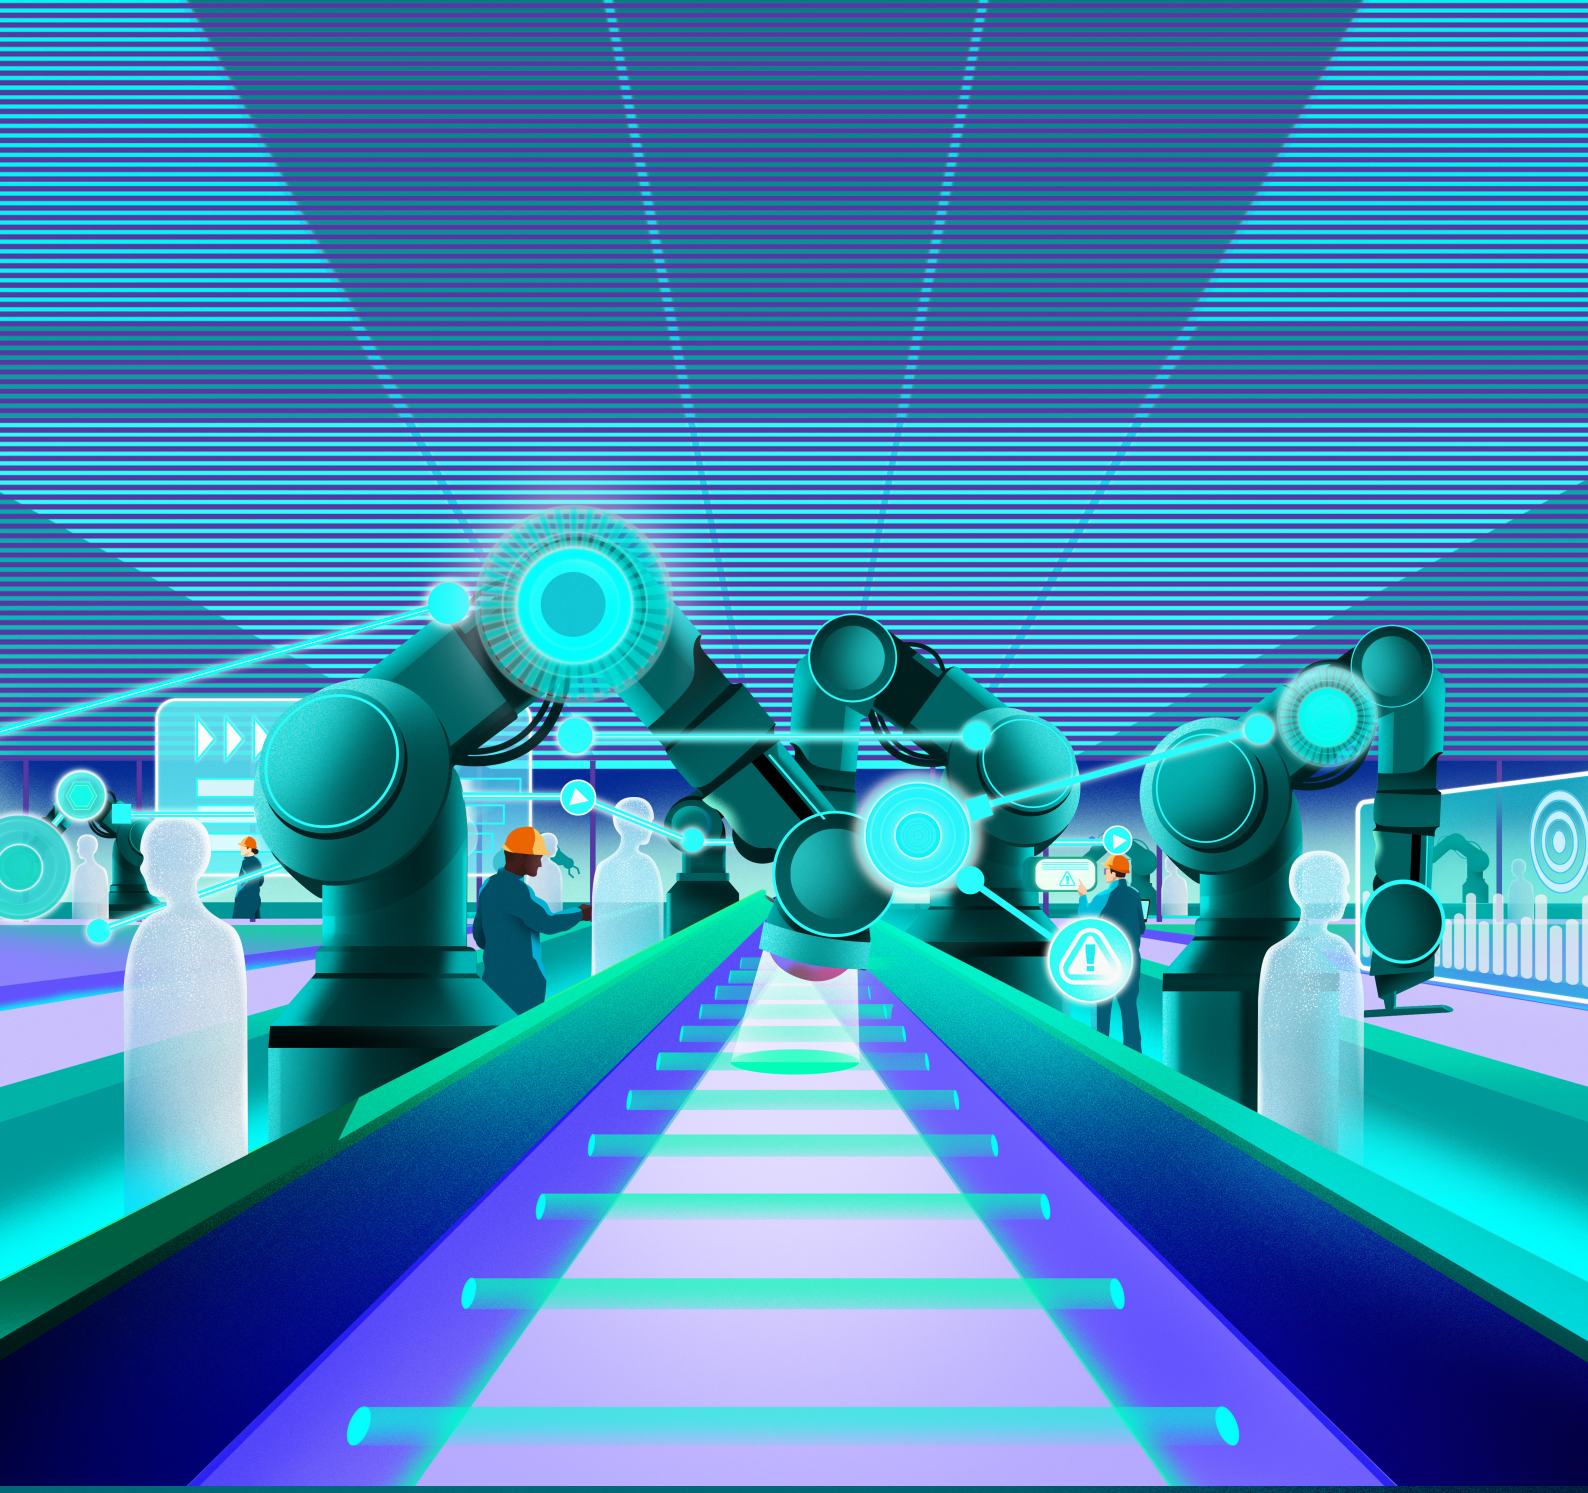

Living Factories

What AI Agents Could Mean For Industry

WIRED Consulting.

SIEMENS

2035. A car plant outside Munich.

The control room buzzes with quiet efficiency, its displays showing real-time data from the symphony of automation unfolding on the floor below.

An alert flashes on the main dashboard—an ultrasound sensor detects a micro-fissure in the casing of a high-voltage battery cell. AI analyzes data from an installation robot, and discovers that it applied uneven pressure during assembly, causing the defect. To stop it happening again, the system simulates a fix via the robot's digital twin then uploads new settings. Since the incident crossed a safety threshold, a human is kept in the loop throughout.

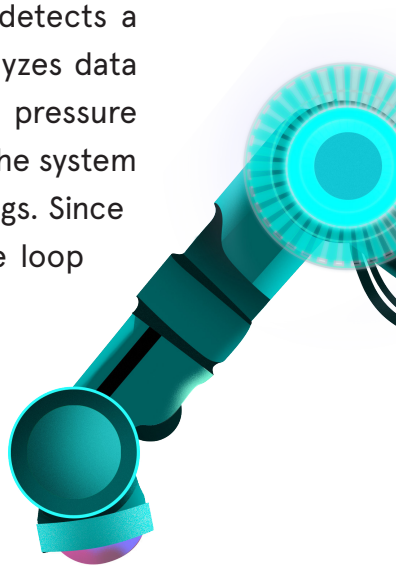

The control team requests that the system runs additional checks on other batteries. The factory pauses battery assembly, and dispatches automated transport to ferry units for scanning. Meanwhile, AI dynamically reschedules factory-wide workflows, and redirects equipment to new tasks.

With all checks passed, production resumes. The incident data is fed back into the system, which is also deployed at the company's other facilities around the world. The factories, collectively, learn and evolve.

This scenario illustrates a vision for intelligent manufacturing that is gaining traction among leaders across industries. It's a future where the physical and digital worlds are more profoundly converged, where humans function as strategic commanders, and where everything from single machines to complex production lines and networks of factories work together optimally and autonomously with minimal disruption. And it's all premised on a big idea with a simple name: AI agents.

Agents of change

