido Jouret, chief digital officer for Swiss industrial corporation ABB. "We've had intelligent rule-based systems. What we haven't had is the equivalent of the human cortex—systems that can learn."

That's what AI technologies represent in the current period of technological change. It is now critical to carefully consider the future of AI, what it will look like, the effect it will have on human life, and what challenges and opportunities will arise as it evolves.

AI gets real

While the concept of AI seems new to many, it has been around for decades. The phrase "artificial intelligence" was coined in 1956, said Genevieve Bell,

Cellini/Shutterstock

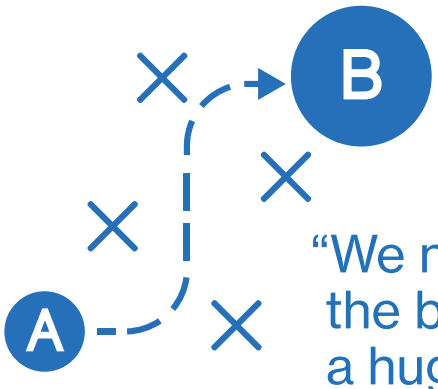

"We need to create a framework that enables the benefits of AI. Society shouldn't discard a hugely beneficial technology just because companies would feel liable."

Guido Jouret, Chief Digital Officer, ABB

distinguished professor of engineering and computer science at the Australian National University, during a presentation at the school in March 2018. The branch of computer science it came to describe was wildly popular for several years, with researchers optimistic that a fully intelligent machine was imminent– "and then not so popular because it turned out it was really hard," Bell said. "It required all this computing that people didn't have."

Today, the supporting technologies have matured to the point where AI is now practical and effective. As it becomes widespread, it brings with it societal and ethical dilemmas. Questions and concerns about the impact of AI are becoming urgent. "What is all that technology going to mean for human beings, our systems, our institutions, our organizations, and our countries?" asked Bell.

AI observers agree it is essential to address how society should approach overarching philosophical and practical questions when planning for the future of AI. Here's how Bell breaks it down in her research:

AUTONOMY: Will AI systems be autonomous, and should they be?
AGENCY: Who will set limits and controls and ensure they are consistently applied?
ASSURANCE: How will humans accommodate safety, risk, liability, trust, privacy, and ethics?
MEASUREMENT: How will we measure the effectiveness of AI systems?
HUMANITY: How will humans interact with AI-driven systems?

Actual autonomy?

There's no consensus on what being autonomous really means for artificial intelligence systems. "Think of autonomy as a matter of degrees," Jouret says. "We are ultimately applying increasing levels of automation. And there are certain areas where that could be delivered." He cites vehicles operating in a mine or cargo ships at sea as controlled areas to test different stages of autonomy.

"I say 'autonomous,' you think sentient, conscious, and self-aware. And if you grew up reading science fiction, you know what happens next," said Bell during a presentation in San Francisco in October 2018. In reality, systems can operate autonomously without being sentient or conscious. Autonomy simply means operating without referring directly to a set of rules or procedures.

That raises questions of ethics and societal impact. "What does it mean to have systems that will be autonomous? How will we feel about that? What will it mean?" Bell explored.

Ste Kelly/Shutterstock

Agency and control

Establishing controls and guidelines on the use of AI-driven systems is in its infancy, but it is critical to consider this now, Bell and other observers believe. "We do need to have governance. We need to create a framework that enables the benefits of AI," says Jouret. But "society shouldn't discard a hugely beneficial technology just because companies would feel liable."

With autonomous vehicles, for example, the rule sets are embedded in onboard technology for faster response and processing times. But for every rule, there are exceptions. "How do we manage and govern what AI can and should be able to do?" says Jouret. "You know how you should conduct or apply yourself; we should have something similar with AI. We should take some of those principles we apply to people and apply them to AI. We need some sort of code of conduct."

And what about when rules need to be overridden? Bell gave the example of getting cars off the road to get a fire truck through. Making that happen becomes immediately complex. "How are those rules going to work across boundaries, across countries, across cultures?"

