Research & Development

Cutting-edge AI to Deepen Lumada 3.0

May 28, 2026

Innovation & R&D, AI

Column: Hitachi's Research and Development Accelerating AI-driven Innovation in the Physical Domain

Hitachi is pursuing innovation through advanced AI (artificial intelligence) technologies in the physical domain. The physical domain refers to the real world that humans can touch, as opposed to the digital domain.

Currently, the fusion of rapidly advancing generative AI technologies and robotics has made “Physical AI,” which perceives and controls the physical world, a major technological trend. A key technology in this field is a new type of generative AI called “world models” that can interpret and infer physical phenomena. Moreover, technological advancements such as automatic programming of robots and machines, generative AI for images capable of interpreting drawings and physical environments, autonomous evolving Agentic AI, and the miniaturization and power saving of edge models are combining to expand the application range of Physical AI. This is expected to address labor shortages on site, enable precise work in hazardous environments, and allow flexible changes to production plans that respond promptly to management conditions.

Hitachi’s AI research history dates back to the 1960s. Since then, efforts have been made to integrate AI with mechanical and control technologies as well as knowledge of design and operation for application to social infrastructure. Hitachi has long-standing accumulation and achievements in the physical domain. Now, aiming for further advancement of Lumada 3.0, Hitachi is strengthening its initiatives in Physical AI even more.

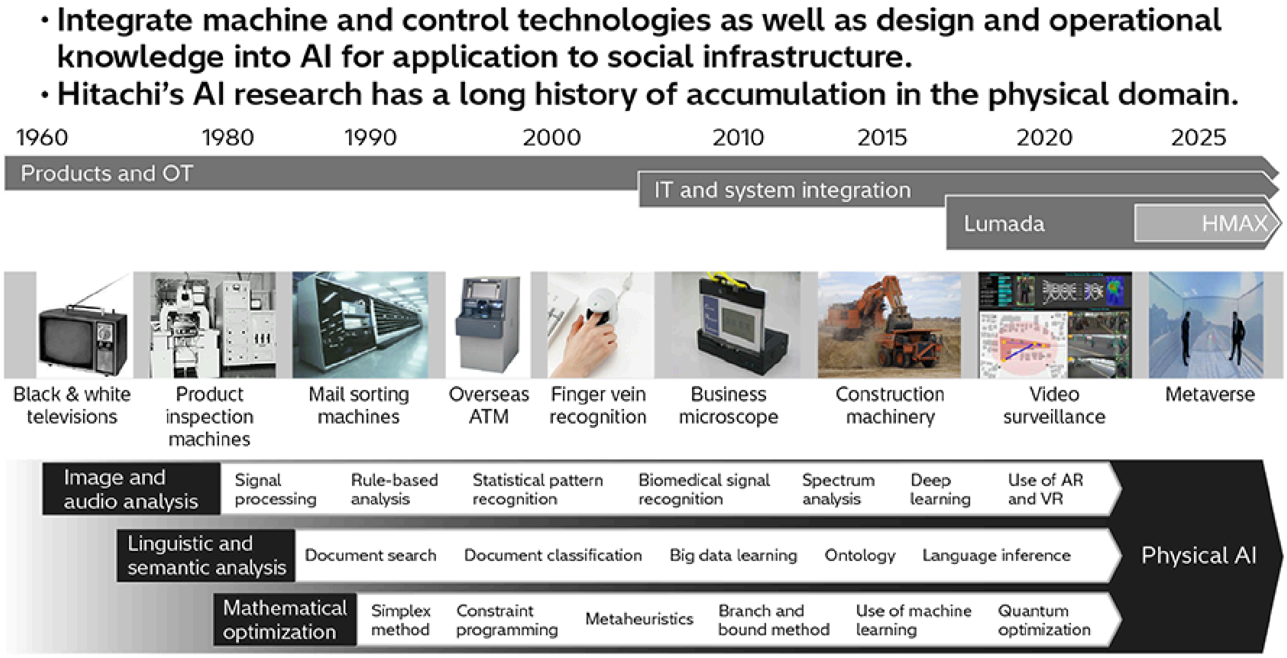

AI Research by Hitachi

OT: operational technology, ATM: automated teller machine, AR: augmented reality, VR: virtual reality

Hitachi's unique Physical AI, which supports social infrastructure, is a technology that autonomously evolves entire digital and physical systems while incorporating real-world environmental changes. The digital domain has the advantage of being able to plan overall optimization using large amounts of information. However, when targeting the physical domain, it is necessary to agilely incorporate real-time environmental changes and update the plan optimally each time. In social infrastructure fields such as power, railways, and industry, Hitachi, which has built systems spanning the digital and physical domains, aims to transform these systems into more adaptive and intelligent ones by integrating AI, creating a world where the entire system evolves through learning functions. Specifically, Hitachi focuses on three areas: (1) AI control of robots and machines through digital-physical cooperation, (2) AI safety assurance technologies to prevent harm to people and society, and (3) knowledge update technologies that incorporate tacit knowledge unique to the physical domain.

Currently, Hitachi is in a position to lead research and development of Physical AI. Applying AI to the physical domain requires inference based not only on sensor data, but also on background factors such as physical phenomena and operational conditions.

Hitachi possesses four accumulated knowledge systems from years of research and development. First, as a base, there is knowledge of physical laws such as thermofluids and electromagnetic phenomena. On top of that, there is knowledge related to the design and control of social infrastructure and products, as well as practical knowledge based on operation and maintenance procedures and simulators. Furthermore, there is knowledge essential for social implementation, including social constraints and causal relationships of infrastructure. These form the foundation for enabling AI to function appropriately in the real world.

