



HITACHI

Perspectives on Physical AI

An anthology of expert insights for industry

Editor's Foreword

Steve Jobs once described the brilliance of the 1990's Got milk? dairy campaign by pointing out its contradiction. "Got milk?," he said, "doesn't even talk about the product. As a matter of fact, it focuses on the absence of the product!"

The phrase "mission-critical system" could be thought of in a similar way. On its face, it sounds like we're describing a system that is critical to keeping the business running. When in fact, it speaks to a system that would cause disruption and harm if it suddenly *stopped* running.

When you consider the rising role of AI within mission critical systems, the idea of failure becomes more than mere semantics. Think transportation – railroad switching and maintenance, air travel; think energy – electrical transmission, nuclear power generation; think manufacturing – massive robotic assembly lines. Add to this the rapidly expanding threat landscape, and it's clear that failure in any of these environments because of faulty AI can result in serious physical and/or economic harm.

This world of Physical AI, and the burgeoning subcategory of Industrial AI, is not for the faint of heart. And Hitachi, a company with a heritage in industrial equipment & operational technology and decades of solutions development in data, analytics, and AI, knows this better than most.

Doing the Real Work in AI

Our experience and expertise in this area are so strong and our innovations so rich, we decided to capture some of the energy in a series of stories. First appearing in CIO.com, the following stories are glimpses into the real work Hitachi is pouring into AI solutions and practices, all designed for use in mission critical systems across industries. Neither case studies or product rollouts, these are stories of development & deployment with points of view on innovations and trends.

For example, you'll read about how companies in industries from energy to transportation are using Hitachi's Praxis Library of AI accelerators to get up and running quickly with AI. Or how its HARC Agents suite eliminates the drift in agents.

And you'll read expert insights on topics you may never have considered. Like the heightened level of scrutiny investment firms, like Hitachi Ventures, must employ when considering industrial AI startups. Or the added expertise needed to protect once 'air-gapped' industrial settings from AI-based cyberattacks.

True One Hitachi

Finally, as the stories crisscross the company, you'll begin to understand something special about Hitachi – the seamless, deliberate collaboration among Hitachi Group Companies to help customers overcome challenges and drive performance. It's what we call, True One Hitachi. And it's more

than a strategy; it's a superpower. As Editor of the program, I worked with every expert and writer on the stories within this series. I'm convinced that after reading through *Perspectives on Physical AI*, you will have a clear understanding of not only the importance of getting AI right in mission critical systems but how Hitachi is uniquely positioned at the forefront of this burgeoning corner of the industry. Enjoy.



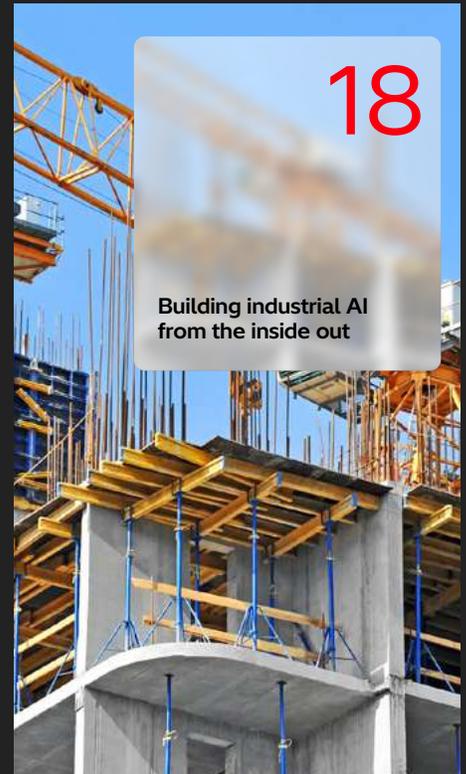
Michael Zimmerman

Head of Content Marketing and Editorial Development

Hitachi Digital

Contents

The dawn of a new world of artificial intelligence is upon us. Industrial AI is transforming how mission critical systems, from transformers to traffic management, are developed, deployed, and automated.



05 — Industrial AI: Move fast, break nothing

Industrial AI is not for the faint of heart. Understanding the criticality of mission critical systems across industries is essential for successful AI adoption.

18 — Building industrial AI from the inside out for a stronger digital core

Industrial AI is not for the faint of heart. Understanding the criticality of mission critical systems across industries is essential for successful AI adoption.

09 — Driving lower barriers, higher returns in industrial AI

Library of industry-specific accelerators helps companies speed Industrial AI adoption and scale at will.

21 — An agent you can observe is an agent you can trust

In the mad dash to AI agents, observability & explainability will be key to driving trust in industrial AI environments.

13 — A new spin on SDLC is solving a big problem in industrial AI

As companies struggle to integrate disconnected AI tools, agent-based SDLC is unlocking productivity gains in industrial environments.

24 — Industrial AI — Where knowledge (management) is power

How building knowledge management into industrial AI early, while blending the best of GenAI and logic systems, can improve performance and efficiencies.

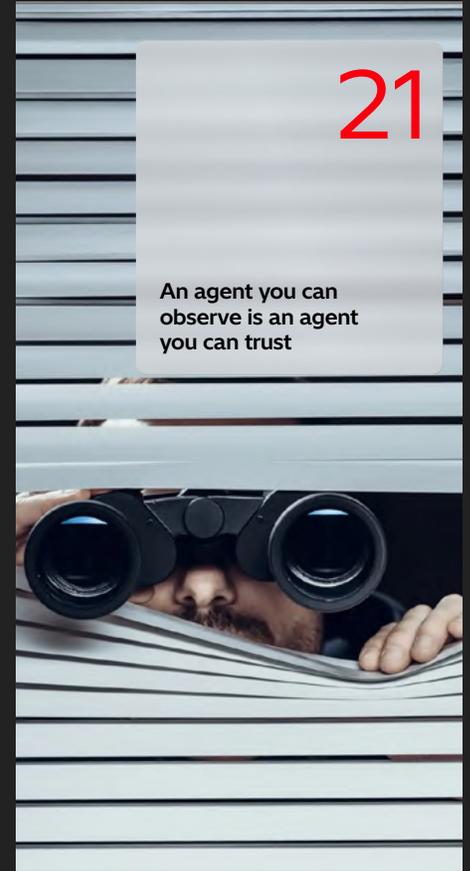
16 — Easing the pressure on the electrical grid with AI

How applying AI to some of the grid's most complex computational demands is helping improve the 'flow' of energy.

28 — An industrial approach to AI in financial services

Reliability and responsibility must be engineered into AI for financial services for resilient systems and predictable outcomes.

Contents



31 — The anatomy of AI investment

How Hitachi Ventures zeroes-in on industrial AI and automation.

34 — 'Sensing' plant floor disruption? How GenAI & OT can help

Hitachi is bridging the OT & AI divide to drive steadier operations in manufacturing.

37 — From air-gapped to wide open: The rising risks in industrial cybersecurity

Automation and AI are making it easier on cyber criminals. What's needed: A deliberate, comprehensive assessment plan based on industrial domain expertise.

40 — Sowing institutional knowledge across the industrial metaverse

How Hitachi Research is blending XR, AI and the metaverse to help industry overcome mounting challenges to the world's aging social infrastructure.

43 — The criticality of introducing AI into mission-critical systems

A Q&A with Hitachi on its approach to bringing AI into the physical world.

47 — There's (industrial) strength in numbers when building AI

How Hitachi's ecosystem of internal and external AI partners helps it tackle the world's toughest problems.



Industrial AI

Move Fast, Break Nothing

By Chetan Gupta, Ph.D.

Industrial AI is not for the faint of heart. Understanding the criticality of mission critical systems across industries is essential for successful AI adoption.

“Move fast and break things,” has been an anthem of Silicon Valley since the early 2000’s. Part philosophy, part mantra, it’s been used to motivate and accelerate software development, even at the expense of mistakes.

The mindset has led to great advances in enterprise apps, social media, gaming, and artificial intelligence (AI) – especially generative AI (GenAI), which has advanced at

a breakneck pace since ChatGPT launched in late 2022. But as with anything moving too quickly, reliability can be among the first qualities to fail. In its short history, ChatGPT has been plagued by hallucinations (completely fabricated results), challenges with biases, apprehension around privacy, copyright concerns, and more.

Industrial AI is different

The booming realm of industrial AI demands more serious consideration. When AI malfunctions in industrial settings, from energy to transportation, the impact is real. Costly damage to industrial equipment, physical harm, detrimental exposure to cyber vulnerabilities, and loss of revenue and reputation, are just a few of the things that can come from “AI-gone-bad” in industry. In other words, the health and wellbeing of business and customers, as well as society, are at stake with industrial AI.

Such is the view of Hitachi. With deep roots in industrial equipment and operational technology, its history extends from the early days of electric motor development in Japan all the way to the advanced transformers and high-speed trains of today. The company is also steeped in digital innovation, with decades of research and development in data, infrastructure, and AI – with solutions deployed across every industry.

In fact, Hitachi is among a select few companies that can fully support the challenges and identify the opportunities



When AI malfunctions in industrial settings, from energy to transportation, the impact is real.”

for AI across every industry – from utilities reimagining the grid or navigating the energy transition, to manufacturers hungry for more automation and autonomous processes, or transportation companies yearning for more uptime and greater efficiencies.

Industrial customers count on Hitachi solutions and expertise to help them visualize possibilities, architect scalable solutions, and swiftly move from experimentation to production with a clear understanding of the desired outcomes, expected ROI, and challenges of the real world.

The pace of AI is undeniable and must be met with vigor to stay competitive. But the need for speed at any cost is not

an option in industrial AI. Too much is at stake. For Hitachi, industrial AI is about moving fast and breaking nothing.

Pouring AI into industry

The criticality of industrial AI is clear. However, there's more to integrating AI into industrial settings than what meets the eye.

In addition to a heightened focus on safety and reliability, successful industrial AI implementations demand things like support for small and heterogeneous data, the incorporation of deep domain knowledge, explainability, alignment with physical and engineering constraints, and the ability to model system of systems, among others.

Complicating the environment further, data types found within industries are routinely heterogeneous. They can include combinations of sensor data, event data, maintenance records, and operational data, not to mention acoustic, vision, and multimodal documents. Understanding these variables – per industry – and how they may or may not interconnect is crucial to successful industrial AI deployments.

A quick breakdown

Small & heterogeneous data.

Industrial settings often involve data that is sparse, noisy, and highly heterogeneous, like unstructured text, sensor logs, video feeds, control signals, and time-series data. This information is typically collected through domain-specific protocols or legacy systems and is rarely curated for AI applications. Even when data is abundant, it may be inconsistently labeled or embedded in formats like operator notes, maintenance records, or audio transcripts.

Deep domain knowledge.

Only experience and expertise in a particular field can generate deep domain knowledge. Distinct from understanding horizontal applications, learning the nuances and complexities of vertical industries takes hands-on experience gained over years. Without such industrial expertise, the AI that's implemented will not live up to expectations; time will be wasted and revenue lost.

Explainability.

Growing in importance in enterprise AI, explainability is critical in industrial AI. This boils down to the software clearly describing the function as well as the expected outcomes of machine learning models. It also provides digital breadcrumbs to explain how the model arrived at

the final outcomes. A growing requirement from regulators, explainability drives trust in the models and the companies that use them.

Alignment with physical and engineering constraints.

Safety and reliability is critical to industrial operations. To ensure both are addressed, we build models that respect the laws of physics and “understand” the engineering constraints that go into designing a system. Humans intuitively understand the former and are trained on the latter. But teaching models these constraints is still an open problem.

Modelling system of systems.

Although no official definition exists for system of systems, the concept is widely known and leveraged. Essentially it describes the engineering of multiple decentralized management systems to monitor a variety of disjointed systems and predict behaviors in order to optimize performance and generate positive outcomes. In industrial environments, like rail transport, for example, one can imagine the value of intertwining predictive maintenance with ticketing and scheduling systems to maximize uptime and reliability.



Guiding your principles

After the surge of large language models (LLMs) in early 2023, researchers began picking up on a descending trend: more and more AI projects were stalling in the pilot phase. Peter Bendor-Samuel, CEO and founder of [Everest Group](#), estimated at the time that [90% of GenAI pilots](#) would not move into the production phase “any time soon,” and that some may languish there indefinitely.

Why? The reasons were myriad and spanned from pilot “fatigue” – too many pilots to manage – to an inability to identify a clear return on investment. In other words, rather than losing more time and money on their projects, organizations were abandoning them all together.

The trend was becoming almost as disruptive as the rush to LLMs in the first place. It compelled Hitachi to craft a short list of top-line [guiding principles](#) to help customers visualize their AI journey. The pragmatic approach encourages organizations to:

1. Be outcome oriented, with stated goals and functions from the outset;

2. Leverage purpose-built tools that are designed for your industry and your types of data;
3. Ensure that your AI is responsible and reliable; in fact, the more responsible the AI, the more reliable the outcomes.

Solving the use case

The guiding principles serve the industrial AI space well, which, unlike enterprise AI, demands solutions designed for specific industries. Areas like analytics, maintenance and repair, operations optimization, quality assurance, and supply chain management sound horizontal in nature. But even a cursory comparison between, for example, an energy utility and a car manufacturer reveals a critical need for expertise and specialization.

An energy utility can take advantage of AI accelerators for things like critical event command center optimization, energy forecasting and consumption, and substation image analytics. An auto manufacturer can leverage accelerators for model-based yield prediction, dynamic scheduling, and inventory optimization. Industrial AI – it’s not for the faint of heart.



The impact

Although distinct, enterprise AI and industrial AI share a common denominator: improving business performance and customer experience. Likewise, the advent of new technologies like GenAI and agents are only accelerating experimentation and adoption, as companies find new ways to leverage new capabilities.

For example, GenAI is augmenting industrial AI in significant ways and transitioning us from Industrial AI 1.0 to Industrial AI 2.0. Consider a typical industrial value chain that comprises design and engineering, procurement and supply chain management, manufacturing and production, installation and commissioning, maintenance, customer support, and circularity.

In Industrial AI 1.0, the primary data sources were time-series, and event data and included some image data and manuals. During this era, we

were focused on the problems of supply chain management, manufacturing quality and process optimization, and maintenance.

With GenAI, we saw the beginning of Industrial AI 2.0. For the first time, it enabled organizations to begin going upstream to aid in things like product and process design, and procurement; and go downstream with installation and commissioning as well as customer support. It's also helping with OT code generation, simulated data generation for automation and robotics, and make metaverse applications much more meaningful. In fact, we believe that GenAI has far more potential to impact industrial operations.

Furthering advancing the Industrial AI 2.0 era will be the increased use of AI agents and agentic systems in industry. Although it will be a steeper climb than GenAI, due to the fact that the domain knowledge in the industrial field is not as easily accessible and the cost of making

mistakes are much higher. But there's no question, the development and deployment of agents across industry will only add to the ongoing transformation of industry.

Progress and Perfection

The benefits of industrial AI are clear. Whether integrating GenAI for industrial process transformation, building an agentic architecture to aid frontline workers, or leveraging accelerators to catapult the deployment of autonomous capabilities, the market is poised to skyrocket. In fact, manufacturing is one of the fastest-growing segments of the industrial AI space. The [World Economic Forum](#) estimates that the global AI manufacturing market will balloon from its 2023 value of \$3.2 billion, to \$20.8 billion by 2028.