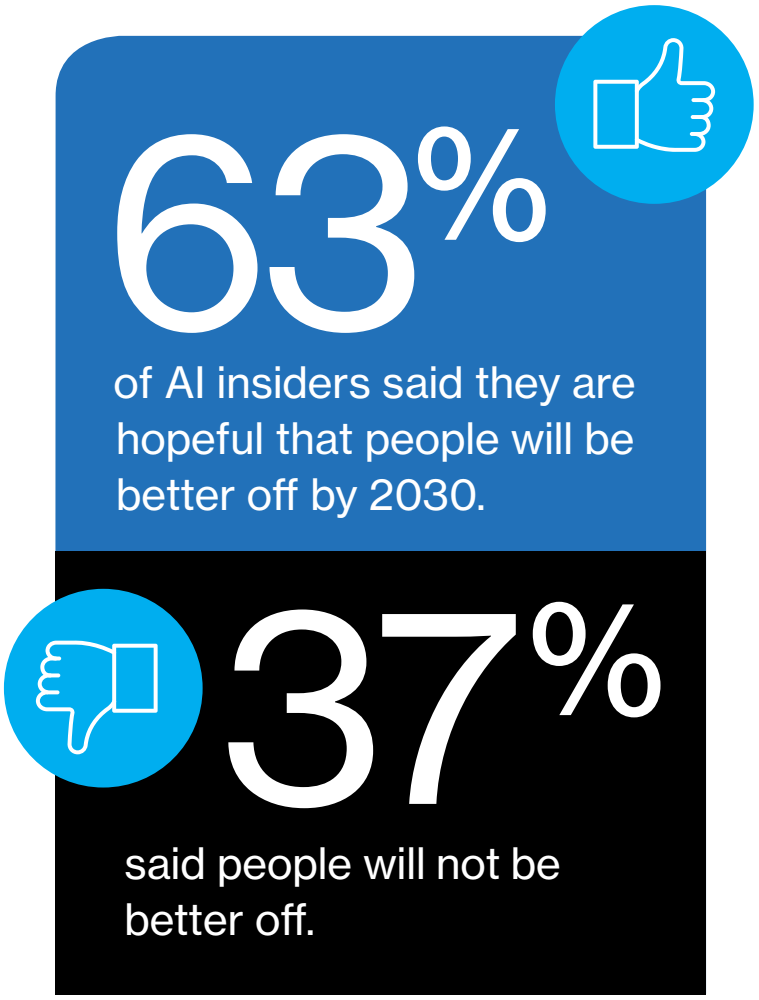

Source: Pew Research Center and Elon University's report "Artificial Intelligence and the Future of Humans," based on a survey of 979 tech innovators, developers, business and policy leaders, researchers, and activists, conducted in summer 2018

Practical matters

Legal matters like risk, liability, trust, privacy, ethics, and manageability are also important. While AI systems maintain a certain degree of autonomy, they must still be managed, governed, regulated, and even insured as a component of society. "The counterpart of acceptable risk could be insurance," says Jouret. "We could ask insurance companies, 'At what point would you insure this?'"

Safety considerations for AI-driven systems manifest most vividly in autonomous vehicles. "With the introduction of safe autonomous cars, people will gradually build confidence in the technology, in the same way that confidence in air travel was established over time," says Sumit Sadana, chief business officer for semiconductor producer Micron Technology.

As the data streams that drive AI grow and evolve, AI systems will function more efficiently, which will affect all use cases. "Which conditions produce good outcomes and of which outcomes should we be distrustful?" Jouret ponders. "The way we do AI and machine learning today doesn't allow for self-inspection. It can't justify its actions. We have to move from, 'How did you make this decision?' to 'Why did you make this decision?'"

A matter of metrics

It's important to measure the effectiveness of technology, especially new or newly expanded systems. "When AI safely goes to scale, how do we decide if it's a good system or a bad system?" Bell quizzed her audience in San Francisco. "Was that system efficient and productive? Did it save time or did it save money?"