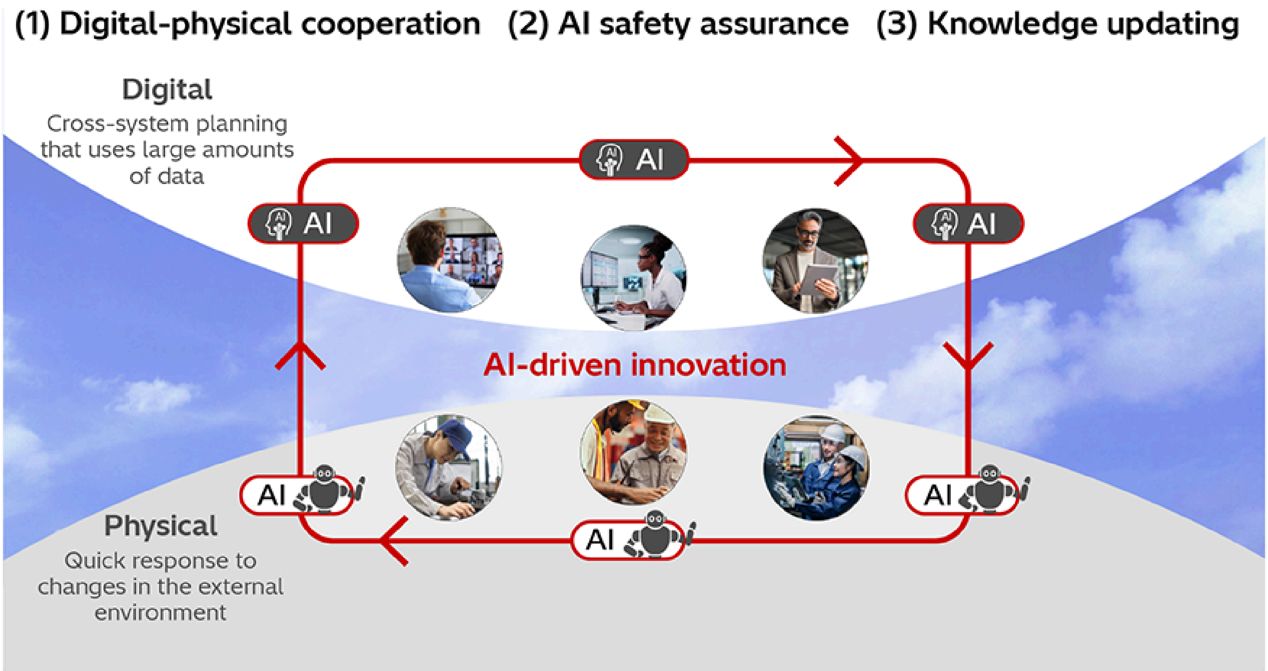

Physical AI for Social Infrastructure that is Distinctively Hitachi

The initiative to integrate this knowledge with AI is called IWIM (Integrated World Infrastructure Model). IWIM consists of a combination of a world infrastructure model trained on physical models and their corresponding digital twins, and a Large Language Model (LLM) trained on on-site documents and experiential knowledge. This enables AI to make decisions not only based on statistical trends in data, but also by understanding the underlying physical and operational meanings, thereby highly supporting decision-making in the field.

Utilization in the field has already begun. In manufacturing, AI agents that reference past equipment data and manuals have been introduced to indicate signs and causes of failures, achieving maintenance that does not depend on skill level and reducing downtime. In railway operation management, AI is operating to assist dispatchers by organizing the situation and suggesting recovery plans when troubles occur. These are good examples of reducing workload while ensuring safety by having humans make the final decisions. This article will explain cutting-edge cases leading Lumada 3.0.