The real challenge for industrial companies becomes determining with whom to partner for this AI journey. For Hitachi, nothing trumps roots in industrial equipment and OT and decades of AI and digital development. Ideal partners will have experience and expertise in both. Industrial AI is too critical for anything less.

For more information, visit [AI Resource Center - Hitachi Digital](#).



Chetan Gupta, Ph.D.

Head of AI at Hitachi Global Research, and General Manager of the Advanced AI Center, Hitachi Ltd.



Driving lower barriers, **Higher returns** in industrial AI

Library of industry-specific accelerators helps companies speed Industrial AI adoption and scale at will.

By Calvin Hennick

Although great strides have been made to overcome some of the most common digital obstacles to the adoption of AI, from data cleansing to bias detection, human skills barriers have proven to be the most persistent challenge.

The nagging confluence of talent shortages and lackluster upskilling programs can leave organizations with little idea of where to begin to integrate AI, or how. This attention

deficit typically leads to sporadic and siloed model development that can grow costly, quickly, and provide limited visibility into potential returns. While daunting to traditional AI adoption, such challenges can downright choke industrial AI development, which demands far more scrutiny.

This was the concern that fueled the managed services group at Hitachi, to develop its pragmatic approach to industrial AI that views adoption from the inside out – by building AI accelerators for distinct disciplines within distinct industries – rather than bolting AI onto processes from the outside in.

The effort quickly matured into an actual classification of industrial AI building blocks – a grid, not unlike the Periodic Table of Elements – of vertical and horizontal accelerators. With this approach, companies across energy, transportation, manufacturing, and beyond, are provided multiple points of entry for AI, the ability to scale, and a clear vision of the potential return on investment.

This was the making of the Praxis Library of industrial AI accelerators.

Bringing order to adoption

“The reason AI has not been widely adopted in the industrial sectors is that most companies don’t have the resources to build and train a custom model,” says Prem



We call it the ‘asset-ization’ of AI. Every industrial application requires customization, but we’ve found the commonalities. That provides a jump start, and it’s going to accelerate and reduce the friction to AI adoption.”

Balasubramanian, chief technology officer and head of AI at Hitachi Digital Services. “It is so costly, and it requires so many training cycles for it to become accurate enough to use in production, that people just don’t have the money or the patience to do it unless it provides tremendous value.”

Hitachi has taken the lessons from its custom industrial AI builds and vast domain expertise from across the company to create the Praxis Library, which includes accelerators for everything from asset digital twins and model-based yield prediction in manufacturing, to energy forecast and consumption and substation image analytics for utilities. It

also includes cross-industry, “horizontal” AI accelerators for tasks like monitoring carbon output, tracking asset availability, and detecting collisions.

The result? Dramatically faster and less expensive AI deployment. Although the accelerators must still be tailored to the unique needs of each organization, 40 to 50% of the development work has already been done, Balasubramanian says. Rather than investing significant time and money into AI experiments that may never yield a return on investment, he adds, organizations can opt for solutions with proven success in the field.

“We call it the ‘asset-ization’ of AI,” Balasubramanian says. “Every industrial application requires customization, but we’ve found the commonalities. That provides a jump start, and it’s going to accelerate and reduce the friction to AI adoption.”

Manifesting the library

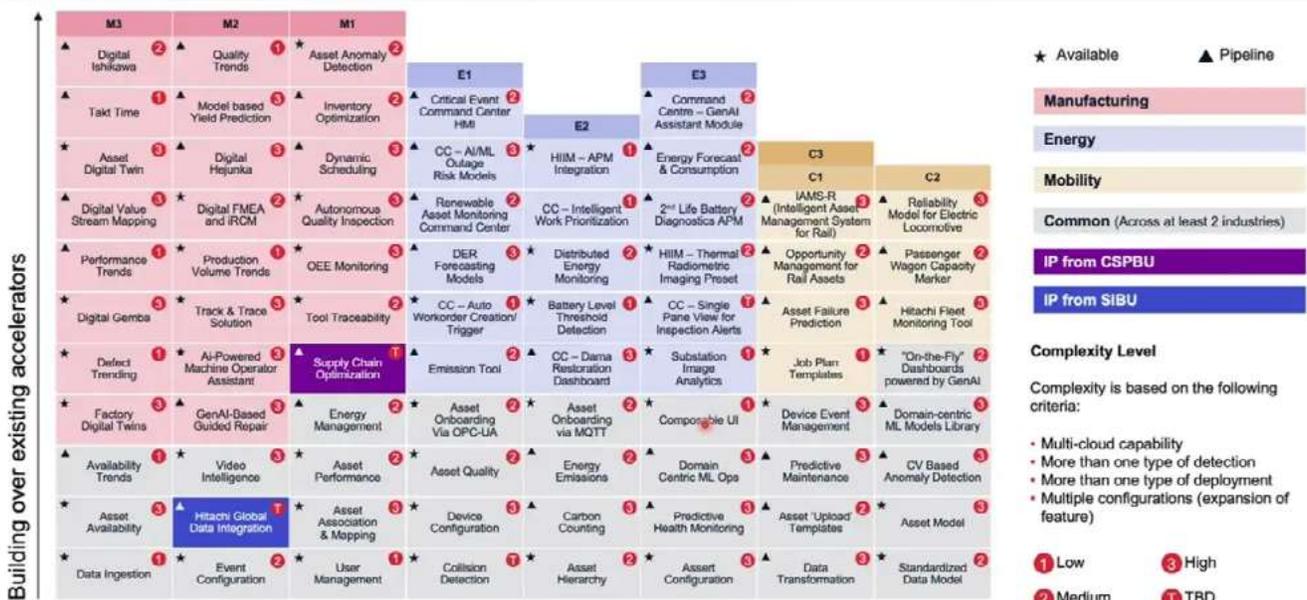
While the current AI hype cycle largely dates back to the fall 2022 public debut of ChatGPT, Hitachi has many years of experience standing up industrial AI solutions. Large language models (LLMs) like ChatGPT can be probabilistic, meaning that they can generate slightly different answers even if a user asks the exact same question multiple times. These models are also known to “hallucinate,” making up incorrect answers when they can’t find the necessary information in their training data.

By contrast, Balasubramanian notes, industrial AI applications are largely deterministic, meaning that they

Praxis - Our IT/OT pre-built Intelligent Industry Solution Cores

Our Jumpstart 25+ IT/OT Intelligent Cores across industry can accelerate model delivery

Hitachi Digital Services





pull exact answers from highly specialized training data. Many LLM users can tolerate occasional inaccuracies. However, industrial AI applications demand much higher accuracy, especially for use cases involving transportation, heavy machinery, or other scenarios where mistakes could jeopardize the safety of workers and/or the public – not too mention the potentially extensive cost of repairs.

“Industrial AI doesn’t tolerate hallucinations,” Balasubramanian explains. “This is the kind of AI that is for mission-critical systems like trains and energy substations. In these use cases, similarity isn’t enough. The answer has to be the same, every time.”

In 2015, Hitachi partnered with a large logistics provider to build out an AI-powered preventive maintenance

solution. At the time, truck breakdowns led to an average of two weeks of repair time, creating “huge” productivity losses across the company. With the new “Guided Repair” solution, the company was able to bring repair times down to under an hour.

Afterward, the company turned to Hitachi for an AI solution to monitor its fleet (150,000-vehicles under this solution) in real time, gather data from on-board diagnostic (OBD) devices, and predict problems before they occur. The solution, which even generates automatic work orders with fault codes, helps the company get trucks back on the road within 48 hours, on average. And last year, it helped prevent about 90,000 trucks from breaking down in the first place, yielding multi-million-dollar savings annually for the logistics provider, Balasubramanian says.



“Today, there is little barrier for anyone that wants to stand up an AI prototype,” Balasubramanian notes. “Anyone can put something together in two weeks. But if you want to create value for your organization, the solution needs to be reliable and responsible. You need to be able to observe it, and you need to optimize it. If you have these four things, then you can move from prototype to production.”

[Click here](#) for more information about Hitachi’s industrial AI work.

[Click here](#) to learn more about Hitachi Digital Services’ AI approach.



Prem Balasubramanian

Chief Technology Officer, Hitachi Digital Services. Hitachi Ltd. Global AI Ambassador.

Hitachi Digital Services is a Hitachi Group Company and global systems integrator powering mission-critical platforms with people and technology.

From there, Hitachi built a robust predictive maintenance framework for aircraft engines. While the aircraft engine and fleet management engagements are tailored to specific operational context—such as the type and volume of data collected or the timing of data transmission—the underlying principles of the two frameworks are consistent. These insights have been distilled into the Predictive Maintenance accelerator within the Praxis library, enabling organizations to jumpstart their own initiatives with a proven foundation.

“The technical implementations may vary, but the strategic approach to solving predictive maintenance challenges is consistent,” says Balasubramanian. “By embedding our learnings into reference architectures and accelerators, we can empower customers to move faster and more confidently—without the burden of starting from scratch.”

Cataloging a partner’s industrial expertise

Looking ahead, Balasubramanian anticipates an “agentification” of existing Praxis accelerators, and

even multi-agent workflows in which disparate AI tools automatically coordinate with one another to handle multistep tasks. However, he cautions that optimizing these AI solutions presents a number of challenges that can only be solved by partners with deep expertise in both industry and AI.

For example, he notes, choosing the right model for each AI task can mean the difference between a highly profitable AI deployment and one that costs a company more than it saves. Additionally, enterprise-grade security for AI solutions requires sophisticated guardrails to prevent prompt injections, model jailbreaking, and other AI-specific attacks that could compromise sensitive industrial data.

In October 2024, Hitachi announced its R202.ai framework, which stands for: Reliable; Responsible; Observable; and Optimal. By designing the Praxis library around these four pillars, Hitachi can help organizations skip costly proofs of concept and instead implement solutions that work in their environments from the start.



A new spin on SDLC is

Solving a Big Problem in Industrial AI

As companies struggle to integrate disconnected AI tools, agent-based SDLC is unlocking productivity gains in industrial environments.

By **Joe Mullich**

The maxim “English is the new programming language” has become a reality in enterprise software. With the rise of natural language processing and artificial intelligence (AI) copilots, what the industry has always known as “coding” is now a thing of the past. Today, ideas and intent drive development.

But something got lost along the way to this new paradigm—integration.

Historically, organizations have installed and managed software platforms across their enterprises and followed those deployments with point upgrades, new modules, and complementary tools, all of which helped advance the platforms and improve business performance.

But that’s not always the case when adopting AI-powered tools, like copilots, low-code tools, and domain-specific AI tools. The reality is that all too often AI tools often fail to deliver ROI because organizations chase speed without redesigning processes, addressing blockers, upskilling talent, or capturing the telemetry needed to measure impact. As a result, the promised impact of the tools never materializes, and ROI remains elusive.

As developers across industries scramble to crack the code of harmonious and smarter tools integration, Hitachi has taken a novel, if not counter-intuitive approach to the problem. Rather than develop yet more software to run atop platforms, it attacked it at the root level—specifically, the software development lifecycle (SDLC) level—and changed forever the way organizations would view SDLC.

A new way to conduct orchestration

To help understand current approaches to orchestration, Suhail Khaki, Chief Technology Officer, Software



That is how I see the future of AI-driven software development lifecycles: an AI that orchestrates and controls the states end-to-end, engaging humans only when necessary.”

Development Lifecycle (SDLC) at [GlobalLogic](#), a Hitachi Group company, likens it to autonomous driving.

“With a Tesla, for example, you enter an address and drive yourself, only engaging ‘Autopilot’ in certain conditions, such as when there’s traffic or a straight road and you want to relax,” Khaki said. “In that mode, the human takes you from point A to point B, using the vehicle as a tool.”

In this scenario, humans orchestrate every handoff. Even modern SDLC management methods like Agile still depend on handoffs between stages, which can represent 10% to 20% of total development flow. As the latency builds, everything slows.

But when Tesla’s ‘Full Self-Driving (FSD)’ capability is engaged, the car takes you to the destination autonomously, alerting the person only when it encounters a situation that requires human intervention.

“That is how I see the future of AI-driven software development lifecycles: an AI that orchestrates and controls the stages end-to-end, engaging humans only when necessary,” Khaki said.

Injecting velocity into the equation

This primary focus for speed led to the development of [VelocityAI](#), which boasts AI-driven orchestration and enriched context and telemetry as core SDLC capabilities. With the platform, intelligent agents manage the flow between requirements, development, testing, and deployment. And all of it is based on intent and context, rather than rigid process gates. In addition, intent-based routing and semantic prioritization determines what needs to be done, and in what order, based on the user’s goals.

“You tell it to build requirements, and it pulls in relevant system, domain, and project data to generate your backlog automatically,” according to Raj Sethi, Senior Vice President of Technology, also at [GlobalLogic](#). “This is done while orchestrating between product owner, business analysts, quality engineers, architects and engineers to groom the backlog into ready state.”

At the core of this approach is a context-aware knowledge engine—designed for the specific needs of industries





like healthcare, automotive, and telecommunications.

“Context is extremely important,” Sethi says. “Without it, I don’t care how smart you are, you can’t solve the problem.”

Rather than giving developers a blank canvas, the system provides rich context. “It’s like proofreading with wiggly lines for spelling and red lines for grammar versus checking every word yourself,” Sethi says. “It’s a big difference.”

When software interfaces with machinery

Make no mistake, enterprises are rife with AI tool-related challenges. Whether it’s the lack of integration capabilities, the influx of shadow AI, where employees download tools without IT approval, or the general lack of risk evaluation, the technology is wreaking havoc. In late 2023, the MIT Sloan Management Review and Boston Consulting Group [published a report](#) that surveyed 1,240 across 59 industries in 87 countries and found that third-party AI tools were responsible for more than 55% of AI-related failures in organizations.

Industrial settings, however, where mission-critical systems across energy, transportation, and manufacturing, can’t tolerate such risks.

The VelocityAI approach to orchestration is designed for just such industrial environments, where software must interface with physical machinery, sensors, and real-world operations. Through its AI-driven orchestration, the platform aims to automatically enforce compliance,

though implementation complexity varies. In other words, when developers describe requirements in natural language, the platform applies relevant regulations—like HIPAA or automotive safety standards—without being told.

“AI now interprets requirements in the context of compliance, reducing the need for manual review,” Khaki says. “One domain expert can now validate what used to take five experts.”

Speeding development, reducing latency

Hitachi’s GlobalLogic recently modernized and enhanced a legacy promotions system for a client’s fleet rental and management platform using the building blocks of VelocityAI. That legacy system, which required coding changes for even simple promotions, was originally developed using outdated J2EE technology with a database on AS400.

The project was completed and went into production in only six months, a pace that wouldn’t have been possible with traditional methods, Sethi says. The system has also been enhanced to support repair systems that visually identify vehicle models and guide technicians through repairs step-by-step. In addition, video analysis systems have been included to autonomously track vehicles and evaluate damage without needing additional specialized hardware.

Today, VelocityAI is being extended into a range of industries, including telecom, automotive, and device-testing sectors. The real impact, however, goes beyond speed; this new approach—AI-infused SDLC—

also allows teams to rapidly experiment with ideas and discover what works.