When ChatGPT launched in 2022, it sparked an AI boom. This has centered around generative AI, a class of AI models such as large language models (LLMs) that can not only create content in a coherent, context-aware fashion but handle an endless range of ad hoc questions—and serve as the backbone for domain-specific tools.

AI agents are widely seen as a key next step in generative AI's evolution. Definitions vary, but typically the phrase is used to describe AI that can 'do' things. "This mainly means it has the ability to use other digital tools," says Armin Hadzalic, Senior Software Developer at Siemens Digital Industries. "To do that well it also has to have reasoning power—the ability to plan stuff—and a memory function. But the concept is evolving, and these days you might also expect it to have a feedback loop inside it, too, so it can learn and improve."

> **"I see a very, very big shift ahead, starting in the next two to five years — the biggest shift I've seen in my career. It will change industry altogether."**
>
> **Armin Hadzalic**
> Senior Software Developer, Siemens Digital Industries

Agents can operate in isolation or in concert with, theoretically, a potentially vast number of other agents to get things done for their human operators. After all, why stop at one agent when software scales so easily?

Agents have become a preoccupation in industry, and their abilities are improving fast. The length of tasks AI agents can reliably handle has been doubling roughly every seven months. At a time when industry faces challenges around costs, time-to-market, supply chain risks and labor shortages, it's not hard to understand why AI agents hold appeal.

"I see a very, very big shift ahead, starting in the next two to five years—the biggest shift I've seen in my career," says Hadzalic. "It will change industry altogether."

AI agents in industry today

Basic versions of agents are already making an impact in the sector.