Matsabe/Shutterstock

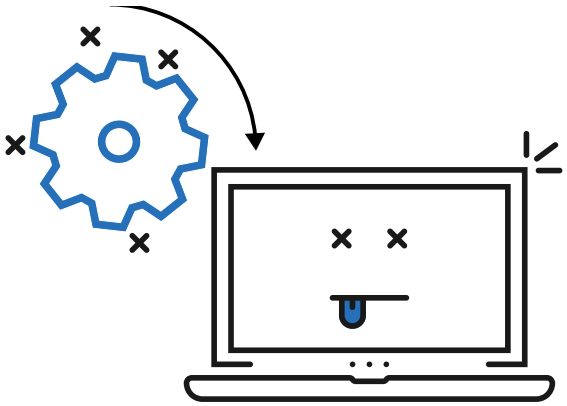

"AI and machine learning are very powerful, but if you apply those algorithms to bad data, you'll only get bad results faster."

Judith Hurwitz, President and CEO, Hurwitz & Associates

When planning for an emerging technology such as autonomous vehicles, for example, metrics typically revolve around safety issues. But that means figuring out what to base metrics on to ensure the technology is operating at expected levels.

Most of the measureable challenges regarding AI output center on the quality of the data—the size and quality of the data sets have a demonstrable impact on the relative accuracy of the outcome. "AI and machine learning are very powerful, but if you apply those algorithms to bad data, you'll only get bad results faster," says Judith Hurwitz, president and CEO of management consultancy Hurwitz & Associates. Data sets can also introduce their own bias or slant depending on the source. "What are the sources of the data? Who has touched the data? Is there bias in the data?"

Human interaction

Then there are lingering questions of how humans will interact with AI-driven systems. Often, these systems lack the same level of manual input to which most people are accustomed. Smart buildings constructed with technology at the core often have controls that are entirely foreign to most people, at least at first. Elevator controls with no buttons that are instead driven by voice commands, lighting controls with no standard on-off switches, and heating and cooling controls with no thermostats are just some examples of AI systems that don't fit current expectations.

For humans to successfully navigate a smart building or ride in an autonomous car, there will be learning curves. In her San Francisco presentation, Bell stressed the need for simple, intuitive ways to interface with AI systems. "You do not want to get into an autonomous vehicle and have to remember whether it's a 10- to 12-character password with an uppercase and a lowercase and alphanumeric."

Manage the revolution

There are straightforward technological questions to address as AI emerges as a technological force. Questions of processing power, connectivity and bandwidth, and security and access control must all be considered. The more complex and nuanced questions concern how society and humanity will manage AI-driven systems.

"We should set certain conditions, like mission-critical situations, that require human validation before acting," says Jouret.

World Economic Forum founder and executive chairman Klaus Schwab agrees it's more important to focus on the human component than on the technological. In his 2016 book, *The Fourth Industrial Revolution*, he wrote, "Shaping the fourth industrial revolution to ensure that it is empowering and human-centered, rather than divisive and dehumanizing, is not a task for any single stakeholder or sector or for any one region, industry, or culture. The fundamental and global nature of this revolution means it will affect and be influenced by all countries, economies, sectors, and people." Though the world is changing fast, he adds, "we can still shape our future in a way that benefits all."

Pranch/Shutterstock

A history of tech revolt

Rapid industrialization started in the 18th century with advances in manufacturing and has continued to the emergence today of the internet of things (IoT), artificial intelligence (AI), and autonomous technology.

The Industrial Revolution

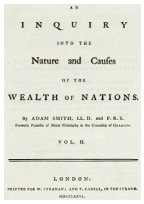

1812 The first gas lighting company was chartered in Great Britain.

1776 Adam Smith publishes *The Wealth of Nations;* Scotsman James Watt's steam engine goes into production.

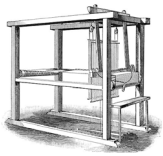

1733 British inventor John Kay patents a weaving machine called the flying shuttle.

The Technological Revolution

1857 Robert Forester Mushet, a British metallurgist, crafts rails out of steel, a more durable material than iron. Construction of a rail transportation network soon accelerates throughout the world.

1881 Englishman Sir Joseph Swan supplies his invention, the incandescent lightbulb, to the Savoy Theatre in London—the first public building in the world to be lit entirely by electricity.