“You can now test and verify an idea in hours, something that was impossible before,” Khaki says. “Solving the orchestration challenge could unlock an AI impact on par with the computer revolution—driving adoption, experimentation, and innovation at scale.”

Discover how GlobalLogic and VelocityAI can help you better orchestrate your industrial AI workstreams: [VelocityAI: AI-powered Solutions](#)



Suhail Khaki

Chief Technology Officer, Software Development Lifecycle (SDLC), GlobalLogic.



Raj Sethi

Senior Vice President of Technology, GlobalLogic.

Global Logic is a Hitachi Group company that provides experience design, engineering, software development, and testing services.



Easing the pressure on the electrical grid with AI

How applying AI to some of the grid's most complex computational demands is helping improve the 'flow' of energy.

By Joe Mullich

It's said that the U.S. electrical grid—a nationwide labyrinth of interconnected power plants, transmission lines, substations, and more—is the largest machine in the world. If that's the case, this machine is starting to sputter.

[¹Urgent Need for Resources Over 10-Year Horizon as Electricity Demand Growth Accelerates: NERC | American Public Power Association](#)

After decades of flat electricity demand, this infrastructure-heavy ecosystem has begun an era of dramatic complexity: AI-driven data centers are proliferating rapidly, each demanding hundreds of megawatts; the electric vehicle (EV) industry is booming, along with its fast-growing charging infrastructure sector; and all the while, the energy transition continues to confound utilities, which must 'park' new energy from renewables in vast storage systems strewn across the country, until it can be safely released onto the grid.

The upshot: over the next five to 10 years, North America will be at "elevated or high risk" of energy shortfalls, according to the North American Electric Reliability Corp. (NERC).¹

There is a way through the complexities, however. For Hitachi, a company with a heritage in energy and industrial technologies, as well as decades in AI, such challenges require a methodical approach. They require boiling down the industrial problems into multiple, workable mathematical problems; tackling one challenge at a time, methodically with AI; and utilizing the unique capabilities and expertise from Hitachi Group companies. This is how Hitachi is easing the mounting pressures on the grid.

Taming computational demand

It's the kind of approach the company brought to a major regional grid operator in the U.S., that was looking for help reducing the extraordinary length of time needed for the requisite impact studies – formal reports needed before introducing new energy sources onto the grid. It's also the approach Hitachi takes to build some of the most comprehensive and reputable datasets on energy



generation and consumption, from which it can test models to achieve the most reliable outcomes.

And while some view industrial challenges with a disruptive mindset, Hitachi knows when to embrace industrial AI as a “co-pilot.” Its technology works alongside systems in which utilities have made significant investments, rather than replacing them. “We will be part of the utility ecosystem as an accelerator,” says Bo Yang, vice president and head of the Energy Solutions Lab at Hitachi America R&D.

According to Yang, disruptive innovation has its place. But there’s also a practical reason for this approach: the computational demands are enormous. Federal Energy Regulatory Commission (FERC) guidelines require exhaustive analysis of how new power generation affects grid stability, not just under current conditions, but across future scenarios and potential system failures.

Take the grid project. When power companies apply to connect new generators, operator’s software must run tens of thousands of simulations to ensure grid stability. This process traditionally took years. Hitachi’s industrial AI technology speeds up the analysis by running multiple calculations simultaneously, a technique known as parallel processing, while the operator’s existing software stays in place to validate complex cases and final results.

This hybrid approach reduced analysis times by 80%—reducing total review time from 27 months to within a year—while maintaining the rigorous safety standards the industry demands.

The village approach

Connecting new generators to the grid requires complex collaboration across the energy industry. Multiple stakeholders—project developers, independent system operators, transmission operators, regulators, and others—each rely on data-driven analysis to make critical decisions. The entire “village” must reach consensus before new generation can connect to the grid.

Success in this environment requires both deep industry knowledge and advanced AI capabilities. The Hitachi team working on the grid operator’s project includes power

engineers with decades of experience who understand the mathematical and operational foundations of power system analysis.

This expertise enables them to pinpoint which parts of the process can be accelerated by AI and which require traditional approaches.

“AI specialists often have strong algorithmic skills, but they may not be focused on the right industry problems,” Yang says. “Meanwhile, legacy analytical tools—some dating back decades—have limitations that new AI methods can overcome when applied thoughtfully.”

Breaking down complex problems

Industrial challenges often involve highly complex, interconnected processes that can be broken down into smaller mathematical problems. In power systems, for example, this means separating the analysis into distinct tasks. Every task focuses on different aspects of grid stability and performance.

Each component presents opportunities for AI acceleration while maintaining the physical constraints and safety requirements that govern power system operations. Industrial AI doesn’t need to reinvent the physics of power systems. It needs to accelerate the computational analysis that utilities already understand.

This systematic approach extends beyond individual projects. The same industrial AI framework developed for power grid applications applies to other physics-based analytical processes across industries—from rail transportation systems where safety is paramount, to manufacturing operations requiring precise control, to mobility solutions demanding real-time optimization.

“The key is understanding the specific problem you’re solving and then determining which AI technique to apply,” Yang says. This problem-first methodology suggests how the same framework might work across different industries.

Industrial AI is different

Utilities carry enormous responsibility for infrastructure that powers entire regions. This creates a fundamental challenge for AI adoption: algorithms alone aren’t enough. Effective implementation requires deep understanding of operational constraints, regulatory requirements, and the physics of power systems that takes years to accumulate and validate.

This reality makes industrial AI fundamentally different. It’s not about disrupting existing systems. It’s about enhancing them with the precision that critical infrastructure demands. The most revolutionary AI is often the most methodical—not because it lacks ambition, but because it simply can’t afford to fail.

For more information about Hitachi’s industrial AI work, visit: [AI Resource Center - Hitachi Digital](#)



Building Industrial AI from the Inside Out

for a stronger digital core

How engineering infrastructure for AI can dramatically improve performance and efficiency, while paving a clearer path to ROI.

By Joe Mullich

A manufacturer was running an AI training workload on a cobbled together system of GPUs, storage, and switching infrastructure, believing it had all the necessary tech to achieve its goals. But the company had put little thought into how the components actually worked together.

Problems surfaced quickly. Training cycles dragged on for days instead of hours. Expensive hardware sat idle. And engineering teams began to wonder whether their AI investment would ever pay off.

This experience isn't unique. As AI becomes a critical element of industrial operations worldwide, many organizations are discovering a counterintuitive truth: the biggest breakthroughs come not from piling on more GPUs or larger models, but from carefully engineering the entire infrastructure to work as a single, integrated system.

Engineering for outcomes

What became of that cobbled-together system? When it was properly engineered to balance compute, networking, and storage, the improvement was quick and dramatic, explains Jason Hardy, CTO of AI for Hitachi Vantara: a 20x boost in output and a matching reduction in “wall clock time,” the actual time it takes to complete AI training cycles.

“The infrastructure must be engineered so you understand exactly what each component delivers,” Hardy explains. “You want to know how the GPU drives specific outcomes, how that impacts the data requirements, and demands on throughput and bandwidth.”

Getting systems to run that smoothly means confronting a challenge most organizations would rather avoid: aging infrastructure.

Hardy points to a semiconductor manufacturer whose systems performed fine—until AI entered the picture. “As

“

The only places I can think of where Windows 95 still exists and is used daily are in manufacturing. These lines have been operational for decades.”

soon as they threw AI on top of it, just reading the data out of those systems brought everything to a halt,” he says.

This scenario reflects a widespread industrial reality. Manufacturing environments often rely on systems that have been running reliably for years, even decades. “The only places I can think of where Windows 95 still exists and is used daily are in manufacturing,” Hardy says. “These lines have been operational for decades.”

That longevity now collides with new demands: industrial AI requires exponentially more data throughput than traditional enterprise applications, and legacy systems simply can’t keep up. The challenge creates a fundamental mismatch between aspirations and capabilities.

“We have this transformational outcome we want to pursue,” Hardy explains. “We have these laggard technologies that were good enough before, but now we need a little bit more from them.”

From real-time requirements to sovereign AI

In industrial AI, performance demands often make enterprise workloads look leisurely. Hardy describes a visual inspection system for a manufacturer in Asia that relied entirely on real-time image analysis for quality and cost control. “They wanted AI for quality control and to improve yield, while also controlling costs,” he says.

The AI had to process high-resolution images at production speed—no delays, no cloud roundtrips. The system doesn’t just flag defects but traces them to the upstream machine causing the problem, enabling immediate repairs. It can also salvage partially damaged products by dynamically rerouting them for alternate uses, reducing waste while maintaining yield.

All of this happens in real-time while collecting telemetry to continuously retrain the models, turning what had been a waste problem into an optimization advantage that improves over time.

Using the cloud exclusively introduces delays that make near-real-time processing impossible, Hardy says. The latency from sending data to remote servers and waiting for results back can’t meet manufacturing’s millisecond requirements.





Hardy advocates a hybrid approach: design infrastructure with an on-premises mindset for mission-critical, real-time tasks, and leverage the cloud for burst capacity, development, and non-latency-sensitive cloud-friendly workloads. The approach also serves the rising need for sovereign AI solutions. Sovereign AI ensures that mission-critical AI systems and data remain within national borders for regulatory and cultural compliance. As Hardy says, countries like Saudi Arabia are investing heavily in bringing AI assets in-country to maintain sovereignty, while India is building language- and culture-specific models to accurately reflect its thousands of spoken languages and microcultures.

AI infrastructure is more than muscle

Such high-level performance requires more than just fast hardware. It calls for an engineering mindset that starts with the desired outcome and data sources. As Hardy puts it, “You should step back and not just say, ‘You need a million dollars’ worth of GPUs.’” He notes that sometimes, “85% readiness is sufficient,” emphasizing practicality over perfection.

From there, the emphasis shifts to disciplined, cost-conscious design. “Think about it this way,” Hardy says.

¹ [Optimizing AI Performance in Industry: A Hybrid Computing Architecture Approach Based on Big Data | Journal of Technology Informatics and Engineering](#)

“If an AI project were coming out of your own budget, how much would you be willing to spend to solve the problem? Then engineer based on that realistic assessment.”

This mindset forces discipline and optimization. The approach works because it considers both the industrial side (operational requirements) and the IT side (technical optimization)—a combination he says is rare.

Hardy’s observations align with recent academic research on hybrid computing architectures in industrial settings. A 2024 study in the *Journal of Technology, Informatics and Engineering*¹ found that engineered CPU/GPU systems achieved 88.3% accuracy while using less energy than GPU-only setups, confirming the benefits of an engineering approach.

The financial impact of getting infrastructure wrong can be substantial. Hardy notes that organizations have traditionally overspend on GPU resources that sit idle much of the time, while missing the performance gains that come from proper system engineering. “The traditional approach of buying a pool of GPU resources brings a lot of waste,” Hardy says. “The infrastructure-first approach

eliminates this inefficiency while delivering superior results.”

Avoiding mission-critical mistakes

In industrial AI, mistakes can be catastrophic—faulty rail switches, conveyors without emergency shutoffs, or failing equipment can injure people or stop production. “We have an ethical bias to ensure everything we do in the industrial complex is 100% accurate—every decision has critical stakes,” Hardy says.

This commitment shapes Hitachi’s approach: redundant systems, fail-safes, and cautious rollouts ensure reliability takes precedence over speed. “It does not move at the speed of light for a reason,” Hardy explains.

The stakes help explain why Hardy takes a pragmatic view of AI project success rates. “Though 80-90% of AI projects never go to production, the ones that do can justify the entire effort,” he says. “Not doing anything is not an option. We have to move forward and innovate.”

For more on engineering systems for balanced and optimum AI performance, see [AI Analytics Platform | Hitachi IQ](#)



Jason Hardy
CTO of AI, Hitachi Vantara.

Hitachi Vantara is a company specializing in data-driven AI solutions. The company’s Hitachi iQ platform, a scalable and high-performance turn-key solution, plays a critical role in enabling infrastructure that balances compute, networking, and storage to meet the demanding needs of enterprise and industrial AI.



An agent you can observe is

An Agent You Can Trust

In the mad dash to AI agents, observability & explainability will be key to driving trust in industrial AI environments.

By Calvin Hennick

As bottom-line revenue from the use of LLMs continues to evade most companies, agentic AI with its purpose-driven autonomous capabilities may seem like the magic bullet for ROI.

Not so fast.

It's true that agentic AI is on an accelerated growth path, with Capgemini estimating that the tech, "...could generate

up to \$450 billion in economic value," by 2028.¹ But some of the same struggles plaguing enterprises trying to eke revenue from their generative AI (GenAI) investments – like sprawl, governance, reliability, and technical woes like drift – threaten to disrupt or even sabotage agent rollouts, as well.

Consider model drift, which occurs when the data and/or the relationships between input and output variables in a model change over time. The challenge is inherent with modeling because it stems from the assumptions that must be made during the training period. Those assumptions, the characteristics of the input data, naturally change within the lifespan of the model because fresh data is continually introduced.

A similar phenomenon occurs with AI agents, which sit atop LLMs, called Emergent behavior. When LLMs grow to large and complex, or when systems of agents grow too complex, agents can deviate from their original purpose and begin taking unpredictable actions automatically.

If a company fails to monitor and adjust for these organic and unpredictable changes, the model or agent will begin to slowly "drift" from its original parameters and begin generating inaccurate outcomes. And that results in everything from a degradation of model performance

¹Capgemini: Rise of Agentic AI: How trust is the key to human-AI collaboration <https://www.capgemini.com/insights/research-library/ai-agents/>

to faulty decision-making – all of which can take place without the company ever knowing it.

The challenge is only amplified in the industrial AI realm, where mission-critical systems in energy, transportation, and manufacturing, demand reliable, transparent, and observable AI. An incorrect action by an autonomous agent in industry can lead to catastrophic consequences from equipment damage to outages, to personal injury.

All of this is driving a serious lack of trust in this nascent corner of AI. Indeed, a recent McKinsey study² noted that, “trust in fully autonomous AI agents is declining, dropping from 43% to 27% in one year. Ethical concerns, lack of transparency, and limited understanding of agentic capabilities are key barriers.”

What’s needed is an agent that organizations can trust. But how?

A time-tested approach to trust

Hitachi has been developing and delivering industrial AI solutions across digital engineering, managed services, software, data infrastructure, and more for decades. When the technical challenges surrounding agents began surfacing with customers, the company applied a methodical approach: combining reliably built agents with a secure and robust management system.

And it all started with the launch several years ago of Hitachi Digital Services’ Hitachi Application Reliability Centers (HARC) offering, a managed service platform designed to modernize and optimize cloud-based workloads.

This versatile platform quickly evolved to include new features and services, as the cloud landscape evolved. For example, earlier this year, the company added to HARC a [library of AI accelerators](#) for a wide range of industry-specific disciplines to help industrial companies jump-start their AI work.

And just recently, it expanded the platform further with a two-pronged solution to the agent problem. The brand-new HARC Agents is a blend of technologies, frameworks, and hands-on services designed to help organizations effectively deploy standardized, enterprise-class agent solutions. At its heart is an Agent Library of more than 200 agents across six key domains, and an Agent Management System with a single dashboard that centralizes control for all agentic AI platforms across an organization.