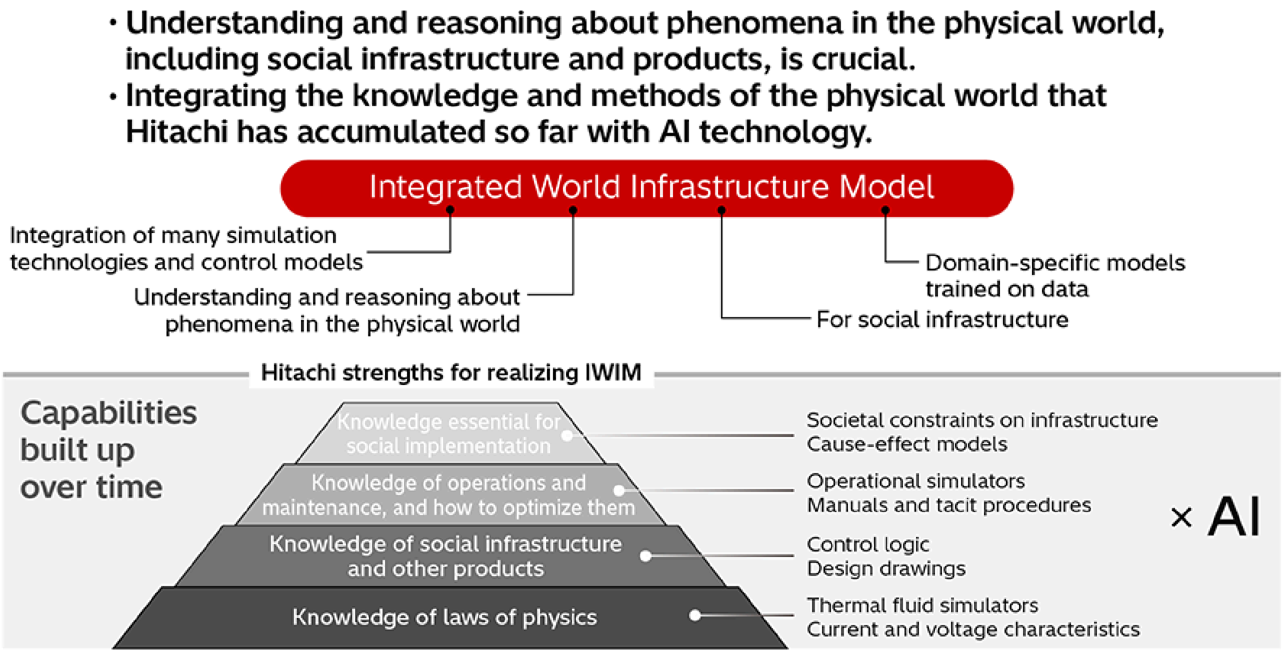

Overview of IWIM

1. AI for Business Enhancement

Artificial intelligence (AI) for business enhancement refers to technologies that support managerial decision-making by simulating business risks and opportunities based on macro-trends and contextual information collected from business operations.

The approach identifies salient business signals by integrating external sources (e.g., investment and market data) with early warning signs of diverse risks captured at the workplace. One representative application is multi-agent simulation, which generates scenarios for investment opportunities and operational or system risks. In such simulations, agents with domain-specific expertise (e.g., procurement, sales, legal, treasury, and so on) can be generated automatically. Simulation topics can be selected through analyses of external sources such as investor relations (IR) materials and economic reports. For example, by structuring challenges common to many corporations into a tree representation to extract industry macro-trends.

Furthermore, enabling workplace-deployed AI agents to participate in the simulations allows primary, real-time signals, such as business sentiment and risk appetite, to be incorporated. The resulting scenarios can be delivered to managers via dashboards, chat-based interfaces, or other decision-support channels.

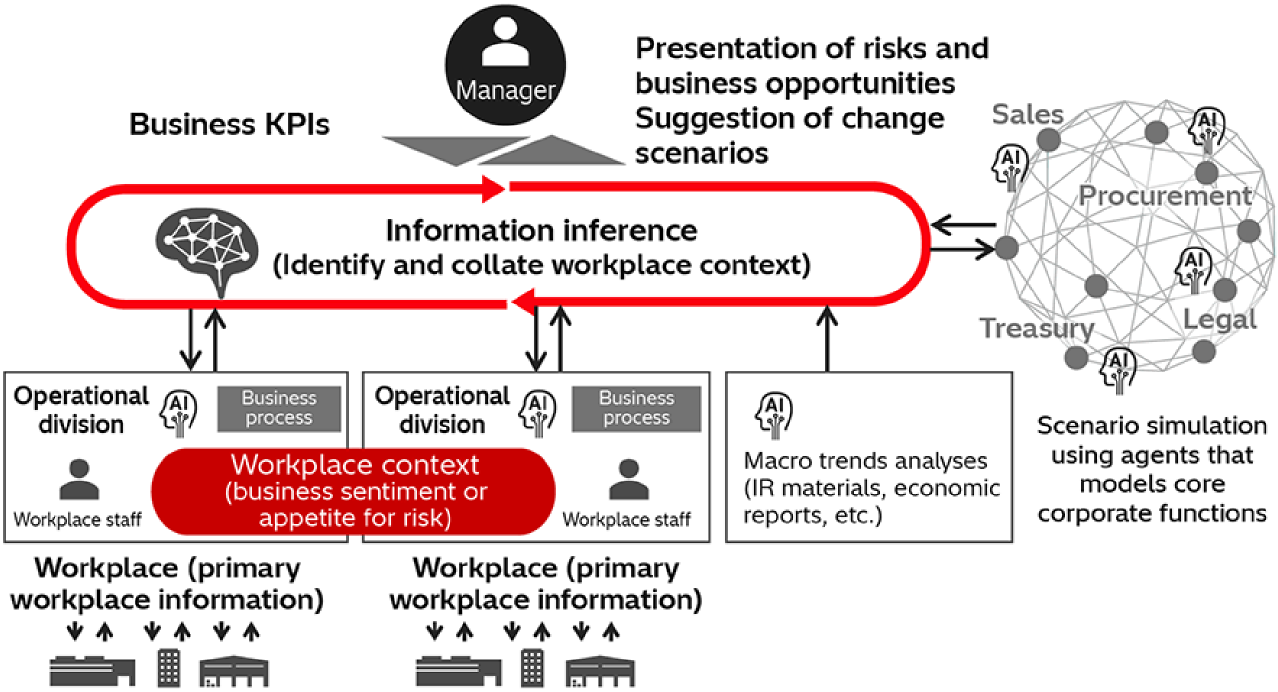

[1] Simulation of Business Risk Scenarios

KPI: key performance indicator, IR: investor relations

2. Domain-specific AI Models

Hitachi is putting generative AI to use in OT tasks like production planning optimization or infrastructure maintenance that involve the control of machinery. As many OT tasks have no tolerance for error and a complexity that derives from equipment condition, physical laws, and workplace rules, they require a generative AI that is customized to a specific purpose (vertical AI) by training on specialist domain knowledge and decision-making processes. Unfortunately, collating the decision-making processes of skilled workers in bulk is difficult, and thus conventional techniques for fine-tuning generative AI using large amounts of data are not always effective. Accordingly, Hitachi has developed a reinforcement learning technique for using small amounts of data to tune the thought processes of generative AI to suit OT applications.