If a fully fledged agent is one that reasons autonomously, selects what digital tools to use for any given task, and iterates its work to achieve a higher order goal, then the ones that manufacturers are piloting today are typically a stepping stone towards that vision. In many instances, they may more accurately be described as an "AI workflow"—an LLM using predefined tools, and operating along predefined paths, with a human responsible for any iteration. However, the language is still in flux, and many businesses call these agents.

The existing use cases tend to fall into two categories.

The first is insight. In the past, if a factory worker wanted to know what the cycle time was on the night shift, or find out if there were any anomalies in a production run, they would have to manually trawl through the data. Now, AI lets them simply ask the question in natural language. The agent then interacts with the relevant tools and data sources, combs through the information and returns a direct answer—it turns a lengthy, specialist process into a quick, accessible one.

The second major use case is recommendation. These are agents which monitor data and proactively suggest actions. If the AI notices an increase in scrap during a shift, for example, it might suggest changing the lens on a vision sensor to improve performance.

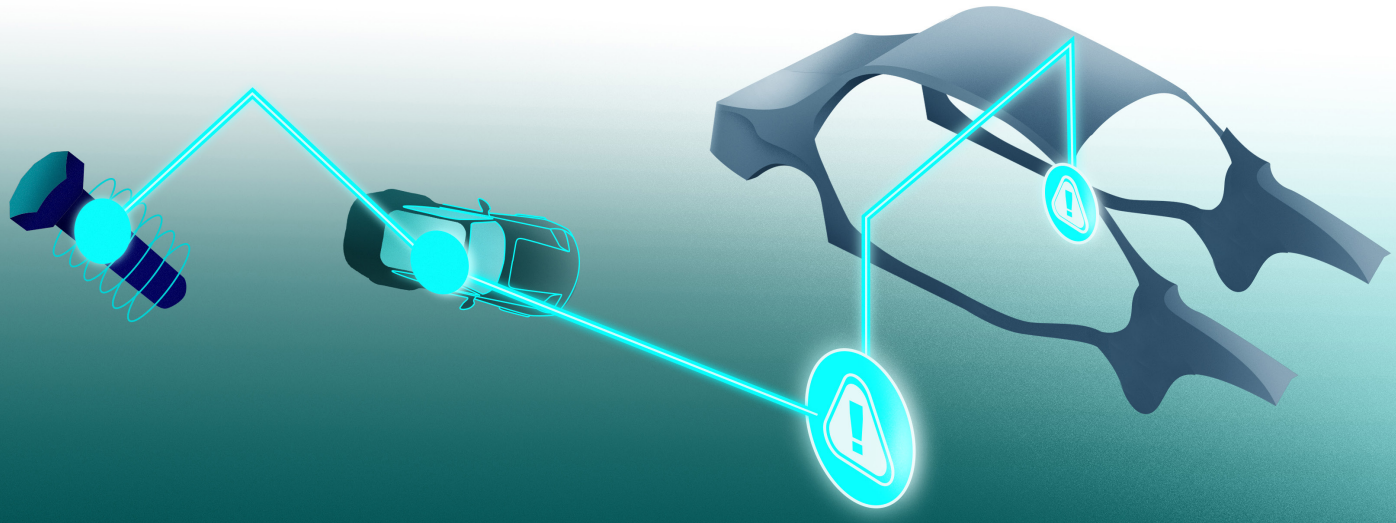

Although these early types of agents exhibit limited decision-making autonomy, the value can still be significant. "We have a customer that manufactures components for the food processing industry. They have maybe 150 factories worldwide, and supply chain is a big issue for them," says Lina Huertas, Industry Executive for Manufacturing at Microsoft UK, which has strong partnerships in industrial AI. In the past, surfacing high-priority supply chain problems would take days of research every week. Now, this customer is trialing a system that ambiently tracks a wide range of data sources and proactively flags issues. "Every Monday you come in and you get a list of, say, 10 problems organized by order of priority. The first thing they're going to do is probably have a team meeting to solve the big problems, so you're in action right from the beginning. It's game changing for organizations."

> **"It used to be quite complex to get all this information together at one point. Via the agent it's very easy to access all this information."**
>
> **Marco Goergmaier**
> Vice President of Enterprise Platforms and Services,
> Data and Artificial Intelligence, BMW Group

BMW Group has had similar success with insight agents. The company is deploying agents across various aspects of its manufacturing operations—from production and logistics to customer service and HR—and says it has derived particular value from an agent that helps collate all the information about a vehicle before it is delivered. This involves fetching information from disparate systems about the vehicle's production, supply route, risk assessment and more. "It used to be quite complex to get all this information together at one point," says Marco Goergmaier, Vice President of Enterprise Platforms and Services, Data and Artificial Intelligence at BMW Group. "Via the agent it's more convenient to access all this information."