1886 Karl Benz, a German inventor, patents the world's first practical automobile.

1971 Intel develops the Intel 4004, an early microprocessor.

1947 The transistor is invented following a series of experiments at AT&T's Bell Labs in New Jersey.

The Digital Revolution

1977 Home computers hit the commercial market.

1989 Tim Berners-Lee invents the World Wide Web. It becomes publicly available in 1991.

The Fourth Industrial Revolution

2004 The US Department of Defense begins sponsoring challenges that fuel autonomous technologies.

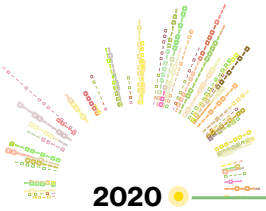

2020 The number of devices connected to the IoT is projected to reach 30 billion.

2017 The Asilomar Conference on Beneficial AI is held in California to hammer out guidelines for producing AI that benefits society.

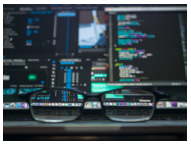

2015 Machine-vision systems for the first time outperform humans in a series of vision-related tasks.

Getting smart about the future of AI is an executive briefing paper by MIT Technology Review Insights. It is based on research and interviews conducted in December 2018. We would like to thank all participants as well as the sponsor, Intel. MIT Technology Review Insights has collected and reported on all findings contained in this paper independently, regardless of participation or sponsorship.

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the US and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. And through its growing MIT Technology Review Global Panel, Insights has unparalleled access to senior-level executives, innovators, and thought leaders worldwide for surveys and in-depth interviews.

About the sponsor

Headquartered in Santa Clara, California, Intel is the world's leading technology company, powering the cloud and billions of smart, connected computing devices. Together, these connected devices are generating millions of terabytes of data every day. Recent breakthroughs in computer and data science now give us the ability to analyze and derive immense value from that data. At Intel, we provide the technological foundation of this new world of data, enabling important emerging technologies like artificial intelligence, autonomous vehicles, and 5G networks. Learn more about how Intel is expanding the boundaries of technology to make the most amazing experiences possible at **www.intel.com.**

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2019. All rights reserved.

Photo credits for timeline: Flying shuttle: Wikimedia Commons. *Wealth of Nations*: Wikimedia Commons. Gaslamp: Cattallina/Shutterstock. Robert Forester Mushet: Wikimedia Commons. Lightbulb: Shutterstock. Karl Benz Patent Motor Car: DaimlerChrysler AG/Wikimedia Commons; Creative Commons license, https://creativecommons.org/licenses/by-sa/3.0/legalcode; no alterations. Transistor: Unitronic/Wikimedia Commons; Creative Commons license, https://creativecommons.org/licenses/by-sa/3.0/legalcode; no alterations. Intel Chip 4004: Thomas Ngyuen/Wikimedia Commons; Creative Commons license, https://creativecommons.org/licenses/by-sa/4.0/legalcode, no alterations. Commodore PET 2001: Tomislav Medak/Flikr; Creative Commons license, https://creativecommons.org/licenses/by-sa/3.0/legalcode; edited by Bill Bertram in Pixel8. Tim Berners-Lee: ITU Pictures/Flikr; Creative Commons license, https:creativecommons.org/licenses/by/2.0/legalcode. Connected truck, factory, machine icons: Buffaloboy/Shutterstock. Machine-vision: Kevin Ku/Pexels.com. Lectern: Mehsumov/Shutterstock. Data burst: Amiak/Shutterstock.