This technique involves the use of trial and error to teach generative AI what thought processes to use to generate accurate production plans and to avoid incorrect responses. Using the technique, Hitachi achieved AI models with high reasoning capability and abstention capability for incorrect thought processes. Applying this technique based on the many different rules of specific tasks facilitates the safe use of generative AI in OT applications that require a high degree of reliability.

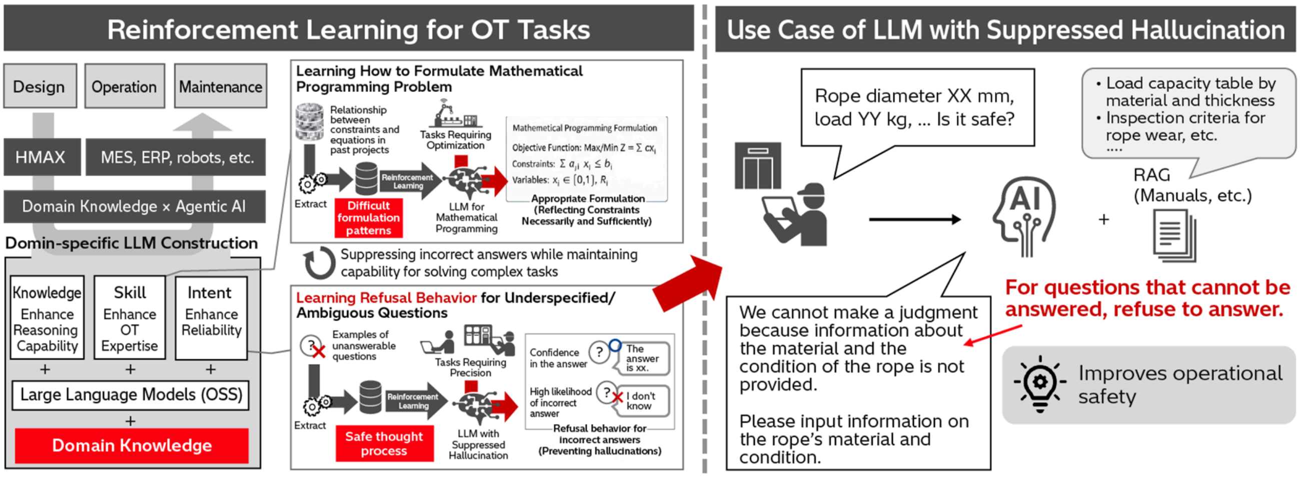

[2] Development of Reinforcement Learning for OT Tasks and Example Use Case

MES: manufacturing execution system, ERP: enterprise resource planning, OSS: open source software, RAG: retrieval-augmented generation

3. Adaptive RAG Framework for Context-aware Q&A

As generative AI becomes integrated into business workflows, ensuring reliable answers and interpreting complex, visually structured documents remain significant challenges. Conventional systems are prone to hallucinations and conventional text-centric retrieval augmented generation (RAG) methods often fail to capture the meaning embedded in diagrams, tables, and intricate layouts. The adaptive RAG framework addresses these issues by preserving page-level context and delivering fast, high-accuracy responses without sacrificing structural or semantic fidelity.

Strategic adaptive processing aligns retrieval and reasoning with the user's intent. The framework identifies whether a question requires factual confirmation, analytical interpretation, opinion, or contextual explanation, and each category is assigned to an optimized processing path that balances precision and latency. A multimodal vision-language model interprets visual and textual elements, modular plug-ins accommodate heterogeneous enterprise data, and an LLM-based re-ranking mechanism enhances relevance across candidate answers.

Through this integrated approach, the framework enables human-level comprehension of technical manuals, research papers, and financial reports. Such capabilities support informed and efficient decision-making across research, engineering, and business domains.