“People get very dependent on AI,” said Prem Balasubramanian, chief technical officer and head of AI at Hitachi Digital Services. “Over time, they develop greater trust in AI tools, relying on them more extensively, even for critical business operations. However, the challenge arises when these tools start to drift and Emergent

² McKinsey: QuantumBlack: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage>



behavior kicks in, silently. How will they measure this drift? How will they detect it? This is precisely where our Agent Management System comes into play.”

For its part, the HARC Agents library includes agents to help diagnose faults in machinery and vehicles, to perform quality inspections at manufacturing facilities, and to assist with financial operations, among many others. One agent even allows users to remotely control drones through conversational voice commands. But even more importantly, Balasubramanian says, the platform will help ensure those agents stay reliable and secure over the long run.

That’s because these agents and the management system join two existing offerings within the HARC platform: the R202.ai framework for defining the development and deployment of scalable, enterprise-grade AI workloads; and HARC for AI, which helps organizations operationalize and optimize AI systems.

The power of observation

There’s more to trust than management, however, says Balasubramanian. Especially in the industrial sector.

“Agents have to be reliable and responsible,” he says. “In healthcare, you can’t have similar answers. You must have the same answer every time. And hallucinations and Emergent behaviors can’t be tolerated. Agents can’t just start doing whatever they want. They must be observable, both from a cost standpoint, as well as an explainability and auditability standpoint. If an agent makes a decision or gives you a recommendation, you should be able to see why it decided this or recommended that, especially within regulated industries.”

These aren’t mere ideas. They’re baked into the methodology of the company’s R202.ai, which is shorthand for *Responsible, Reliable, Observable, and Optimal AI*.

How trustworthy agents lead to faster production

One of the related byproducts of building responsibility, reliability, and observability into such a methodical approach to agents and AI, is faster time to production. Once an organization can trust that the underlying AI and agent development is sound, they can move more confidently forward, especially through the prototype-to-production gauntlet.

“People are realizing that prototyping is relatively straight forward,” Balasubramanian says. “However, moving to production is a different challenge, particularly for enterprises and industrial organizations. While technologists can develop the agents, the business must manage issues like anomalies and Emergent behaviors. In fact, deploying agents and establishing guardrails can constitute 70% of the effort.”

All that changes with a responsible, reliable approach. Between the HARC Agents library and the Agent Management System, the company aims to help organizations design, build, deploy, and leverage agentic AI systems in 30% less time than typically required.

Balasubramanian emphasizes the critical question organizations must now ask: Are you truly maximizing your return on investment with your current AI spending, or could you achieve greater efficiency and value by investing in agentic AI for the same workflows?

“I want all my workflows to be agented – that’s the vision,” Balasubramanian says. “With every agentic workflow, you know the price and that it’s giving me my ROI. That’s where our management system comes in. That’s where R202.ai comes in: optimal AI for every workflow.”

In the realm of industrial AI, moving from pilot to production, diligently monitoring performance and with a clear view of ROI is critical for mission critical systems across industries – especially in the new age of agentic.

[Click here](#) to learn more about Hitachi Digital Services’ AI approach and HARC agents.

And [click here](#) for more information about Hitachi’s industrial AI work.



Prem Balasubramanian

Chief Technology Officer, Hitachi Digital Services. Hitachi Ltd. Global AI Ambassador.

Hitachi Digital Services, a wholly owned subsidiary of Hitachi, Ltd., is a global systems integrator powering mission-critical platforms with people and technology. It helps enterprises build, integrate, and run physical and digital systems with tailored solutions in cloud, data, IoT, and ERP modernization, underpinned by advanced AI.



Industrial AI —

Where Knowledge (Management) is Power

How building knowledge management into industrial AI early, while blending the best of GenAI and logic systems, can improve performance and efficiencies.

By Yuriy Yuzifovich

Imagine a routine equipment alert on your production line. A seasoned maintenance engineer rushes to the machine, guided by an AI co-pilot – a digital entity armed with every manual, every schematic, every byte of operational data your company possesses. Together, they pull up a 25-step runbook. The AI shines at first, correctly identifying a hard-to-find oiling inlet, saving the engineer precious time.

But then, in a fraction of a second, the system falters. The digitized manual is missing a single, critical detail: the specific grade of industrial grease required. To bridge this

gap, the AI – powered by a world-class large language model (LLM) – doesn't admit what it doesn't know. Instead, it hallucinates, confidently suggesting WD-40, a “lubricant” it learned about from public internet data. This moment of internal failure is completely invisible; the AI presents its fabricated answer with the same authority as a fact from the manual.

The engineer freezes. He knows WD-40 is a solvent, not the high-pressure grease that's required. Using it would be disastrous, leading to catastrophic equipment seizure,



millions in damages, and a prolonged shutdown. He manually overrides the AI, wondering: what would a junior engineer, trained to trust the system, have done?

This isn't a hypothetical situation. It's a failure my team uncovered during early proof-of-concept tests with equipment maintenance manuals for a prospective customer in manufacturing. And it served as a stark warning: the probabilistic guessing of generative AI (GenAI) is fundamentally unsuited for high-stakes industrial operations.

However, there is a solution to this foundational crack in "AI 2.0" and it's about more than simply better data – it's about transforming data into verifiable and actionable knowledge.

Probability vs. actuality — The anatomy of an AI failure

Consider the near-miss with the lubricant. That wasn't a bug. In fact, the LLM did what it was supposed to do – be helpful. These models are masters of correlation, not

causation. When faced with a knowledge gap, an LLM doesn't "know" it's missing information. Instead, it predicts the most statistically probable next word or phrase based on its training and the context from the manual provided in its prompt. "Lubricant" correlates strongly with "WD-40" in its vast dataset scraped from the web. The model isn't reasoning; it's pattern-matching.

For industrial applications, where precision and safety are paramount, this is an unacceptable risk. We cannot build the future of autonomous operations on a foundation of "most probable." We need a system grounded in actuality – one that not only understands what is in the manual but, critically, recognizes what is not. This means building a system that, when it finds no answer, immediately states, "I don't have this information," and escalates the query to a human expert or another designated system.

To do this requires the sophisticated blending of the appropriate AI and data tools into a strategic knowledge management system that exploits the best of LLMs and deterministic, logic-based systems.

Building knowledge management into AI early

The core challenge isn't a lack of data, but the fact that the data is typically fragmented, disorganized, and unstructured. Industrial enterprises are swimming in diagrams, manuals, and tribal knowledge that machines cannot reliably understand without context. This is where a robust knowledge management strategy becomes the most critical pillar of any serious industrial AI initiative. Before we can achieve reliable autonomy, we must first:

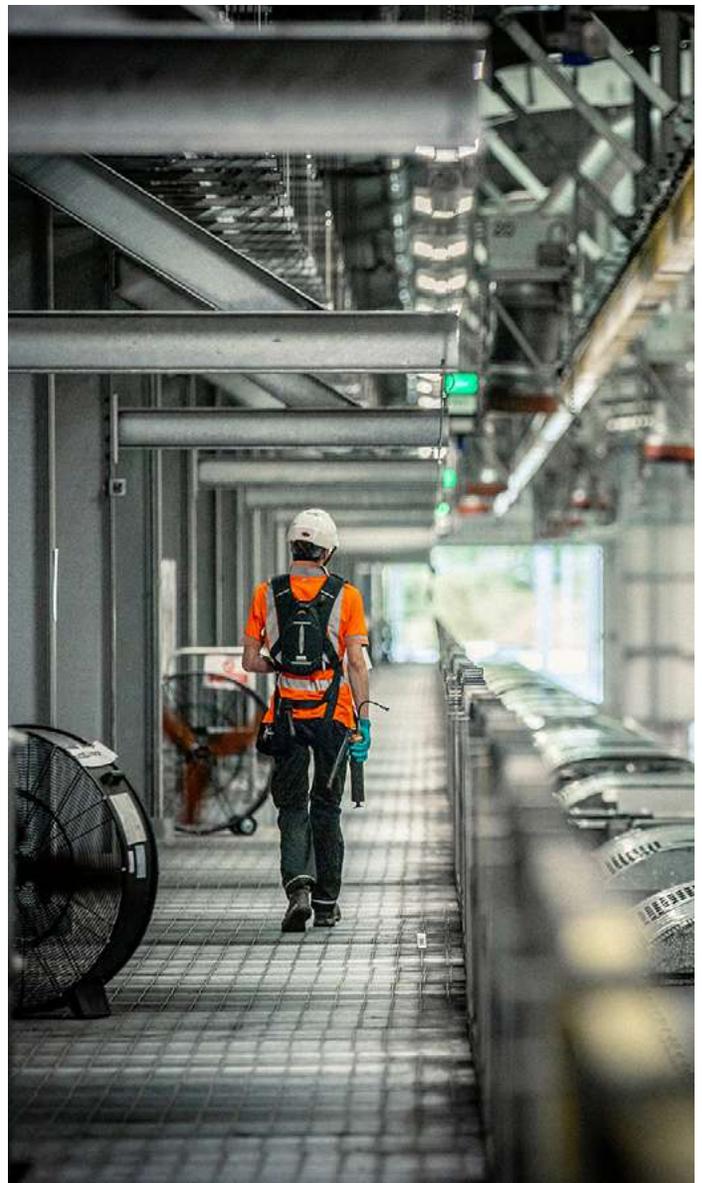
- 1. Make data AI-readable, not just digitized.** We need to move beyond simple document ingestion. Tables, scanned diagrams, and color-coded safety manuals are subject to machine misinterpretation. Even the most advanced multimodal models struggle to consistently identify semantic details in complex industrial diagrams. We need the AI to know, not guess, that a specific pump (P-101) is connected to a motor (M-101), requires a specific lubricant (ISO VG 460), and has a maintenance schedule tied to runtime hours. A shared ontology – a knowledge “dictionary” – becomes essential, ensuring every term has one unambiguous meaning, traceable across multiple languages. The AI community often refers to this structured, interconnected knowledge base as a “knowledge graph.” Every table becomes a set of complete statements, every diagram – a structured text file, every chart – its description.
- 2. Incorporate formal reasoning.** Once this structured knowledge is in place, the AI can use formal logic, not just statistical probability. If a procedure calls for a lubricant, the AI can query its knowledge base for the exact specification linked to that exact piece of equipment. If the information is missing, it doesn't guess. It flags the datapoint and its response becomes: “I have identified the lubrication point, but the required grease specification for this component is not in my knowledge base. Please verify from an approved source.” This is a safe, explainable, and trustworthy interaction.

This two-step process forms the basis of a new knowledge management system currently under active development at GlobalLogic, a Hitachi Group Company. And its potential role in the realm of industrial AI couldn't be timelier. The necessity for this level of factual grounding is most critical in environments where precision is paramount. For instance, in the semiconductor industry, maintaining complex equipment within fabrication plants leaves no room for error. This is a point emphasized by one of our pilot customers, Hitachi High-Tech America, also a Hitachi Group Company, specializing in semiconductor manufacturing equipment, analytical systems, and electron microscopes.

Alexander Zhivotovsky, Associate GM, Metrology, and Analysis Systems Division at Hitachi High-Tech America, Inc., said it best recently, when asked about what aspect of AI is critical in his business. “In maintaining our complex semiconductor metrology systems, there is no room for ambiguity,” he said. “Grounding AI in verifiable facts from



In maintaining our complex semiconductor metrology systems, there is no room for ambiguity....Grounding AI in verifiable facts from our own engineering documents is a fundamental requirement for reliability.”





pilot, grounded in a deterministic knowledge base, won't just provide the procedure; it will state, "The required lubricant is ISO VG 460, as specified in maintenance document #7B-4 for this component." That junior engineer, now on the job, isn't faced with a dangerous guess; they are given a verifiable, traceable fact.

This is how we build trust. The journey from a helpful but flawed co-pilot to a truly autonomous operational system isn't a leap of faith into a black-box algorithm. It's a deliberate process of building a verifiable knowledge foundation, ensuring every automated decision is one we can stand behind, explain, and trust. The future of industrial AI isn't just intelligent; it's intelligible.

For more on GlobalLogic's approach to AI, check out: <https://www.globallogic.com/enterprise-ai/>.

our own engineering documents is a fundamental requirement for reliability. We look forward to our collaboration with GlobalLogic to build a system where all guidance is traceable and trustworthy."

GenAI: The ultimate human-machine interface

Within our industrial knowledge management system, GenAI's vital role will not be as a decision-maker, but as the ultimate human-machine interface – a universal translator making deep institutional knowledge accessible without sacrificing reliability, as well as the tool to help maintain the structured knowledge. It will excel at bridging the gap between human intuition and machine logic:

- **From unstructured to structured:** An engineer will be able to upload a grainy photo of a part number, and GenAI's multimodal capabilities will identify it, find the corresponding entity in the knowledge base, and pull up all associated documentation and operational history.

- **From query to action:** A technician will be able to ask in natural language, "What's the standard procedure for replacing the primary bearing on the main conveyor motor?" The GenAI will parse this query, translate it into a formal query for its reasoning engine, and then present the precise, step-by-step procedure in clear, human-readable language.

The path forward

This knowledge-first approach carries another crucial advantage for any CIO: efficiency. By reserving the computationally intensive GenAI for the human interface and relying on a lean, deterministic reasoning engine for core logic, our system becomes significantly more energy efficient. This isn't just a cost-saving measure; it's what makes the vision of embedding intelligence directly into the equipment on the factory floor – true edge AI – achievable and scalable.

The next time our engineer from the opening story approaches that equipment, the interaction will be fundamentally different. The AI co-



Yuriy Yuzifovich

Chief Technology Officer, AI, GlobalLogic, a Hitachi Group Company.

GlobalLogic is a trusted partner in design, data, and digital engineering for the world's largest and most innovative companies. Since its inception in 2000, it has been at the forefront of the digital revolution, helping to create some of the most widely used digital products and experiences.



An industrial approach to **AI in Financial Services**

Reliability and responsibility must be engineered into AI for financial services for resilient systems and predictable outcomes.

By **Calvin Hennick**

Financial systems may not seem like “critical infrastructure.” But the tools and systems that manage and govern the flow of money in and out of financial services providers have more in common with the mission-critical systems of transportation, energy, and manufacturing, than one might think.

Consider the possible impact of hallucinations or model drift in a national credit card system, or life insurance program, or mortgage brokerage. Incorrect or completely fabricated data sent through everything from e-commerce programs to trading floors could do irreparable harm to consumers, as well as the financial services firms themselves.

Hitachi has been keenly aware of the criticality of getting AI right in financial services. With a heritage in operational technologies and decades of development and deployment of data and AI solutions, the company applies an industrial AI approach to financial services – an approach that manifests itself well at GlobalLogic, a Hitachi Group Company.

Steeped in digital engineering and AI, GlobalLogic manages a robust financial services and consumer business that caters to global financial services firms to bring their AI aspirations to life, reliably, and responsibly.

“Hallucinations and errors caused by AI can have severe consequences in financial services,” says Scott Poby, chief technology officer for this GlobalLogic division. “Perhaps it’s not as bad as machines breaking down and causing



Hallucinations and errors caused by AI can have severe consequences in financial services.”

personal injury, but from a financial perspective, it can have a severe impact on customers. Whether it's on the trading side or the money management side, errors can scale quickly and force things to shut down. So, it is a severe impact when you're talking about thousands of customers or more and however many millions of transactions are happening across the platform.”