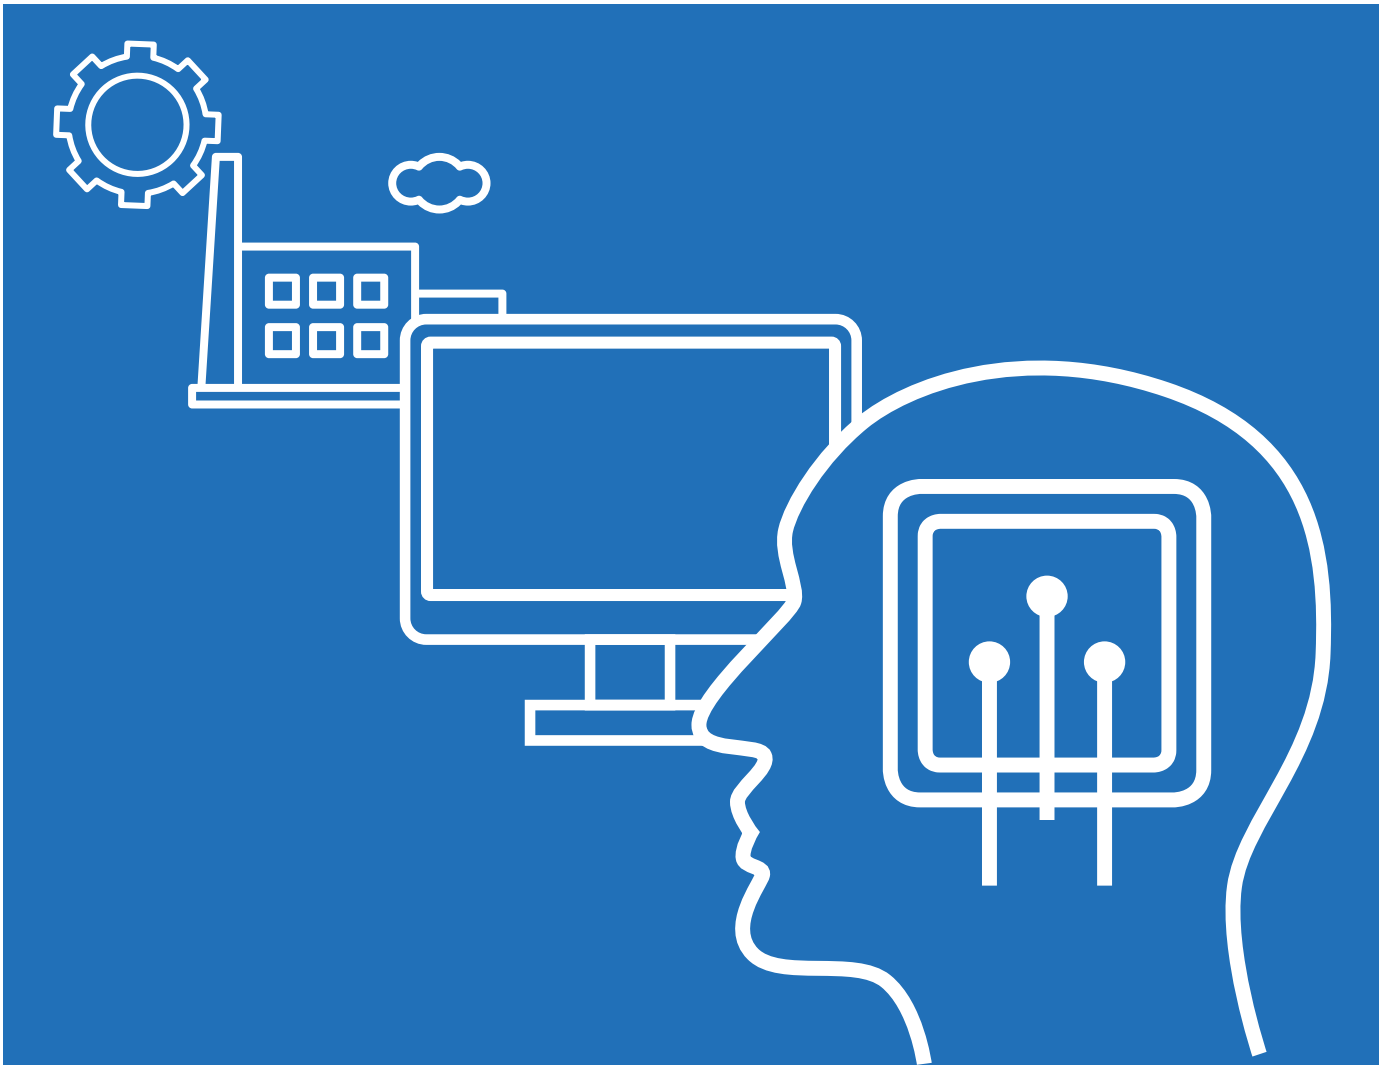

MIT Technology Review Insights

insights.techreview.com
@techreview @mittr_insights
insights@technologyreview.com