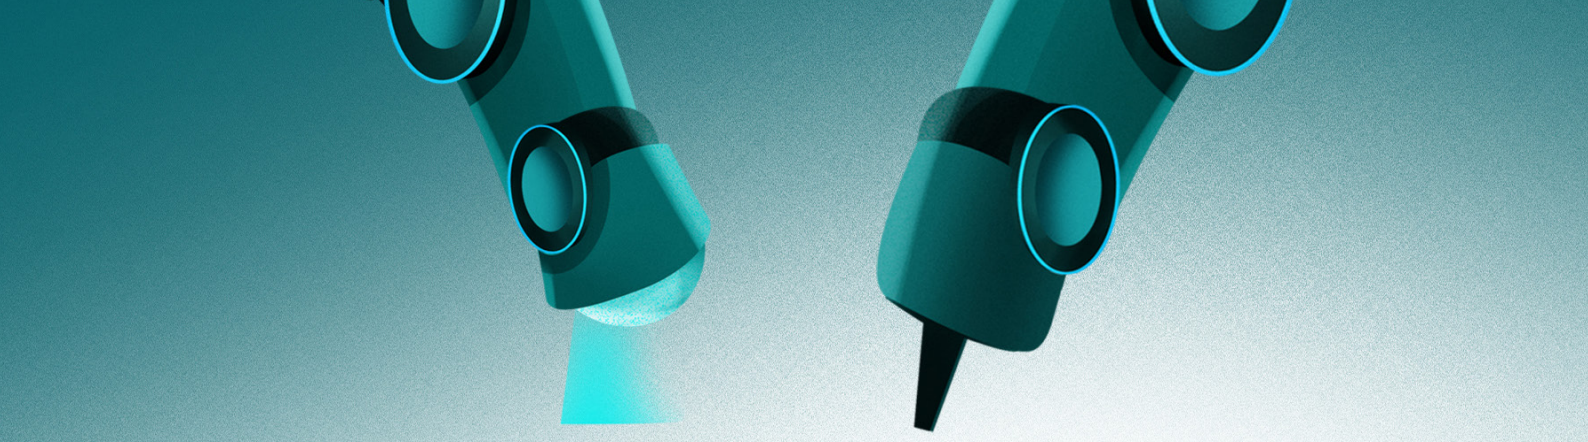

But there is also a third, more sophisticated use case that's emerging: execution. This is where the agent doesn't merely generate an insight or recommendation but performs an action in the real world. This is a more frontier idea, especially for industry, since the sector deals not just in bits but also atoms. "If you are a data analyst and you ask AI to analyze your spreadsheet, and the result is wrong, there's no physical damage," says Hadzalic. "But in industry, there is a real, physical effect on a machine."

Siemens is currently exploring the idea with its Siemens Industrial Copilot: a suite of AI-powered assistants designed to optimize workflows and enhance human-AI collaboration. It comprises specific AI tools for different areas of the industrial value chain, from design and planning to engineering, operations, and service.

An example of an execution agent can be found in the development pipeline for the suite's Production Engineering Copilot. This is a tool that helps customize the functioning of industrial machinery, whether that's a conveyor belt, CNC machine, a robot—anything you might find in a factory. Traditionally, doing so involves a number of stages. First there's designing the new process, then there's creating the computer code to make it happen, and then there's testing before, finally, the code is uploaded to the machine. Now, Siemens is developing an execution agent to handle this whole process. "Imagine you want to tell your conveyor to adapt its speed depending on whether plastic or metal components are on it—you need to do some image recognition on the conveyor and you need an algorithm to change the speed," says Hadzalic. "Now we can use agents to do everything: describe what we want to have in the machine, generate the code, and so on."

Currently, a human checks and approves the agent's work at every stage. "But this might change in the future, because we simulate everything with a digital twin approach, and based on the outcome we can automate the whole process." And at that point, true agents would start to come into view.

What's holding agents back?

The vision for the future of industry sketched at the start of this piece is clearly far ahead of what is in deployment right now. So how do we get from here to there?

There are a range of challenges to address. One is integration. Although agents can handle data in different forms and modalities, a lack of standards remains an issue. Mission-critical systems can't accept the risk of agents translating syntax on-the-fly each time they communicate with other platforms or agents; they need consistent, state-persistent data rails to guarantee the required level of accuracy and dependability. "When you look at the enterprise landscape of a company, it's complex: software-as-a-service stacks, self-developed applications, legacy systems," says Goergmaier. "In order to have a real autonomous agent, you need to be able to write back information to the systems, and it needs to be done in a way that is absolutely secure and correct. Solving this will be decisive for scaling trustworthy, production-grade AI." What's more, if an agent is handling sensitive data, it needs to be able to communicate that data to other agents or systems in a way that is leak-proof and regulation-compliant.

> **"In order to have a real autonomous agent, you need to be able to write back information to the systems, and it needs to be done in a way that is absolutely secure and correct."**
>
> **Marco Goergmaier**
> Vice President of Enterprise Platforms and Services, Data and Artificial Intelligence, BMW Group

Emerging communication standards for agents are expected to provide the universally accepted digital plumbing that will resolve this issue. Examples already in play today include Google's Agent2Agent standard for AI-to-AI interactions, and Anthropic's Model Context Protocol standard for AI tool use. Beyond enabling structured communication, these standards offer in-built security, and should make it easier for organizations to implement consistent policies across systems.

Connected to the need for security are issues of control and observability. Getting the most out of an agent means loosening the reins, yet its decision-making may not be fully transparent. It's important to avoid trusting it more than is safe. "As AI capabilities continue to evolve, it's easy to imagine simply deploying an agent and expecting it to autonomously manage whole functions," says Huertas. "But true innovation comes not just from powerful technology, but from responsible implementation. That's why we emphasize and enable the importance of strong governance and clear guardrails—ensuring AI systems are focused, transparent, accountable, and aligned with human oversight."