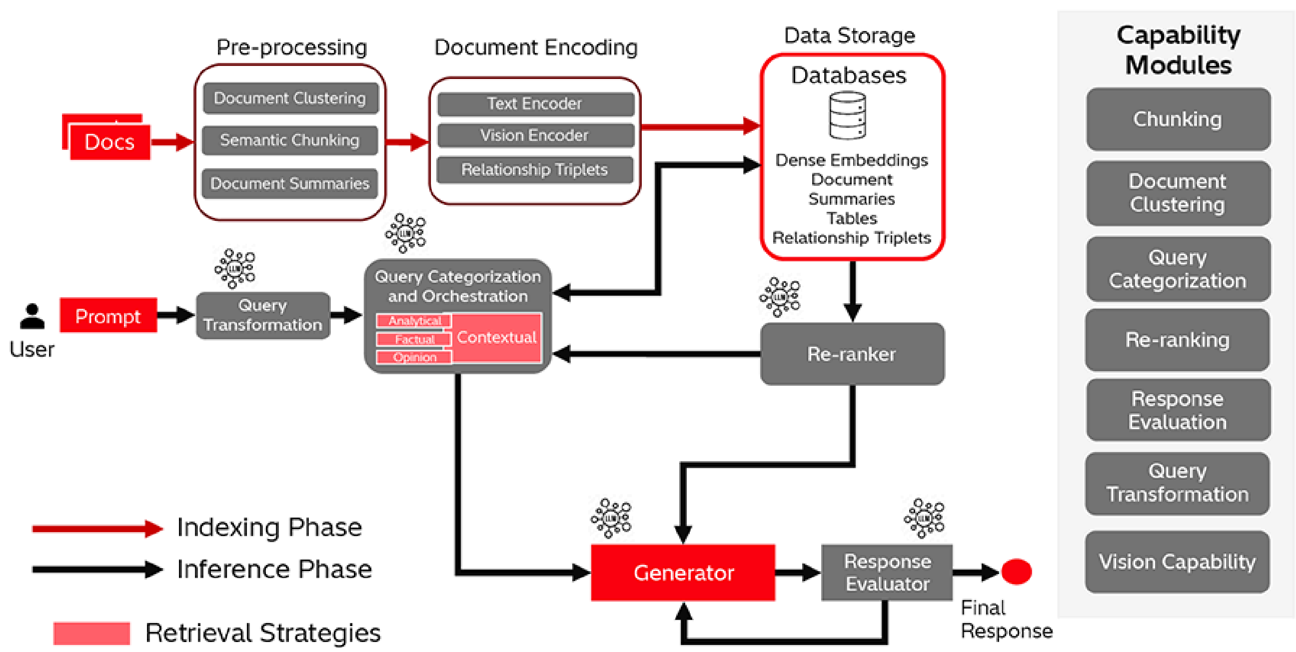

[3] Adaptive RAG Framework for Context-aware Q&A

4. AI Agent Platform

The Research & Development Group at Hitachi has been working on an AI agent platform to support Hitachi's "customer zero" strategy through the development, use, sharing, and deployment of advanced AI agents that are highly specialized and equipped for reliability and safety. By building an ecosystem to bring together agents that incorporate the diverse expertise and skills of Hitachi Group and developing multi-agent techniques that can cope with the complex and multi-faceted work of a large corporation, Hitachi has undertaken in-house verification trials aimed at implementing this technology and reforming business processes for the AI era. In addition to boosting operational efficiency throughout the group, Hitachi is also using this work to look at how it can deliver value through the flexible deployment of agents in customer businesses.

In the future, Hitachi intends to engage in further research and development combined with in-house implementation to encourage agent reuse and to get successful initiatives deployed more widely. It will also accelerate improvements in operational efficiency and the creation of new value while ensuring rigorous protection of data through the sharing and deployment of agents in ways that transcend organizational boundaries.

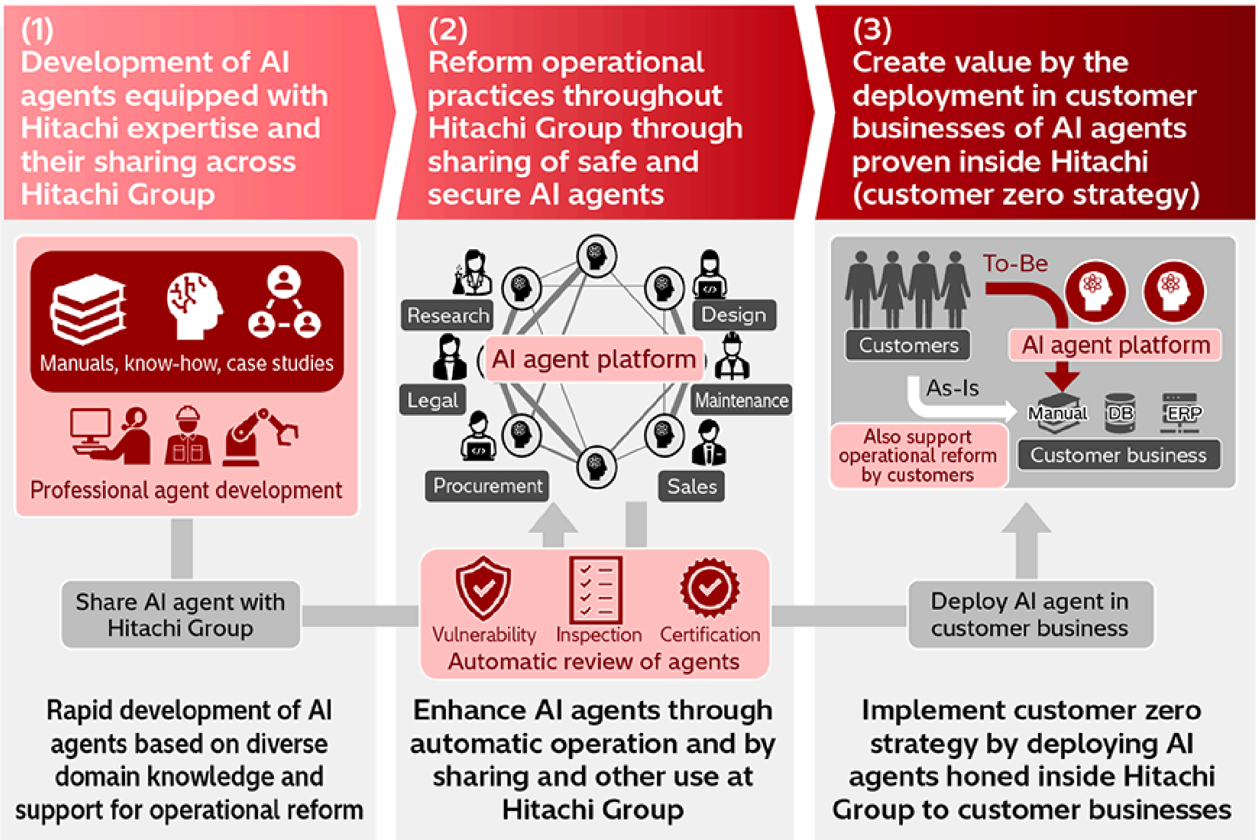

[4] Use Cases for AI Agent Platform

5. Strategy Support AI

The strategy support AI system uses backcasting to support planning for new business initiatives or policies at organizations such as local government, government agencies, or manufacturing businesses. To provide this planning support, the system uses multiple AI agents working in tandem.