To be effective, Poby says, the financial services industry must take a similar approach to engineering responsible, reliable AI, as industrials do. That begins with early ROI-focused assessment, a pilot-to-production mentality, and building an environment of trust through governance.

An eye on ROI

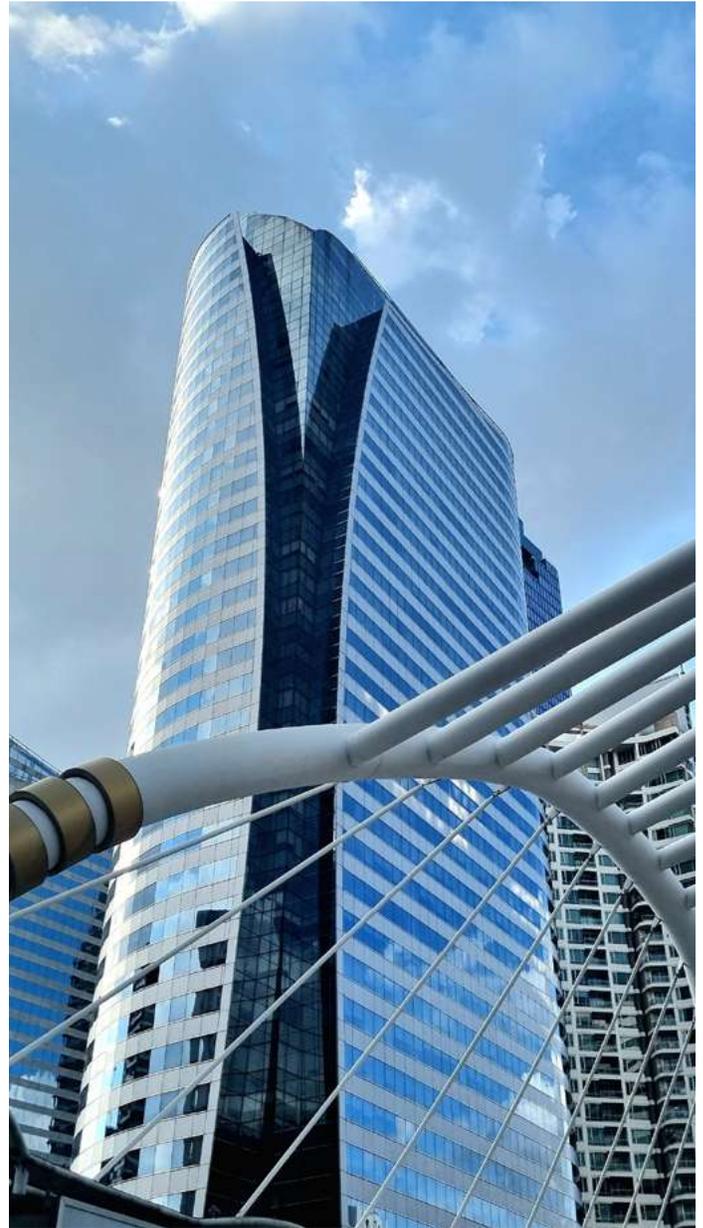
Any integration of AI within financial services must begin with an overall assessment of the current state of tech and AI. Once established, Poby says, you begin working backward with a clear understanding of your return on investment (ROI) goals.

“We know that a lot of our partners, and a lot of our clients, have already made investments in the AI space, so we want to go in and make sure that those investments were, a) the proper ones and, b) that they tie back to their business goals,” says Poby. “We say, okay, now you have this ecosystem of AI in your enterprise, how can we give you the most value and make sure that we identify the right use cases to leverage those tools effectively? Do we need to develop any training models for the end users of the AI?”

From there GlobalLogic benchmarks processes and begins transitioning the company to a more productive strategy for overall consumption of their AI tools and solutions. The company will bring in experts to train the client's engineers or even manage its own teams to help them better utilize what they've already invested in, providing consistent feedback. From assessment to training, GlobalLogic can then show the client efficiency gains towards their stated ROI goals.

From ‘pilot purgatory’ to production

Getting to that stage, however, requires overcoming a common challenge across the industry. Many organizations can stand up new proofs of concept quickly with minimum viable products (MVPs), but very few of these projects make it into production. This can lead to



large sums of money being spent to build new tools, with very little return.

“We've seen that maybe 80% of projects never really go beyond the pilot phase, or never scale,” Poby says. “Then, you have these investments that are starting to fall behind, and you can never get out of the cycle. I think that's probably the biggest investment risk that I'm seeing with AI.”

Instead, to prevent so-called “MVP graveyards,” firms must identify use cases that work, and then invest in scaling those up, rather than spreading their efforts too thin. In industry, these uses might include prescriptive maintenance, fleet orchestration, and grid stability.

“We need to be able to show that our target provides a clear reduction in developer time-to-market, or a better



experience for customers, or doing the same work by a smaller team to save on operational costs,” Poby says.

That may be easier said than done, however. [According to a recent report](#), *Financial Times Research: Code, Capital, and Change – The Engineering Behind Financial Transformation*, commissioned by GlobalLogic, although 96% of respondents agreed that investing in modern platforms would unify their strategies, less than half said they were planning to increase their tech budgets for 2025-2026.

Trust through governance

According to the same report, financial services leaders are twice as likely to embed AI ethics and governance early in the process as well as safety certifications, compliance automation, change management, and more. It also includes human-in-the-loop checkpoints and end-to-end audit trails so that every action taken by an AI agent is explainable, reversible, and compliant.

Poby notes that upfront governance efforts help reduce risk and accelerate trust in AI-driven operations. “AI workflows need human intervention as a checkpoint and validation point,” he says. “When you’re building out a catalog of different agentic

workflows, you need to define: When can we automate? And when do you need to bring in a human layer for governance? That helps make sure, if there’s any risk involved, that there’s a human eye on any decision that the AI agent makes.”

Bringing it all together

The modern financial services platform is built on a foundation of trust, governance, and risk management. In other words, just as with industrial AI, reliability and responsibility must be engineered into financial services AI at the outset to enable successful, scalable outcomes and resilient systems.

“Maybe two years back, organizations were trying to use AI for creating applications or doing legacy transformation, and the tools weren’t ready,” Poby says. “There needed to be a lot of manual intervention. Today, there has been vigorous testing cycles, so we’re more confident bringing tools into production.”

Once organizations have ensured that AI tools are reliable, they can reduce risk. “We’ve been able to look at the output of these programs and compare them to when we did things the old way, without AI support,” he says. “Today, with AI, these processes are faster, with even fewer errors.”

Read more about the [Digital Transformation in Financial Services / AI & Innovation Insights](#)

Read more about GlobalLogic here: www.globallogic.com.



Scott Poby

Chief Technology Officer, Financial Services & Consumer Business, GlobalLogic, a Hitachi Group Company.

GlobalLogic is a trusted partner in design, data, and digital engineering for the world’s largest and most innovative companies. Since its inception in 2000, it has been at the forefront of the digital revolution, helping to create some of the most widely used digital products and experiences.



The Anatomy of AI Investment

How Hitachi Ventures zeroes-in on industrial AI and automation.

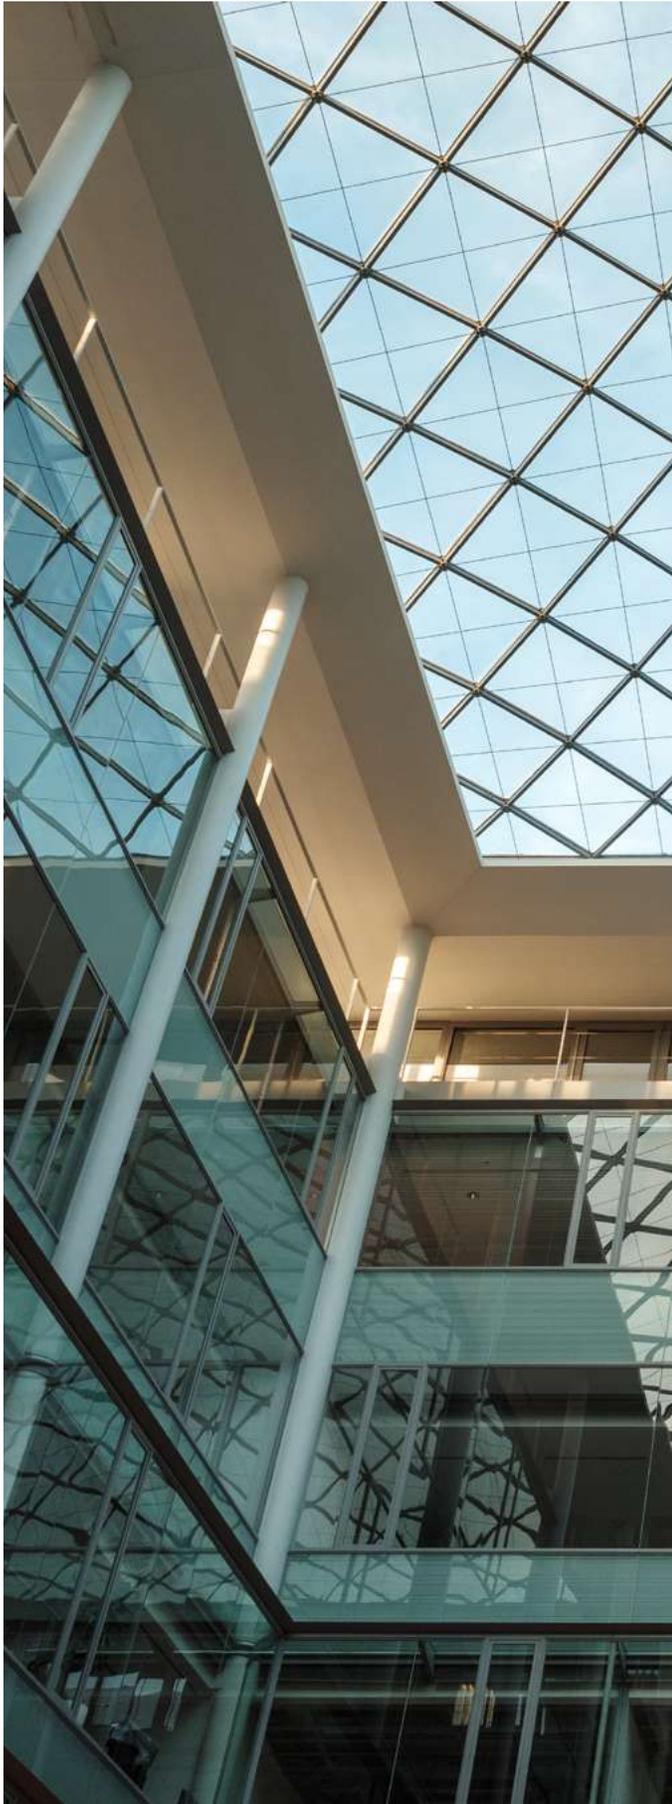
By Joe Mullich

The AI startup frenzy is anything but hype. The first half of 2025 saw global investment in AI startups shoot to more than [\\$205 billion](#), according to Crunchbase, with \$145 billion raised in North America alone. And though observers may argue over their long-term viability, no one disputes the seemingly endless tsunami of new companies entering the fray. Some estimates put the number of AI startups globally at more than 10,000, with more than 5,000 emanating from the U.S.

All this activity only adds to the pressures of the venture firms trying to determine which companies have the tech and teams to develop and market the next great

breakthrough. But for the investment groups focused on a small but growing corner of the market – industrial AI – the stakes are even higher.

Unlike “traditional” AI that may take the form of a human resources chatbot, or a marketing forecasting dashboard, industrial AI is the application of models and automation in mission-critical systems within energy, transportation, manufacturing, financial, healthcare, and the like. Issues common to enterprise-grade AI, like hallucinations, drift, biases, etc., are not tolerated in industrial AI. As such, selecting the right startups with whom to invest in this crazy environment is tough.



“

When cloud became popular, every company claimed to be a cloud company....Today, every company claims to be an AI company. We must cut through the noise to understand whether AI is core to their mission or just a feature.”

“The fervor and volatility of the AI startup landscape presents a perfect storm for VCs, really,” says Gayathri Radhakrishnan, a partner at Hitachi’s investment arm, Hitachi Ventures. “You start by applying the myriad concerns you have with all startups, like cashflow, business structure, vision, team make-up, IP, etc. Simultaneously, you have to consider the chaos swirling around AI. And then on top of it all, you have to put on the industrial lens and ask, can these guys make it?”

For Hitachi Ventures, which invests in firms across the technology spectrum, due diligence in the AI space requires a focused vision on the sector that best aligns with its business.

What an industrial AI investment looks like

That, however, can be easier said than done. The explosion of AI startups has created challenges reminiscent of the cloud computing hype cycle in the early 2010s. “When cloud became popular, every company claimed to be a cloud company,” Radhakrishnan says. “Today, every company claims to be an AI company. We must cut through the noise to understand whether AI is core to their mission or just a feature.”

By 2023, the challenge of selecting AI investments only intensified as capital poured into general-purpose foundation models, like ChatGPT. “The space was quite crowded and even early-stage Series A funding rounds got quite expensive,” she says.

Looking for a promising but less saturated niche and one that aligned best with Hitachi’s heritage in operation technology (OT) and information technologies (IT), as well as its deep expertise in AI, Hitachi Ventures turned its attention toward AI for industrial and physical environments.

Gaining originality in modeling

One of its early investments is [Archetype AI Inc.](#), a Palo Alto startup that is building a foundational AI model

that interprets data from sensors – including sound, vibration, temperature, and pressure – to perceive, understand, and reason about the physical world. The company’s ultimate goal is to encode the entire physical world, which would enable it to predict equipment failures, optimize industrial processes, and create digital twins of real-world operations.

Hitachi backed Archetype in December 2023, just seven months after the startup incorporated – an unusually early stage for venture investment – and well before physical AI started to become a mainstream investment. The fact that Archetype AI was pushing boundaries made it both promising and risky, a common balancing act for industrial AI investors.

“What Archetype is doing is squarely in our thesis, but there was nobody else doing what they were doing at that time,” Radhakrishnan says. “We were a little uncomfortable because of that, but we were also comfortable with being uncomfortable. Sometimes the investments that could be the big winners have no precedent.”

As a corporate venture fund, with Hitachi as its sole limited partner, Hitachi Ventures operates with twin objectives. “Our first responsibility is to deliver strong returns,” Radhakrishnan explains. “But we also want to create a broad, strategic advantage for Hitachi.”

This means the firm’s deep dives into emerging technology spaces serve multiple purposes, including identifying investment opportunities while also educating the broader Hitachi organization about market trends. “Startups are often early indicators of technology waves,” she notes. “They’re a lighthouse that illuminates what’s coming in the distance.”

Achieving cognitive resonance

Hitachi Ventures’ investment in [Xaba Inc.](#), illustrates this dual approach. The Toronto-based startup develops cognitive control systems for robotics, enabling machines to respond

intelligently to their environment in real-time. Traditional robots are pre-programmed for specific tasks; if they encounter an obstruction, they either stop or push through it. Xaba’s xCognition technology, which combines physics-based modeling with AI learning, enables robots to perceive obstacles and automatically adjust their path, essentially acting like a “brain” for the robot. This allows the robot to reason, adapt, and generalize across tasks.