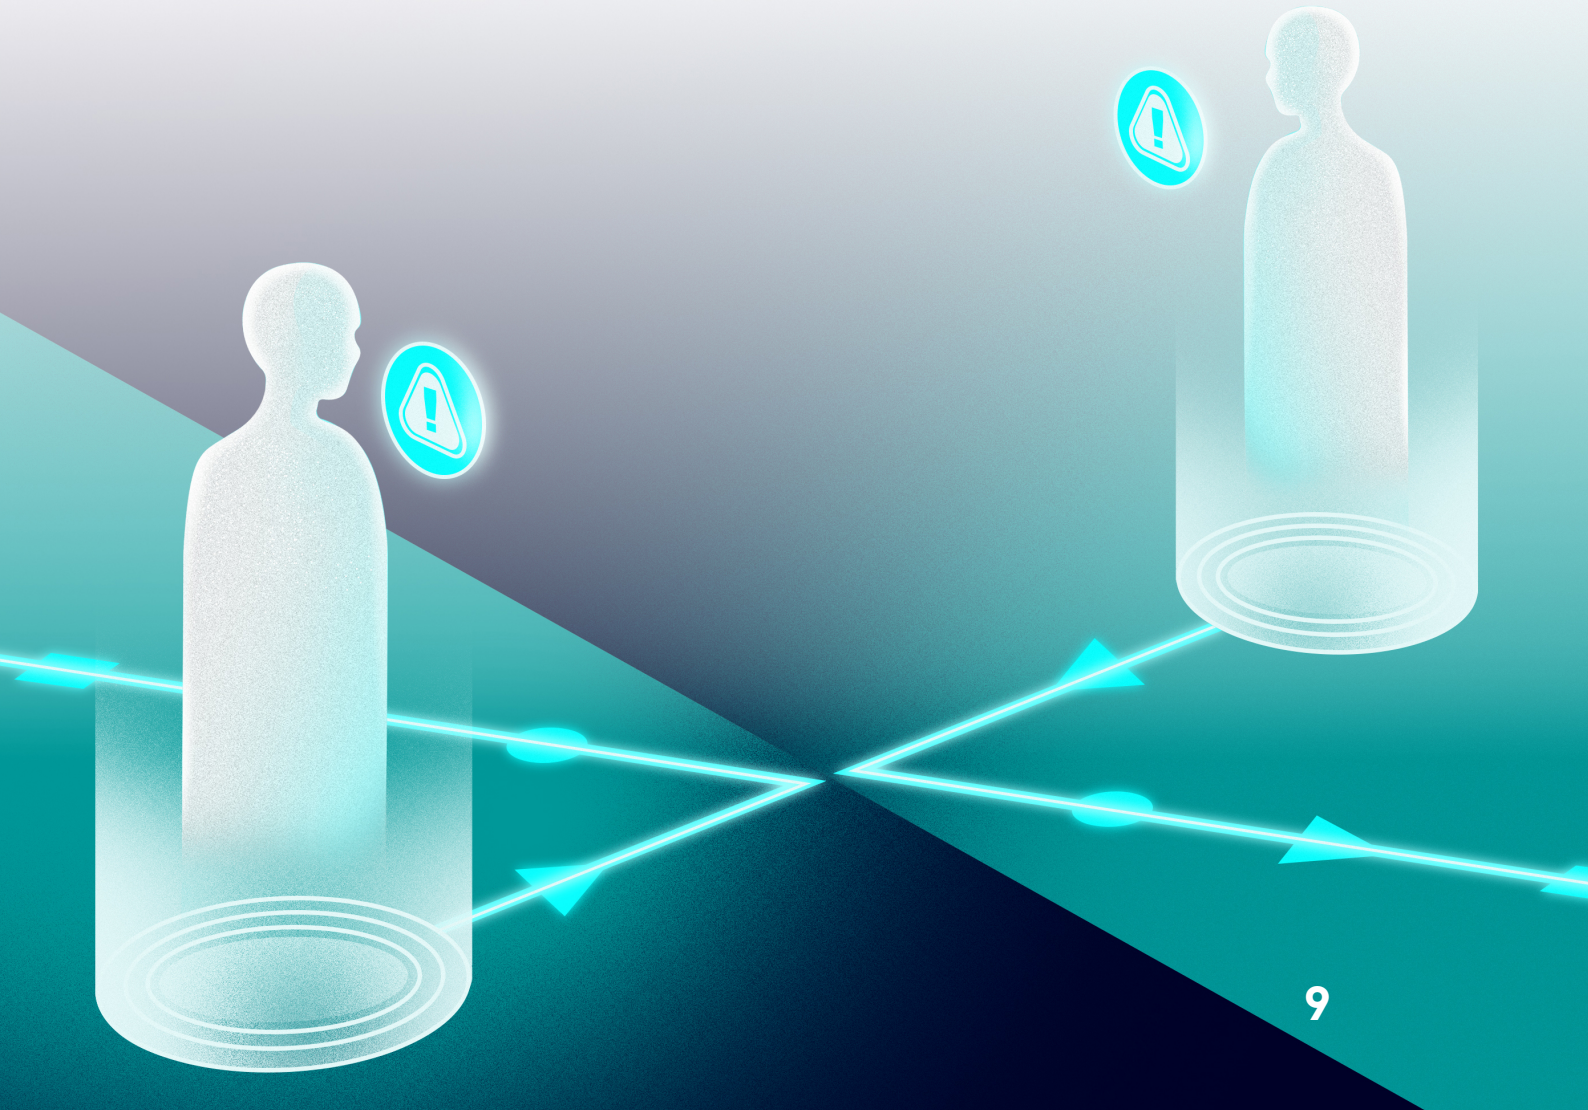

While issues like these don't seem fatal, there is a tougher challenge facing true industrial agents—one that is common to all forms of generative AI. Understanding intent and comprehending information is an issue even for basic agents. If they are asked to help with a really complex scenario, the number of variables can be "too blurry" for the system, says Hadzalic, "so it may not find the information you want to extract, whereas a human would be able to distinguish it."

More generally, consistency, errors, and hallucinations—where the error is plausible, making it hard to spot—are core problems with anything running on generative AI models. Some of this is the result of the models' probabilistic nature; some the result of how they're built. "Agents' greatest limitation today is that they lack general intelligence," says Peter Koerte, Managing Board Member, Chief Technology Officer and Chief Strategy Officer at Siemens. "People can generalize and understand context. AI agents are heavily dependent on the data they're trained on. If there are gaps or biases, this will show up as inaccuracies."

"At BMW Group, AI never flies solo; a human expert reviews almost every decision we automate. This is very important."