The first step is to use a future scenario AI agent that can understand the vision of future society generated by Hitachi's future scenario simulator to analyze how this differs from today and to identify which societal challenges to address. Next, a societal challenge AI agent with knowledge of social trends performs repeated cause-and-effect analyses of these challenges to uncover the numerous technical issues that accompany them. A user technology AI agent that understands the technologies possessed by the user then produces numerous suggestions for how the user's technologies can be used to resolve these challenges. As a result, the AI system is able to offer many options for new businesses or policies in a short period of time.

The generated solutions can then be honed and narrowed-down by another AI to identify which ones are practical, thereby accelerating the creation of new businesses or policies for overcoming societal challenges.

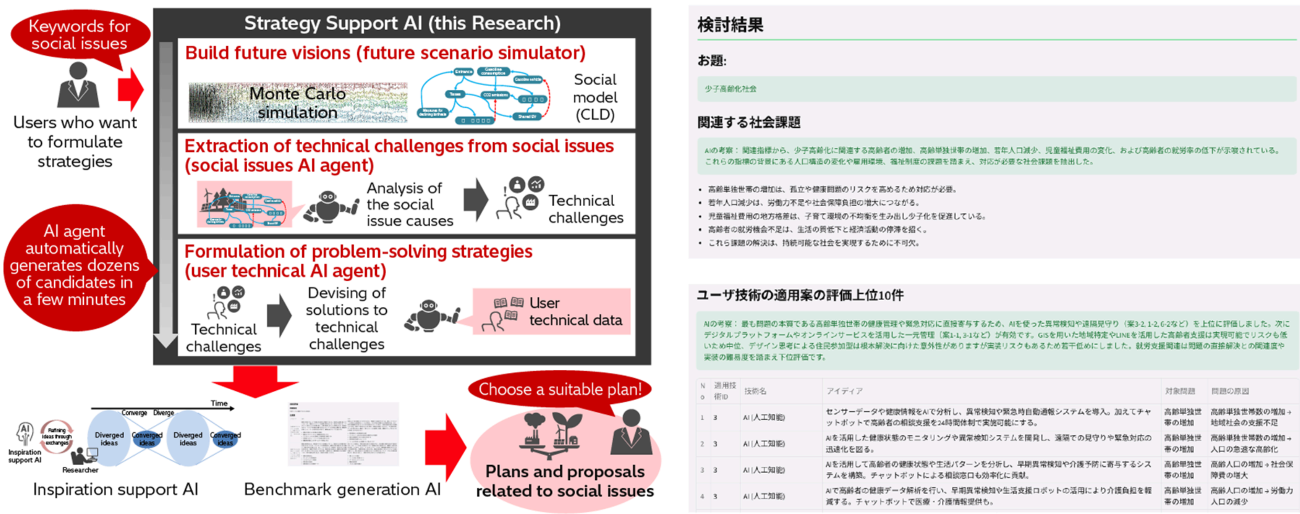

[5] Overview of Strategy Support AI and How it is Used

CLD: causal loop diagram

6. Underwriting Agent for Life Insurance

While life insurance underwriting requires an accurate assessment of the applicant's state of health, the fact that most applications are filled out by the applicant themselves and tend to be expressed in vague terms makes the process difficult to automate. Utilizing data from approximately 100,000 people, Hitachi has developed a risk simulator for insurance that estimates the risk of hospitalization due to lifestyle disease based on more than 200 assessment parameters or pre-existing conditions and is working to utilize the simulator in underwriting. While it enables automatic assessment by providing a natural language processing technique that utilizes dictionary matching to convert the disease names mentioned in applicant declarations to standard codes, the cost of keeping the dictionary up to date and the difficulty of dealing with complex written language has limited the proportion of applications that can be assessed automatically.

Instead, Hitachi has looked to the contextual comprehension capabilities of LLMs and combined this with a knowledge database of disease names to develop a method that minimizes hallucinations. This should dramatically improve the accuracy and scope of disease name resolution. By drawing on specialist knowledge and underwriting expertise to maintain appropriate control over generative AI, it is hoped that this will lead to more advanced assessment practices.