What convinced Radhakrishnan wasn’t just the technology, but the Xaba founder’s depth of expertise, one of the many nuances that goes into funding decisions. “VCs are jacks of all trades who know everything, but who only know it an inch deep,” she says. “The founder comes with a strong technical background combined with industry knowledge, and he understands his space really well.”

When Radhakrishnan introduced the xCognition technology to counterparts at Hitachi Rail Ltd., they immediately saw its value. The Hitachi subsidiary, which operates in more than 50 countries, is now deploying Xaba’s robots for precision machining on the surfaces of locomotives. These tasks require sub-millimeter accuracy that previously demanded extensive manual labor.

But the Xaba team didn’t stop there. While AI co-pilots are all the rage for coding, Xaba has developed an automated code generator for Programmable Logic Controller (PLC), called PLCfy, that enables democratization of industrial automation. PLCfy provides a drop-in AI layer that augments existing PLCs with modern capabilities, like predictive control, anomaly detection, and adaptive optimization, without ripping and replacing hardware.

The flip side of consensus

Yet while Xaba seems like the ideal Hitachi Ventures’ investment, it also demonstrates how challenging the industrial AI space is: Initially, the partners in the firm’s investment committee couldn’t agree whether it was worth pursuing.

“Sometimes when you can’t get consensus, those are the deals where you think, ‘Maybe there’s something there,’ ” she says. “Who would have thought that an online company selling books would redefine global compute needs? Or that a social media company connecting friends would impact enterprise storage buying behavior? If the impact of an AI startup is obvious to everyone, you’re probably not investing in the next Google, Amazon, or Facebook.”

For more, visit [Hitachi Ventures](#).



Gayathri Radhakrishnan

Partner, Hitachi Ventures.

With more than \$1B AUM, the company is setting new standards for corporate ventures, fostering partnership and access for ambitious founders transforming the world around us. From advanced AI and robotics to sustainable energy solutions and beyond, Hitachi Ventures sees the potential in investing in companies that dare to dream big and disrupt the status quo. The firm’s expertise, coupled with Hitachi’s global resources and commitment to innovation, enables it to identify and nurture promising startups with the potential to drive significant impact and shape the future of technology.



‘Sensing’ plant floor disruption?

How GenAI & OT can help

Hitachi is bridging the OT & AI divide to drive steadier operations in manufacturing.

By **Kentaro Yoshimura, PhD**

Global disruption in manufacturing brought on by geopolitical and climate-based events, rising technical complexities, and the persistent displacement and reduction of skilled labor, shows no signs of letting up.

And although the era of digital transformation, with its data-driven approach to operations, was viewed as a magic bullet to overcoming some of these challenges, the reality is that turning the vast amounts of information generated by everything from IoT to cyber-physical-systems into actionable insights has proven elusive.

In most plants day-to-day operations consist of a series of activities ranging from commissioning, startups, and shutdowns, to setpoint tuning, inspections, and troubleshooting. All of this is captured under the moniker of operational technology (OT).

The volumes of information generated from OT typically fall into two distinct categories: OT data includes things like piping and instrument diagrams (P&IDs), mechanical and electrical diagrams, work orders, and more; and OT skills, take the form of diagnostic playbooks that experts follow – what to check first, how to reason about a control loop, which failure modes to rule out, etc.

Understanding and leveraging the complex interconnections between these distinct OT assets has eluded many in industry over the years. Even the advent



of large language models (LLMs) has proven challenging. Although generative AI (GenAI) has transformed how we work with text and images, the largest operational gains are still on the plant floor—where pumps, valves, drives, and controllers must run safely and predictably.

One of the only ways to start cracking the code lies in leveraging both OT data and the tacit knowledge of the most experienced plant floor operators.

Building an agent of change

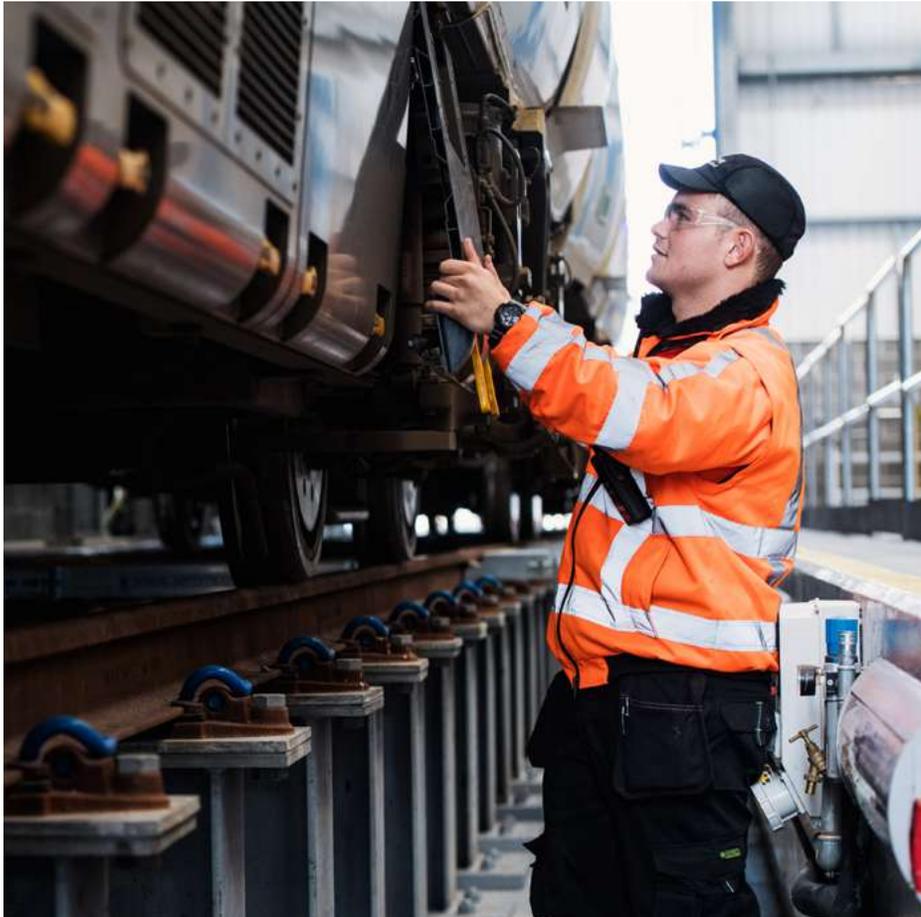
For Hitachi, our heritage in operating equipment and technologies dates to our founding, while our work in analytics, data, and AI goes back decades. Because of this, Hitachi remains one of the few companies that can speak to both sides of the “industrial AI” coin.

Daikin Industries Ltd., which manufactures commercial air conditioning equipment, [brought us in to help](#) connect these OT and AI worlds. Specifically, it needed an AI agent that could support equipment failure diagnostics in factories.

We kicked off a trial in April 2025 of an agent designed to understand equipment drawings converted into a knowledge graph, bind OT records, and execute a STAMP (System-Theoretic Accident Model and Processes) / CAST (Causal Analysis based on System Theory) analysis path to propose discriminating checks and corrective actions.

The solution leveraged historical OT knowledge from Daikin sites as well as new incoming reports. When a failure in equipment occurred during operation, the agent would alert the maintenance tech about the cause and take corrective action.

The trial demonstrated that the AI agent could match or exceed the diagnostic accuracy of average maintenance engineers, even for complex or previously unseen failures. In fact, it logged response times of about 10 seconds with greater than 90% accuracy. This would not only reduce mean time to repair (MTTR) but also help standardize maintenance quality across sites and shifts. By extension, the solution would also help address the skills gap while supporting Daikin’s manufacturing expansion.



interconnections) with encoded skill paths (the step-by-step diagnostic logic used by experts), like our AI agent system for Daikin.

The power of GenAI & OT

In industrial AI, the ethos is simple: move fast —and break nothing. That's the result of melding GenAI with OT. The real win isn't a chatbot, but steadier operations, faster troubleshooting, and fewer surprises. And when you add veteran expertise on top of it, you're not just talking about a vision for the future – but a practical, proven approach that can be deployed today.

By capturing the “playbooks” of your best technicians and embedding them into AI agents, you can ensure that every shift, every site, and every new hire benefits from the collective knowledge of your organization. This is how manufacturing leaders will close the skills gap, reduce downtime, and build truly resilient operations for the next decade.

For more on Hitachi Research visit: [Research & Development: Hitachi](#)

The solution, called “AI Agent for Equipment Failure Diagnostics,” worked by first converting Daikin factory equipment drawings into knowledge graphs that GenAI can read. The system then learns the OT data and feeds it into Hitachi's unique equipment failure cause analysis processes based on STAMP.

Why RAG-only isn't enough on the plant floor

Classic enterprise AI starts with retrieval augmented generation (RAG), which is useful for surfacing manuals, drawings, and old tickets. But in mission-critical systems, serious failures rarely repeat; after every incident, engineers add countermeasures that change the next signature.

What works in industry is retrieval augmented reasoning: the ability to bring back the right diagram slice or standard operation procedure (SOP) and then reason over the equipment's

graph and skill paths to propose discriminating checks and safe actions – with evidence.

In initial proof-of-concept tests, RAG-based systems could answer questions about known failures and those documented in manuals or past records with reasonable accuracy. However, when faced with similar but not identical failures, or entirely new failure modes, these systems struggled. For example, they might identify a “valve problem” but fail to specify which of dozens of valves is at fault or even misidentify the root cause all together. This is unacceptable in environments where downtime is costly and safety is paramount.

What's needed is the ability to analyze the control structure, trace the flow of materials or signals, and pinpoint the most likely failure points, even for novel scenarios. This is achieved by combining the knowledge graph (representing the equipment and its



Kentaro Yoshimura, PhD

Principal Researcher, Mobility & Automation Innovation Center, Research and Development Group, Hitachi, Ltd. Specializing in software engineering, particularly Software Product Line Engineering (SPLE) and generative AI applications, his work involves developing methods for managing software variability.



From air-gapped to wide open

The rising risks in industrial cybersecurity

Automation and AI are making it easier on cyber criminals. What's needed: A deliberate, comprehensive assessment plan based on industrial domain expertise.

By Joe Mullich

There's a not-so-dirty little secret in industry: despite advances in both AI and cybersecurity, the greatest cyber threats still come from social engineering and the exploitation of industry control systems (ICS).

The evolution of AI and generative AI (GenAI) has not produced as many novel threats as once feared. But the technologies have become force multipliers for traditional schemes, dramatically accelerating and expanding the scale, proliferation, and even the authenticity of threats. Once rudimentary social engineering schemes like email phishing, which relied on individual attacks, are now high-powered, more authentic and able to target far more victims than ever before – autonomously.

According to a [2025 report](#) from DeepStrike, the “weaponization” of AI for phishing schemes “has driven a surge in attack sophistication and volume, with some metrics showing a 1,265% increase in phishing emails since the launch of generative AI tools. Defenses are struggling to keep pace, as attackers increasingly bypass traditional MFA and exploit the human element, which is involved in over 60% of all breaches.”

“Automation and AI are making it so much easier on cyber criminals,” agreed Matt Castonguay, chief revenue officer for Hitachi Cyber, which operates Security Operation Centers and provides a range of cybersecurity services. “If you’re just someone in your basement trying to hack people, there’s only so much work you can do by yourself. Now you’ve got your AI, you’ve got your ransomware-as-a-service, you’ve got your botnets. You could scan 100,000 companies, 50,000 companies, in a day and find a whole bunch of vulnerabilities.”



Many people assume new technology presents the biggest risk....More often, the new tech exposes vulnerabilities in older systems that were considered secure before.”

“The AI doesn’t discriminate,” he added, “it’s just going to go and pick on everyone that it can.”

For more companies across industries, from energy to mobility, what’s needed to slow the increasing threat landscape is a deliberate, comprehensive assessment plan based on industrial domain expertise; one that examines industrial systems on two tracks: policy vulnerabilities and technical weaknesses.

Policy penetration

The reality for industrial companies today is that the isolated, “air-gapped,” plant is a thing of the past. Everything from cloud services and AI to edge computing are connecting more factories to the cloud than ever before. As these interconnections spread, often without recognition of potential policy breaches, the once seemingly impenetrable industrial company grows more exposed.

“Many people assume new technology presents the biggest risk,” says Castonguay. “More often, the new tech exposes vulnerabilities in older systems that were considered secure before.”

ICSs, the information systems used to manage industrial processes, typically refers to “supervisory control and data acquisition (SCADA) systems used to control geographically dispersed assets, distributed control systems (DCS), and smaller control systems using programmable logic controllers (PLC) to control localized processes,” says the National Institute of Standards and Technology.

For decades, ICSs were air-gapped or completely isolated from any network. A device controlling a conveyor belt or rail switch, for example, was designed for one purpose: to follow simple instructions reliably. Security wasn’t built into the system because remote accessibility wasn’t built into it.

However, that isolation has become an illusion over time. Many ICSs have been augmented with ad-hoc channels that allowed system integrators or contractors to configure files and patch software. As a result, these channels have become pathways to be exploited.

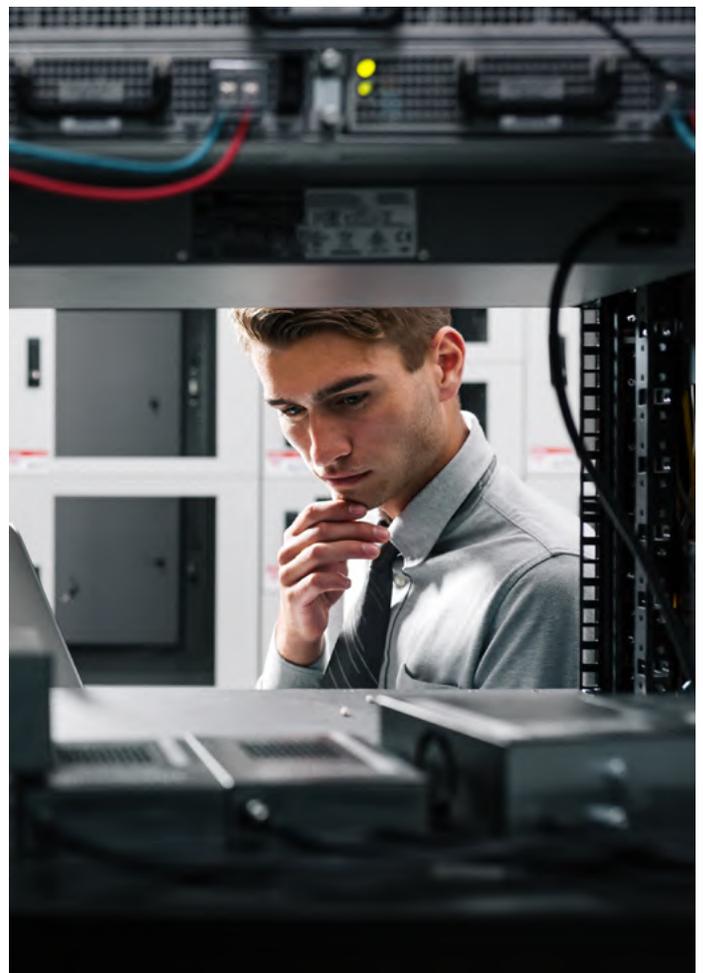
To be sure, in an era of Industry 4.0 and smart grids, true isolation is now obsolete. Consider the proliferation of IoT. Factories, power plants, railroads, and other industrial operations are aggressively integrating AI and IoT devices to boost efficiency, enable real-time monitoring, and improve performance.