Kathrin Berkel
AI Self Services Lead, BMW Group

For now, unreliability is a key limitation for agentic systems. The picture is steadily improving but even a small error rate—perhaps even just one percent—can prove critical once scaled across a long chain of tasks. For now, industrial organizations deploying AI in complex processes mitigate the risk by keeping humans involved in decision-making processes. "At BMW Group, AI never flies solo; a human expert reviews almost every decision we automate," says Kathrin Berkel, who leads AI Self Services at BMW Group. "This is very important."

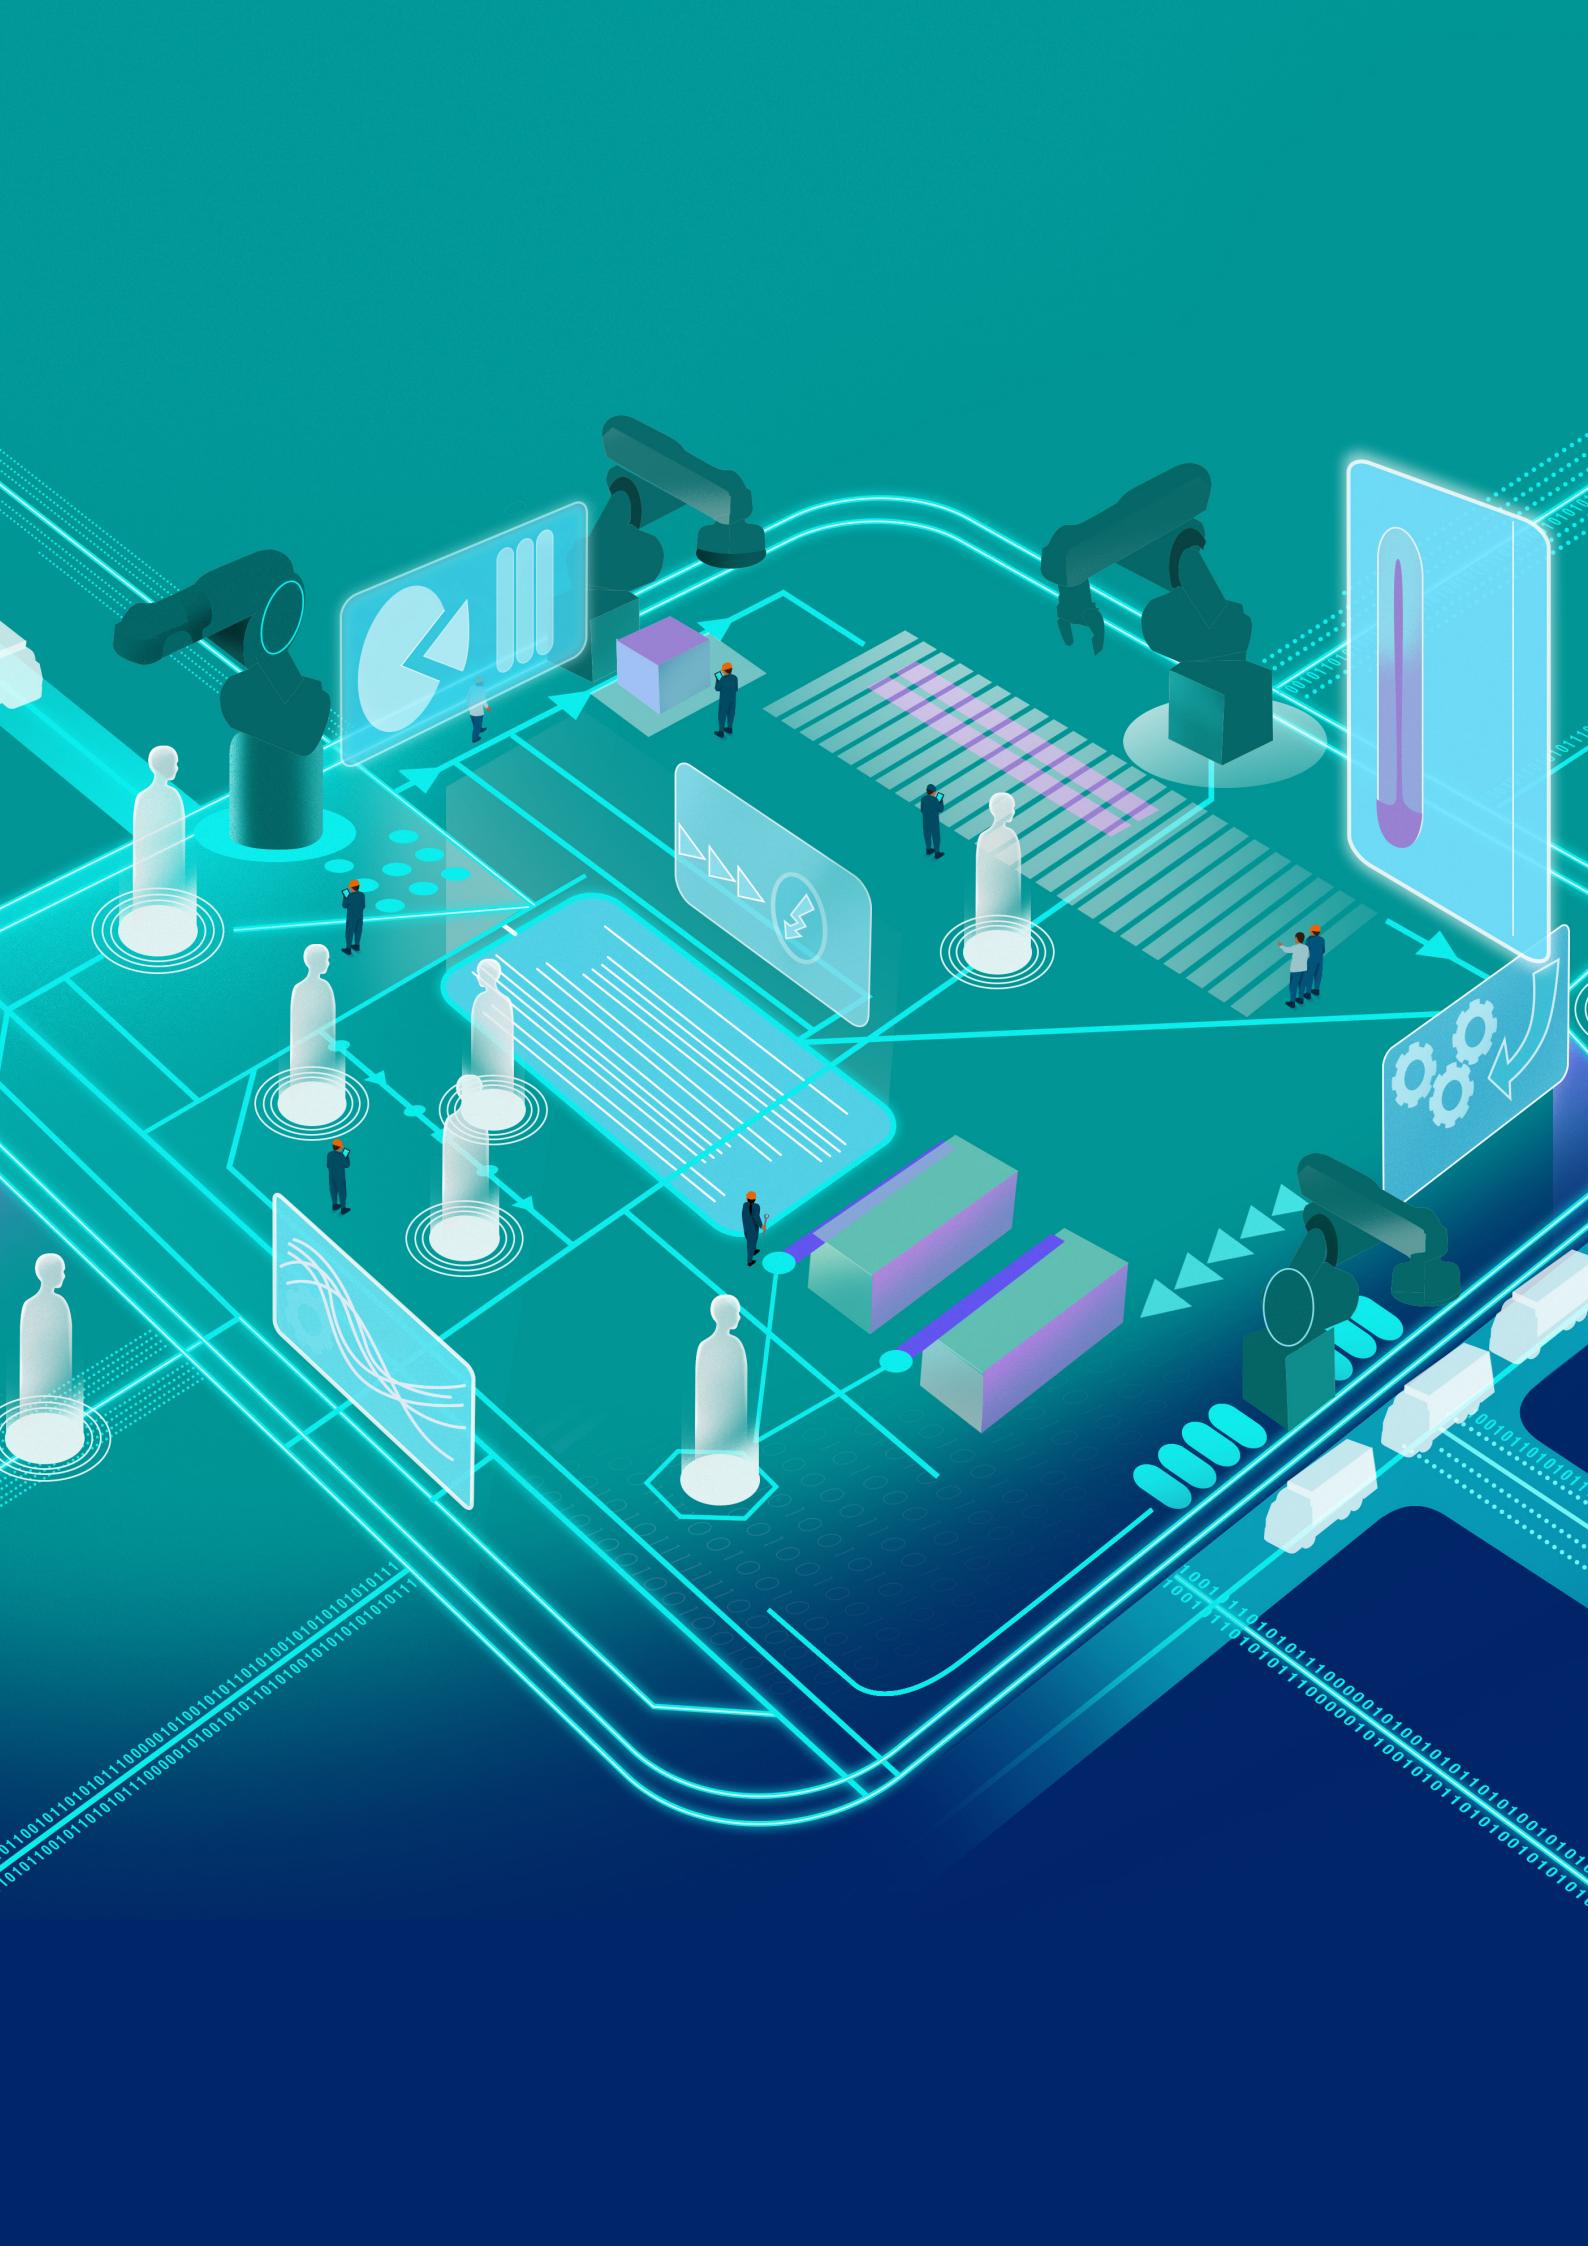

The next generation of agents

Agents are a central focus for all the frontier AI labs, every major tech company, and any enterprise with developer chops. As such, there is plenty of work going into helping agents become more reliable and able to handle tasks with greater autonomy. Key progress areas include developing advanced planning frameworks, self-evaluation mechanisms and simulations, and rigorous benchmarking frameworks for testing prior to deployment. There is also an emerging body of good practice. Berkel says that BMW, for instance, "provides agents with a lot of context and puts guardrails in place to limit the agents' actions to that context" as a method of improving reliability.

Interest has also coalesced around the idea of orchestrating agents. Today's agents are task-specific agents, often referred to as 'narrow' agents, but—as Koerte says—there is a need for general intelligence. One answer? "You can create a large team of agents, each with their own capabilities. Many one-trick-ponies can be very useful, if you have a 'manager' to delegate tasks and keep them in check," says Koerte. "That's the so-called 'orchestration agent'. It's the digital equivalent of a project manager."

Basic versions of multi-agent systems are already reality, and are critical for highly complex systems. Dr Jay Lee, Director of the Industrial Artificial Intelligence Center at the University of Maryland, cites the example of a system he helped create for a manufacturer of semiconductors. The manufacturer uses a technique called extreme ultraviolet lithography (EUVL) to make its chips, and wanted to reduce costly downtime on the EUVL production line to zero—an extremely challenging request. To do it, says Lee, "you divide and conquer". Lee conceived of separate AI workflows to monitor each component in the EUVL system—one for the laser beams, one for the motion system, one for precision feedback, and so on—and then combined them together to "create a mission-focused hierarchical coordination function overseen by an orchestrator that could talk to a human user."