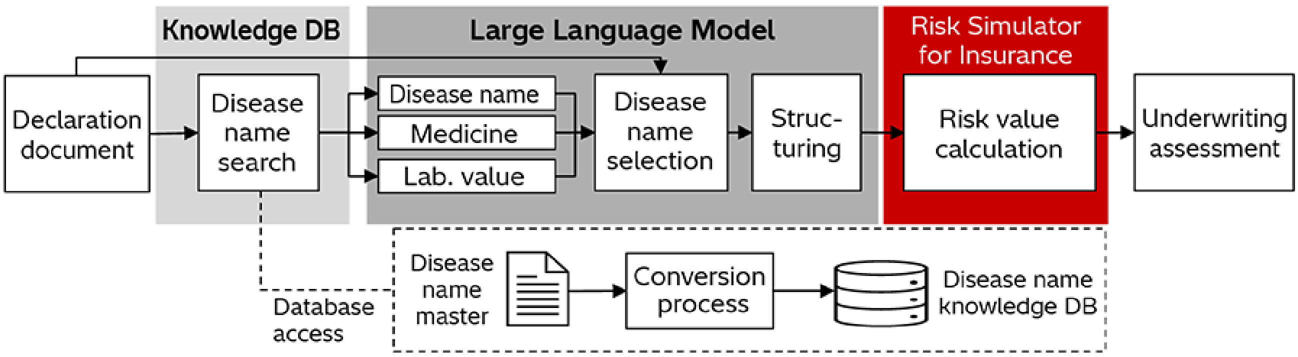

[6] Underwriting Agent for Life Insurance

7. MA-ATRIX for Assessing AI Adaptation to Drive Business Transformation

Expanding the use of AI across an organization and transforming business processes require a systematic assessment model. Such a model must objectively assess the status and progress of AI adaptation, visualize the current situation, and support the development of specific improvement measures and actionable plans.

Hitachi has jointly developed and published its maturity model titled "Maturity Assessment & AI TRansformation IndeX; Generative AI Adaptation Roadmap (MA-ATRIX)" with Gen-AX Corporation. It is a model for systematically assessing AI utilization and adaptation in an organization and supporting the staged implementation of business process transformation*1.

MA-ATRIX systematically assesses AI adaptation in terms of seven maturity levels across seven assessment dimensions. Based on Hitachi's extensive domain knowledge and more than 1,000 generative AI use cases, the model defines "goals" to be achieved and "practices" to be followed for each assessment dimension and maturity level. The model supports ongoing measures for business process transformation by providing a comprehensive visualization of the factors essential to getting from the proof of concept (PoC) stage to actual business outcomes, presenting this in terms of assessment dimensions that include "organizational readiness," "data governance," and "integration into business processes" as well as technical considerations.

*1. Hitachi and Gen-AX release MA-ATRIX maturity model for accelerating business transformation through generative AI