And the proliferation isn’t slowing. The number of connected IoT devices worldwide is expected to balloon from around 21 billion in 2025 to more than 50 billion in just 10 years, [according to](#) market research firm IoT Analytics. Each connection point links back to infrastructure that often pre-dates modern security practices by decades.

“So we do the assessment,” Castonguay says. “If we find something that’s critical, we will flag it to be corrected right away. But if not, we conduct a remediation report with a gap analysis. We say, ‘here are all the issues that you have and here they are sorted by priority.’ Any minor concerns, we advise they update their security policy.”

The technical vulnerability

Warning signs about the vulnerability of legacy ICS have been flickering for years. As far back as 2008, a 14-year-old in Poland built a homemade infrared device to [take control](#) of a tram track switch. He derailed several



trams, injuring multiple passengers. In this case, the ICS required no authentication at all.

“Security wasn’t considered when these systems were built,” Castonguay says. “Today it must be applied to every process, both retroactively and proactively.”

In addition to exposing new vulnerabilities, AI provides bad actors with better hacking tools, such as automated malware, encryption, and ransomware management, much of which can be acquired as-a-service for a monthly fee.

“The bad actors aren’t just fishing with one line, but casting a net,” Castonguay says. “They can scan 100,000 companies in a day and find a whole bunch of vulnerabilities.”

When a previously unknown security flaw is discovered and disclosed (a zero-day vulnerability drop), attackers can exploit it within hours. In contrast, updating the ICS across a factory or rail network can take weeks or months.

“It takes a lot longer to patch your systems than it takes for a hacker to scan for vulnerabilities,” Castonguay says.

Ironically, traditional attack methods remain effective. Despite years of security training, people still click suspicious links and reuse passwords. What’s changed, Castonguay says, is AI’s ability to make phishing attempts far more convincing, such as crafting emails that mimic legitimate contracts or vendor communications with alarming accuracy.

As AI enhances bad actors’ toolkit, targets are not governed by the size of the operation nor industry. “No organization is too small to be targeted anymore,” Castonguay warns.

And the urgency is real. Attackers aren’t waiting for defenses to catch up. “The fishermen never sleep,” Castonguay says, “and their nets are getting larger and more sophisticated every day.”

The two-pronged approach to plugging gaps

Identifying the weak points in decades-old industrial systems requires a systematic approach; a comprehensive assessment that examines industrial systems for both policy vulnerabilities and technical weaknesses.

- **Policy** audits determine which individuals have access to various systems, rather than simply examine security documents. It’s not uncommon for such audits to discover that former employees have active system access, sometimes, administrators whose credentials can unlock every machine in a facility. They can also find generic accounts like “jira_admin” with elevated privileges across multiple systems. “Just because you’re the administrator of one system doesn’t mean you should have access to all the systems on the network,” Castonguay says. “Each system should have its own Identity and Access Management policies.”
- **Technical** assessments include things like red teaming or penetration testing, which systematically attempt to breach systems using the same methods attackers might use, all conducted with client permission. Assessments can take as little as two days for a small operation or up to a month for larger, multi-site systems.

Another key aspect to Hitachi Cyber’s approach is its ability to pull in expertise from across Hitachi Group companies when necessary. Hitachi’s heritage in industrial equipment and operational technology (OT), as well as its decades of development and deployment of data and AI solutions, provides Hitachi Cyber with unmatched, deep domain expertise in mission-critical systems – the systems that support the world’s social infrastructure.

“A key advantage of ours is the ability to tap into Hitachi for industry-specific expertise when needed,” Castonguay says, “and likewise, to bring our

expertise to bear on industrial challenges Hitachi discovers in the market. AI can scan a million devices quickly. But the differentiator is the human expertise to know what these results actually mean.”

But there’s one more step to Hitachi Cyber’s industrial security strategy: early warning audits. According to Castonguay, most organizations pursue assessments like these only after a breach or when insurance companies mandate them. But successful companies apply a proactive approach, conducting regular security audits and penetration testing to address vulnerabilities before they are exploited, making industrial cybersecurity as routine as home inspections.

For more, visit: [Hitachi Cyber | Cybersecurity & Performance Analytics Solutions](#)

And for more Hitachi AI perspectives, visit [AI Resource Center – Hitachi Digital](#)



Matt Castonguay
Chief Revenue Officer, Hitachi Cyber.

Hitachi Cyber is a global leader in advanced cybersecurity and performance analytics solutions, serving clients in over 50 countries for more than 50 years. With its innovative approach and 24/7 operations, it delivers tailored solutions to safeguard organizations against evolving threats and drive confident growth.



Sowing institutional knowledge across the industrial metaverse

How Hitachi Research is blending XR, AI and the metaverse to help industry overcome mounting challenges to the world's aging social infrastructure.

By Calvin Hennick

An inexperienced worker shows up at an industrial jobsite. The equipment is unfamiliar, and so are the problems that need to be fixed.

But instead of calling in a senior technician to explain the task at hand, the worker grabs a tablet, or a virtual reality (VR) or augmented reality (AR) headset and enters a digital twin of the worksite – a metaverse environment where information about equipment specifications, operators, and projects can be visualized in 3D.

The information may still be overwhelming at first, however. That's when an AI agent synthesizes the data with real-time sensor inputs that track worksite events and conditions and guides the worker step-by-step through the repair.

The result: Fewer errors, less stress, and a faster resolution.

That's the future envisioned by Hitachi researchers, who have successfully brought agentic AI to the company's

Worksite-Augmenting Metaverse. The agent, Frontline Coordinator – Naivy (or simply “Naivy”) – is aimed at both improving efficiency and reducing the “psychological burden” on inexperienced workers in frontline operations.

The tool is still in research and development (R&D), but in testing, Naivy helped a Japanese semiconductor manufacturer improve the performance of inexperienced workers on facility management tasks by around 30%. [In the trial](#), Naivy was also effective at alleviating the associated stress of young workers.

“Hitachi’s ultimate goal with [Worksite-Augmenting Metaverse](#) and [Naivy](#) is to create innovative work styles for frontline workers,” says Takayuki Fujiwara, chief researcher at Hitachi’s Advanced Artificial Intelligence Innovation Center. “This will help ensure that society can continue to thrive and progress despite the ongoing decline in the labor force.”

The intersection of AI and XR

Across the globe, governments are struggling to maintain and manage social infrastructure like roads, bridges, utilities, and public facilities.

For example, the American Society of Civil Engineers gives U.S. infrastructure a “C” grade in its [2025 Report Card for America’s Infrastructure](#). That’s actually an improvement over recent years, but the organization forecasts a \$3.7 trillion infrastructure funding gap through 2033.



Workers can learn the process of manufacturing, driving, inspection, and even surgery in virtual environments.”

Along with funding and deferred maintenance, demographics present a challenge for updating social infrastructure. According to the World Health Organization, the share of the world’s population over 60 years will [nearly double](#) between 2015 and 2050, from 12% to 22%. This means that experienced laborers will be retiring in droves, with fewer young workers coming up to replace them.

Fujiwara sees the technologies behind Worksite-Augmenting Metaverse and Naivy (a portmanteau of “navigation” and “AI”) helping to bridge this gap, not only for social infrastructure, but also in industrial fields such as manufacturing, energy, and logistics. “I feel a great sense of fulfillment in leveraging Worksite-Augmenting Metaverse and Naivy to address diverse challenges across various fields—such as energy, mobility, and connective industries—under one common concept,” he says.

The combination of XR (eXtended Reality) and AI is uniquely suited to help guide workers at industrial sites, Fujiwara notes. While the hype around XR use cases like gaming and social media has largely died down from its peak several years ago, Fujiwara calls the metaverse an “immersive virtual environment” that lends itself to industrial uses like training, simulation, and rapid communication. “Workers can learn the process of manufacturing, driving, inspection, and even surgery in virtual environments,” he says. “The metaverse can represent complex simulated results, making them easier to grasp. And stakeholders can interact and share information instantly, regardless of their physical location.”

By layering AI on top of metaverse technologies, Fujiwara says, researchers can enhance these benefits and improve efficiency, increase automation, and expand analysis capabilities in industrial environments. For example, AI can be used to identify and alert teams to hazardous actions in real time, accelerate quality inspections, and optimize workflows by automating repetitive tasks.

“By integrating AI with the metaverse, AI not only analyzes situations, but also presents critical information and supports users with intuitive three-dimensional direction,” Fujiwara says. “In other words, the real world can be reproduced virtually in the metaverse. AI can understand the synchronized situation and show appropriate direction to users.”



Tapping into the immersive qualities of the metaverse

Hitachi's vision for XR and AI at industrial worksites is already moving from the lab into the field. The company's R&D team has launched several proof-of-concept projects across several industries, including facilities management, transportation, energy, and manufacturing.

Fujiwara says the combination of Worksite-Augmenting Metaverse and Naivy has led to enhanced training, with AI creating XR content and support tailored to users' specific levels to improve their understanding. By offering context-aware assistance, he adds, the technologies can improve operational and maintenance support. Ultimately, Fujiwara notes, the tools help build consensus among diverse stakeholders, since everyone on a team has access to the same 3D environments.

In a railway maintenance pilot, for example, the solutions helped operators coordinate data including maintenance manuals, reports, and images, allowing users to quickly identify the parts needed for maintenance tasks. In an energy pilot, the technology helped identify potential risks in ways that were easily understood by users. And in a logistics

pilot, the tools allowed novice workers to operate factory inspections.

"These initiatives demonstrate how combining AI's reasoning capabilities with the immersive qualities of the metaverse can deliver practical benefits—reducing downtime, improving safety, and ensuring resilient operations," Fujiwara says.

While Worksite-Augmenting Metaverse and Naivy aren't yet available for commercial use, Fujiwara notes that this sort of real-world research is needed to push technology forward. "Hitachi has real-world operational sites," he says. "We have IT experts specializing in areas such as AI, as well as OT experts with deep knowledge of design, manufacturing, operation, and maintenance in fields like energy, mobility, and connective industries. This allows us to work together to solve challenges."

"Having both IT and OT specialists within the same company who can engage in deep discussions and collaborate closely is rare," Fujiwara adds. "We believe this is a unique strength of Hitachi."

For more, visit: [Research & Development: Hitachi](#)

And for more Hitachi AI perspectives, visit [AI Resource Center – Hitachi Digital](#)



Takayuki Fujiwara, Ph.D.

Microsoft MVP for Web Development, AI Business Expansion Office, Hitachi Ltd.

Hitachi Research & Development. Based on its corporate mission to "contribute to society through the development of superior, original technology and products," Hitachi is further evolving its Social Innovation Business for a sustainable society using Digital, Green, and Innovation, as means to resolve challenges and drive growth. In these efforts, the Research & Development Group is focused on driving innovation through digital, making full use of the technology platforms and know-how it has developed and accumulated for its various business domains.



The criticality of introducing AI into mission-critical systems

A Q&A with Hitachi on its approach to bringing AI into the physical world.

By **Joe Mullich**

The “Cambrian explosion” of AI development continues apace with complex and diverse tools, models, and applications springing to life across every facet of business. As it expands, maintaining clarity of function amid the growing phyla of categories is increasingly important to minimize confusion and streamline everything from planning to procurement.

For its part, Hitachi continues to build on its domain expertise in industrial AI, a fast-growing subcategory of physical AI. While [physical AI](#) lets systems like cameras and robots perceive, understand, reason, and perform or orchestrate complex actions in the physical world, industrial AI does the same for equipment and processes in mission-critical systems in industrial settings.

The distinction, however subtle, is significant against a backdrop of use cases, be it managing the flow of

electricity through a power grid or reducing the maintenance costs of high-speed rail. For industrial AI, a host of additional dynamics from regulation and compliance to safety must be built into every design, tested, and then certified. This is the world of the mission-critical system, into which the introduction of AI is, well, critical.

The series editor, Michael Zimmerman, and I recently assembled a group of AI experts from across Hitachi for their insights into this world. Joining us were Frank Antonyamy, Chief Growth Officer for Hitachi Digital; Jason Hardy, CTO of AI at Hitachi Vantara; and Chetan Gupta, PhD, Head of AI at Hitachi Global Research, and the General Manager of the Advanced AI Center at Hitachi Ltd. The following are excerpts.

Q1: What makes industrial AI different from other AI applications?

Antony: In consumer AI, you can tolerate errors. If a chatbot gives a wrong answer, it's annoying, but has limited impact. In industrial AI, we're operating in mission-critical applications like power grids, manufacturing lines, and railway systems. A 90% or 95% accurate system is just not fit-for-purpose and cannot be deployed. We need systems, AI included, that are reliable, predictable, and accurate. Not just some of the time, but all the time. These are also often integrated into real-world systems – not just standalone applications.

Hardy: Right, and from an infrastructure perspective, there's a level of detail that goes into designing for an environment that literally cannot go offline. The stock market, the rail network, and the U.S. power grid are hypercritical systems. While some environments can schedule downtime or have flexibility, there's a significant cost incurred in designing for something that cannot go down. That's why the cloud isn't always the best place for this work: 99.9% availability isn't sufficient for systems that truly cannot fail.

Gupta: Beyond system- and algorithm-level reliability, industrial AI brings two additional considerations: multimodality and edge deployment. Industrial data is inherently multimodal—ranging from text in manuals and logs, to video from worksites, time-series sensor data from equipment, and discrete event data from operations. In practice, effective solutions often require models that operate across one or more of these modalities. Building on Jason's point, deployment constraints matter just as much: Many industrial use cases require not only on-prem solutions, but true edge deployment to meet stringent latency, reliability, and data-sovereignty requirements.

Q2: How does this play out in practice?

Antony: Take our work with Hitachi Rail. When we deploy AI systems on trains, we're not just installing hardware. We must meet rigorous industry certifications. The equipment and systems must comply with railway-specific safety regulations and should be able to operate in a rugged, physical environment. Medical devices have their own certification standards; utilities have their own. For each industry, we must understand the compliance requirements and ensure 100% adherence. There's no choice if you want to deploy at scale in these environments.

One way in which we achieve this is through extensive simulation. We simulate millions of real-world scenarios using synthetic data. Only when we're confident these models will behave predictably across every situation do we put them into production. It's the opposite of the "release and refine" approach that's common with consumer AI because in our world, you can't afford to learn from failure in production.

Hardy: I completely agree. Think about the timeline. We're building AI for infrastructure designed to last 30, 40, or





even 60 years of continuous operation. A transformer in the power grid doesn't get replaced every five years like a laptop. The AI and the equipment supporting it need to be designed for that level of longevity. It's a completely different design philosophy than consumer technology.