A fully self-directing factory will require more sophisticated multi-agent systems—and not only in terms of executive autonomy. Perhaps, in years to come, 'orchestration agents' may be overseen by 'meta-orchestrators' that ensure seamless interaction between different groups of agents, creating a cohesive ecosystem that might even extend out of the factory and across the supply chain, unlocking a new level of macro efficiency.

Teaching AI the language of industry

To make this ambition reality, agents will have to develop greater literacy in the world of industry. One approach is to create specialized models. Off-the-shelf foundation models are trained on publicly accessible data such as the internet. This makes them good at things such as writing marketing copy, as there's plenty of that online. But it means they struggle with specialist work—say, designing custom, regulatory compliant parts for a unique machine. The answer is to extend the model by putting it through a further training run using a specialist data set. This technique, called fine-tuning, allows the model to think like a domain expert. "The goal here in one word is: precision," says Lee.

Siemens is developing what it calls an 'industrial foundation model' by fine-tuning a multimodal model. "It will understand the languages of industry," says Koerte. "This means it will be able to reason with complex industrial data that is hard to put into words - 3D models and sensor readings, for example. Building industrial copilots and AI tools on the back of an industrial foundation model would give them a much broader understanding of industrial processes. This would make AI agents even more reliable, trustworthy, and cost effective."

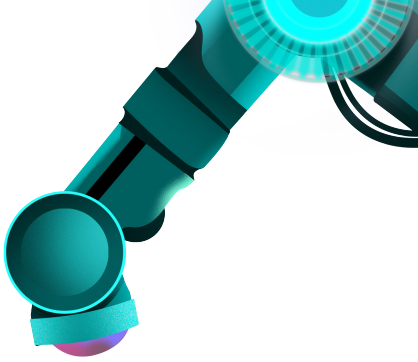

Training this kind of model would require a lot of data—more than any one company necessarily has within its own four walls. Creating a highly capable industrial foundation model may therefore require a new attitude to data. If the old formula for success in industry centered on protecting your data to realize your interests, the new age of AI may incentivize sharing data with partners to accrue enough useful information.