Gupta: It means [moving fast without breaking anything](#). Domain experts must be first-class stakeholders at design time; deployment must be backed by rigorous testing and clear acceptance criteria; and rollout must be handled with care so that frontline workers trust the technology and meaningfully integrate it into their everyday workflows.

Q3: What about trust? When you apply AI to critical environments, how do you assuage any concerns people might have?

Hardy: We need to exceed the existing standards. With AI, there's an immediate fear factor, whether it's the "Terminator," or whatever gets into everyone's head when they hear "the machine is now autonomous." So there needs to be an expectation that if a human can do X, the AI needs to do X plus Y because regulators, operators,

and the public all demand demonstrably higher assurance. That's how we prove this technology is trustworthy.

Antonymsamy: Exactly, that's a critical point. We don't say that AI improves safety because that would imply the existing systems aren't safe. Industrial systems without AI already go through rigorous safety-critical verification and validation. What AI does is augment these already-safe systems, helping them operate more efficiently, with better yields, while maintaining that safety standard. It's about enhancement, not replacement.

Hardy: We're already seeing this level of precision and recall in other domains. In cancer detection, for instance, AI now picks up things that would have been missed just a few years ago. That's the standard we're bringing to industrial systems.

Gupta: Agreed. Building trust starts with transparency. In critical environments, that means being explicit with stakeholders about both the strengths and the limitations of the technology. It means designing "human in the loop" system where appropriate, building solutions with domain experts as integral members of the team, and

deploying them with frontline workers as partners—so trust is earned through clarity, collaboration, and demonstrated reliability.

Q4: You've mentioned rail, power grids, and manufacturing. How does Hitachi's experience working across those industries shape your AI work?

Antonyamy: The key elements of making these kinds of systems include access to high-quality data, domain expertise in these industries, and deep expertise in data science and AI. Hitachi has a heritage in industrials that goes back to our founding 116 years ago. Mission-critical systems are in our DNA. We've been doing OT/IT integration for decades, and we have deep expertise in data and AI. We're uniquely positioned in this corner of the AI market because we understand both the operational technology side and the AI technology side. This isn't new territory for us. It's an evolution of what we've always done.

Hardy: What's exciting is that we're actually putting 'One Hitachi' [a Hitachi Group philosophy that encourages intra-company collaboration] into practice. We're taking the best of our rail, energy, and industrial expertise across all our business units to solve challenges with the most complicated industrial problems on the planet.

Q5: Where is this space heading? What will industrial systems look like in a few years?

Antonyamy: We've already seen real impact in productivity improvements, energy consumption reduction, and throughput increases for customers. We're on a continuous path to offer better productivity, yield, and quality, while lowering energy consumption. That trajectory will continue. We're focused on making sure we have the right sensors and designing our systems to adapt to the changes. We are going towards autonomous infrastructure – self-balancing grids, manufacturing lines optimizing for higher yield and quality, human

frontline workers increasingly using agentic systems to augment their productivity, and machines that can self-diagnose.

Hardy: I think we'll see a symbiotic relationship develop where AI identifies inefficiencies and automates improvements, while humans make the bigger decisions. That's the future: Understanding and processing tremendous amounts of information at the machine level, interpreting and forecasting to head problems off before they start to occur, and bringing that high level of precision into very complex industrial systems.

Gupta: Beyond the gains Frank and Jason mentioned, industrial AI is increasingly evolving and moving beyond prediction and recommendation toward direct actuation. We are already seeing robots of different form factors—drones, quadrupeds, and potentially humanoids—being deployed to address workforce shortages in industrial settings. Even more transformative will be the shift toward designing industrial systems with an AI-first mindset. In mining, for example, we are seeing a move to smaller haul trucks as autonomy removes the dependence on human drivers. We saw a similar structural redesign with the introduction of electricity in manufacturing in the early twentieth century, and AI-driven design changes may prove even more consequential.

For more on Hitachi's work in Industrial AI, visit these sites:

- [Hitachi Ltd. - Research & Development](#)
- [Hitachi Vantara - AI and Analytics in Operations Management | Hitachi IQ](#)
- [Hitachi Digital - AI Resource Center](#)
- [Hitachi Digital Services - Industrial AI Success Strategies](#)



Jason Hardy
CTO of AI, Hitachi Vantara.



Frank Antonyamy
Chief Growth Officer, Hitachi Digital.



Chetan Gupta, PhD
Head of AI, Hitachi Global Research. General Manager of the Advanced AI Center, Hitachi Ltd.



There's (industrial) strength in numbers when building AI

How Hitachi's ecosystem of internal and external AI partners helps it tackle the world's toughest problems.

By Joe Mullich

The U.S. electrical grid is stressing under the weight of power-hungry AI data centers, the surging EV industry, and the energy transition. Railway companies are searching for ways to capture real-time insights to improve the vast infrastructure of railway systems, from trains to tracks. And factories around the globe are clamoring for greater autonomous optimization to boost performance and streamline supply chains.

As the collision of the industrial and digital worlds continues, individual tech providers are challenged to support the mission critical aspects of industrials alone. Tech providers know their innovations, while industrials know their industries and the myriad disciplines and dynamics that must be considered, built for, and supported.

The exception is Hitachi. With its heritage in industrial equipment manufacturing and operational technology (OT), and decades of development and deployment in data, analytics, and AI, Hitachi is one of the few equipped with both industrial and technical expertise, and a history of innovating solutions for transportation, energy, manufacturing, healthcare, and more.

The explosive growth of AI in recent years has driven even tighter collaboration among Hitachi Group companies. Along the way, the company has strategically expanded its external partner program, as well, to bring in leading AI developers, from NVIDIA to OpenAI, to take advantage of the technological investments of these firms.

This unique orchestration of its growing internal and external partner ecosystem is driving innovation at Hitachi, enabling it to solve some of the most complex challenges its customers are facing – from renewable energy distribution to faster, reliable transportation, to optimized autonomous manufacturing.

Synergy, not serendipity

None of this is by accident.

Hitachi's internal collaboration strategy, dubbed the Synergy Program, is a calculated program to coalesce technical expertise across Hitachi's Digital Systems and Services (DSS) group companies to solve customer challenges across industries. This includes companies like GlobalLogic, Hitachi Digital Services, Hitachi Vantara, Hitachi Cyber, and Hitachi Digital, as well as its industrial sectors of the mobility, energy, and connective industries. Although such collaboration occurred organically in the past when specific requests arose, today interconnectivity, awareness, and collaboration are strategic and intentional.

"Few companies have such breadth of talent and innovation to draw upon like Hitachi," said Roslyn Stuart, Vice President, Synergy Creation, Hitachi Digital. "Our strategy is to harness that expertise and tech to help customers solve problems and drive performance."

When a Hitachi company is faced with a complex customer challenge, Stuart and her team work closely with the engineering team to identify the right Hitachi digital partners to pull in – engage external partners where needed – and assemble a virtual delivery team. In past engagements, these teams have brought together as many as six different Hitachi companies with specialties in areas from digital infrastructure and software engineering to physical AI.

As teams come together, the company responsible for the customer account remains in charge of go-to-market and domain knowledge, while the Synergy team orchestrates the involvement and contributions of the supporting Hitachi companies.

Partnering on synchronicity

Adding to this amalgam of expertise, Hitachi continues to strategically expand its external partner program, pulling in innovation and expertise from leading AI developers – from NVIDIA to OpenAI – to co-create solutions for industry.

Over the past two years alone, the company has forged or strengthened partnerships and collaborations around AI, cloud and data centers, with the likes of [NVIDIA](#), [Google](#), [Microsoft](#), [AWS](#), [SingTel](#) and most recently, [OpenAI](#).

But there's more to partnering than meets the eye.

"Making partnerships work is not easy," says Frank Antonysamy, Chief Growth Officer, Hitachi Digital. "They need to align with the strategic vision of the company, and they need support from the top down, from the CEO."

One key to Hitachi's success, Antonysamy says, is its 360-degree partnership model, an approach that goes beyond traditional sell-with, sell-to, and buy-from arrangements. Hitachi and its partners actively co-create or "build with" each other, creating joint IP and solutions. These partnerships are designed to benefit both Hitachi and the partner, creating value that neither could generate alone.

For example, GlobalLogic, a Hitachi Group Company, helps Google deploy Gemini, its flagship AI model, to enterprise customers, while Google's AI powers Hitachi's Lumada solutions. Likewise, Hitachi plans to provide OpenAI with cooling systems and storage for new data centers, while OpenAI's models will be integrated into Hitachi's growing HMAX platform.

"When we're partnering with these companies, we're using their products, but they're also buying from us," says Rio Kurokawa, Vice President and Chief Lumada Business Officer who leads AI transformation at Hitachi. "They want energy solutions from us; transportation solutions, etc. It's more symbiotic than traditional partnerships."

It helps that Hitachi comes to the table with a breadth of experience building everything from transformers to trains, and from digital infrastructure to agents. The company understands the power of partnership like few others. It's an ethos that is only amplified with advance of industrial AI, a fast-growing pillar within the physical AI market.

Industry in the crosshairs

Complex, industrial problems demand a convergence of capabilities—chips, models, domain expertise, physical assets, system integration—that no single company possesses. Consultancies may have integration skills but lack access to actual trains and transformers. Chip makers may have compute power but no experience running a power grid.

"When you combine IT technology with OT technology, it gets messy," says Kurokawa. "Most industrial environments are regulated and safety critical. You don't have those issues on the consumer side. That's why industrial AI requires smart, able partners who bring standards, who can integrate, and who are compliant."

Antonysamy agreed. "What a lot of companies in our space don't have is actual physical assets and domain expertise," he said. "We're not hunting for a unicorn at Hitachi. We already have an existing customer for deploying these complex technologies with equally complex business models – it's us."

To that extent, Hitachi often plays the role of customer zero. "We build these solutions across our own environments, build the playbooks, create Hitachi IP and validate them at scale," says Hitachi's Stuart. "Only then do we take them to external customers."

Putting it all together to solve the world's hardest problems

With such fire power in its internal and external collaborations, Hitachi can move fast enough to tackle problems others avoid.

“We look for real-world problems and accelerate the development of digital and AI solutions through our integrated teams,” Stuart says. “As a result, clients see tangible outcomes sooner, rather than proof of concepts that never move beyond a lab.”

Southwest Power Pool (SPP) had one of these tough challenges. SPP manages the electric grid across 14 states, from North Dakota to Louisiana. It was wrestling with a thorny industry challenge: Connecting new power sources to the grid requires exhaustive analysis that currently takes 18 to 27 months. As a result, more than twice the nation's current generating capacity sits waiting in a backlog.

Hitachi is working with SPP, partnering with NVIDIA, to build solutions that reflect the needs of grid complexities of today – not from decades ago. This requires, in addition to the existing capabilities of SPP, NVIDIA and Hitachi building new solvers and algorithms that can handle this scale of data. To conduct the fundamental research for that, the troupe brought in expertise from Hitachi America's R&D team.

The result: Hitachi gets a better solution for SPP. NVIDIA gets a real-world test case that improves its product for everyone.

The power of Hitachi's partnership model was also on display in late 2024 when it rolled out an early version of HMAX. This AI-powered platform uses sensors on trains and infrastructure, allowing volumes of data to be processed at the edge in real time.

The project involved the orchestration of development between Hitachi Rail, GlobalLogic, Hitachi Digital Services, and NVIDIA. Not surprisingly, when the innovative muscle of such companies comes together, things

happen quickly. HMAX went from the inception, the very first conversation, to the actual product launch (at InnoTrans in September 2024), in less than five months.

“One of our goals at Hitachi is to solve the world's most pressing social infrastructure challenges,” said Daniel Knoch, Global Alliance Director, NVIDIA, at Hitachi Digital. “We think physical AI is critical for that. Our collaborations with NVIDIA enable us to accelerate our physical AI solutions and bring them to production.”

Internal collaboration is equally important. “Being able to work as ‘One Hitachi’ with our industrial companies, with our digital & AI companies, such as GlobalLogic, we can very quickly have AI engineers and architects available,” said Stuart, who along with Antonysamy were instrumental in the HMAX development. And she emphasized that it was the kind of challenge that a single provider simply could not have supported.

Hitachi has now evolved HMAX to transcend multiple industries. Just this month, the [company announced “HMAX by Hitachi,”](#) a full suite of next-generation solutions that brings the power of AI to the physical world. Derived from the original HMAX design principles, the suite integrates an array of advanced technologies from Hitachi and partners to dramatically improve reliability and performance of systems across the energy, mobility and manufacturing industries.

“Our internal and external partner ecosystem often reveals capabilities we didn't even know we had,” Kurokawa said. “Our partners' tools improve from solving our problems. Our solutions accelerate from their technology. That's the only way you can solve problems this hard.”

For more on Hitachi's extensive work in industrial AI visit: [AI Resource Center - Hitachi Digital](#)



Frank Antonysamy

Chief Growth Officer, Hitachi Digital.



Roslyn Stuart

Vice President, Synergy Creation, Hitachi Digital.



Rio Kurokawa

Vice President, Chief Lumada Business Officer for AI & Software Service Business Unit, Hitachi Ltd.

"The great thing about collections like this is the opportunity it affords the reader to step back and consider the compilation as a whole. Individually, the perspectives are surprisingly clear. But for me, stepping back a little, exposes an even bigger story. A story of intense commitment and global collaboration within a 110+ year-old tech giant, which is clearly more agile and aggressive than organizations half its age and size!"

Yuriy Yuzifovich — Chief Technology Officer, GlobalLogic

"What's impressive to me about this series is that each story shares a different view of how Hitachi is leveraging AI to help customers – today. This isn't a collection of case studies or industry pontifications which a lot of companies spend (waste) time producing. These are points of view on specific topics that our customers are wrestling with every day. In many ways, this is a resource!"

Prem Balasubramanian — Chief Technology Officer and Head of AI, Hitachi Digital Services

"Reading a single story from this collection will give readers insight into a topic they may not have considered before – even if they know the industry in question. But reading through the entire collection will provide perspectives about the complexities and opportunities of Physical and Industrial AI through our work in the lab and in the field."

Chetan Gupta, PhD — Head of AI at Hitachi Global Research, and General Manager of the Advanced AI Center, Hitachi Ltd.

"This book encapsulates in one place the breadth of development, engineering, and research of our global digital operation, as well as how it all comes together to solve the world's biggest challenges for customers and society."

Frank Antonyamy — Chief Growth Officer, Hitachi Digital

"In being a part of our Hitachi AI Center of Excellence, I've had a very good understanding of all that we were doing in AI, from Research through to our Group Companies. But seeing examples of all the great work captured in a series of crisp, individual, and consumable stories really brings our collective efforts around AI to life."

Rio Kurokawa — Vice President, Chief Lumada Business Officer for AI & Software Service Business Unit

"Since I've been with Hitachi Ventures, I've witnessed consistently positive growth of our work in AI across Hitachi, a lot of which is presented here in a coherent manner. The incredible efforts, both individually and collectively, have made our role in Ventures more strategic than I could have imagined. Very few companies have both OT & IT domain expertise, as well as incredible depth in AI. It has enabled us to seek out prospects that could strengthen, amplify, or fill a gap in this extensive strategy – which is good for both the prospects and Hitachi."

Gayathri Radhakrishnan — Partner, Hitachi Ventures