From doers to directors

If the reliability and data challenges are overcome, and agents enter industry at scale, this could bring about a paradigm shift in industrial operations. You could imagine, for example, factories making important, autonomous decisions: coordinating with the rest of their supply chains, automatically procuring parts and scheduling production in response to supply and demand.

Many in the field are confident that agents will reach this degree of sophistication. It raises a question: What would this mean for the people who work in industry—what kinds of roles will they have in the future? "New technologies have a history of changing job profiles—be it in industry, infrastructure, or public administration," says Koerte. "In some areas you will need less people to do repetitive tasks, but scaling up industrial operations and growing your digital workforce will create a much greater need for people to manage and optimize this digital workforce."

Koerte has some advice for anyone preparing for that future. "I would encourage everyone to try out the AI tools at their disposal. Experiment, find out how you can use AI safely in your job and explore where it can free up time for more creative tasks and problem solving," he says. "This is where we do best."

SIEMENS

Siemens AG is a leading technology company focused on industry, infrastructure, mobility, and healthcare. The company's purpose is to create technology to transform the everyday, for everyone. By combining the real and the digital worlds, Siemens empowers customers to accelerate their digital and sustainability transformations. A leader in industrial AI, Siemens leverages its deep domain know-how to apply AI to real-world applications.

Discover more at **siemens.com**.

WIRED Consulting.

WIRED is where tomorrow is realised. It is the essential source of fresh thinking and deep expertise on the technological, scientific and societal trends that are changing our world. Consulting is a division of WIRED that brings the unique WIRED network, insights and brand to commercial organisations—helping them to build internal knowledge, develop strategy and create thought-leading content that positions them at the cutting edge.

Discover more at **consulting.wired.co.uk**.