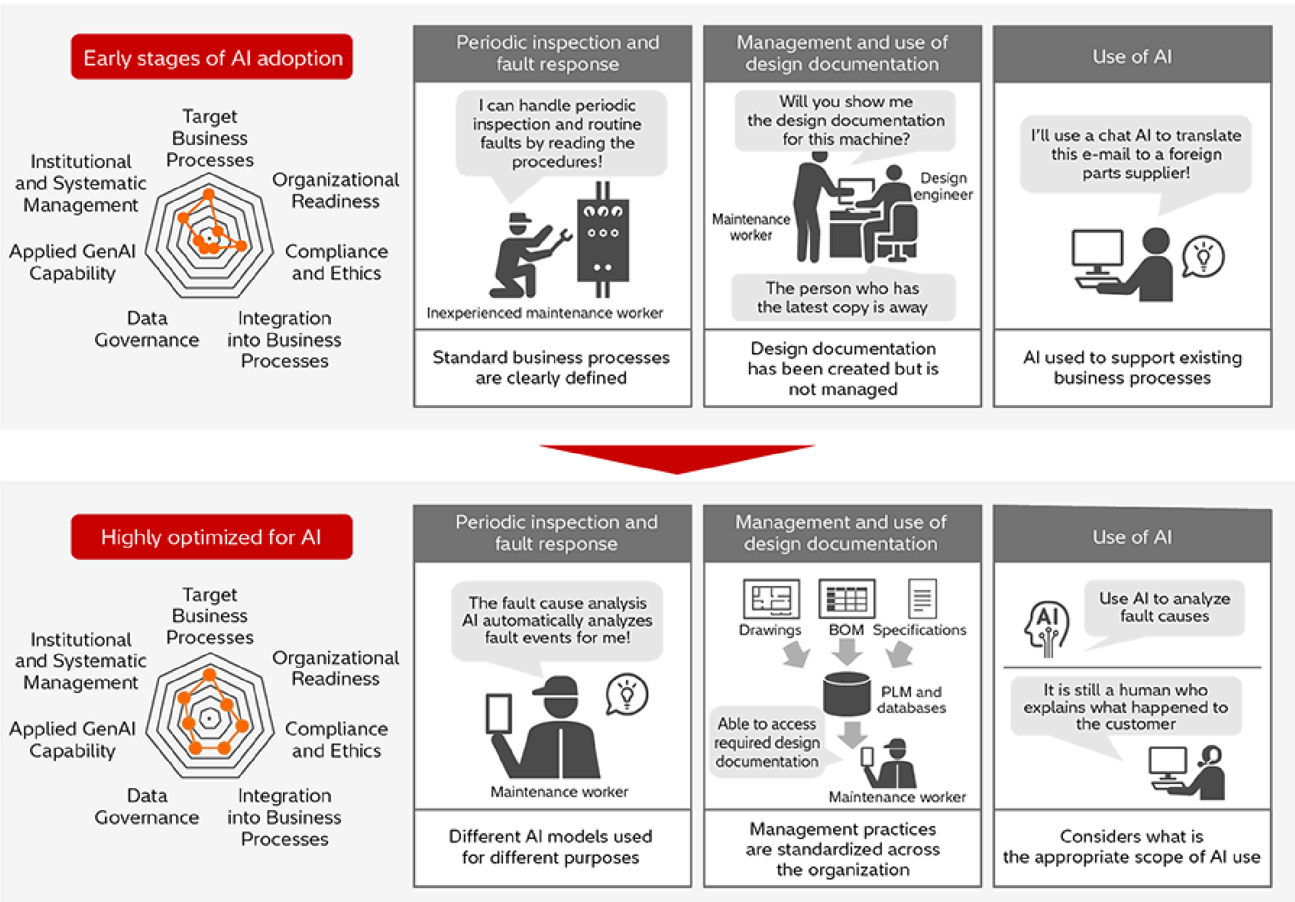

[7] Growth in Maturity of AI Use

BOM: bill of materials, PLM: product lifecycle management

8. Next-generation Naivy AI Agents for Reforming Workplace Safey and Efficiency

To address increasingly severe workforce shortages and the scarcity of skilled workers, Hitachi has developed its Frontline Coordinator – Naivy AI agent that combines generative AI with metaverse-augmented workplaces. The goal of the system is to resolve the pressing challenge of making work safer and more efficient while also reducing the mental stress of workplace tasks.

When used in manufacturing, Naivy integrates the OT knowledge of skilled workers with 3D spatial information about the workplace, and uses this knowledge for training. In addition to intuitive navigation using a tablet computer, the interactive capabilities of the system for planning work, consulting procedures, and assisting with troubleshooting provide powerful support for workers.

Naivy also serves as the core of a risk and hazard prediction (RKY) support system at construction and installation sites. By providing a realistic metaverse replica of the workplace and having the AI analyze instances of similar issues in the past, it performs hazard prediction at a level that previous methods struggled to achieve. Hitachi has demonstrated that this provides a dramatic improvement in workplace safety and efficiency.

In addition to delivering effective action on skills transfer and the improvement of work efficiency, upgrading knowledge of safety in the workplace, and driving digital transformation (DX), use of Naivy to acquire domain knowledge will also contribute to the implementation of IWIM.

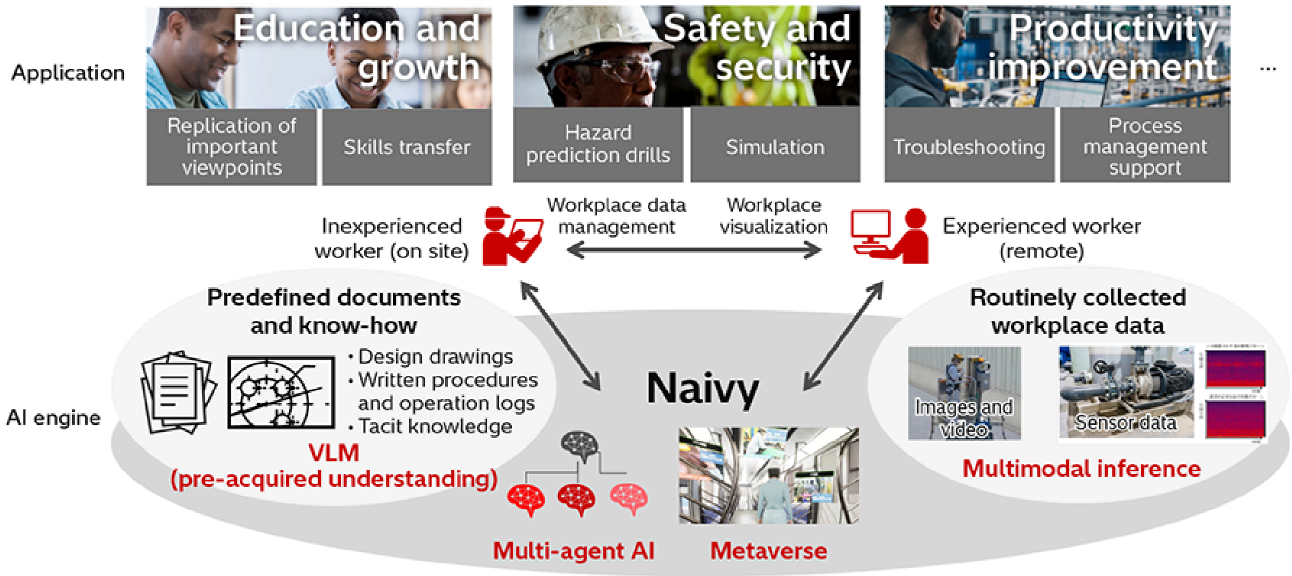

[8] Naivy AI Agents and Applications for Formalizing and Visualizing Workplace OT Knowledge

VLM: vision language model

9. AI Agent Incorporating Workplace OT Knowledge

Passing on the knowledge and expertise of skilled workers is a major challenge for the manufacturing industry. Hitachi has formalized OT knowledge in the form of OT skills and OT data, developing AI agents that are equipped with both specialist domain knowledge and knowledge of the workplace. This replicates thought processes by treating OT skills as formalized knowledge based on methodologies used by expert engineers such as system-theoretic accident model and processes (STAMP) and causal analysis using system theory (CAST).

Also, workplace-specific information has been converted into AI-readable format, including equipment drawings, work logs, and maintenance records. This is collated into a knowledge graph and used for context training of the AI so that the AI agents can deal with the characteristics of the actual equipment used in a plant and respond flexibly to unknown faults. In a trial conducted with Daikin Industries, Ltd., an AI agent was able to identify the causes of a fault rapidly (less than 10 s) and accurately (23% points better than previous methods (from 67% to 90%)). As a means of distilling the knowledge of workplace experts into digital form based on the IWIM architecture, it is hoped that this technology will serve as a new platform for boosting workplace performance and achieving operational autonomy.

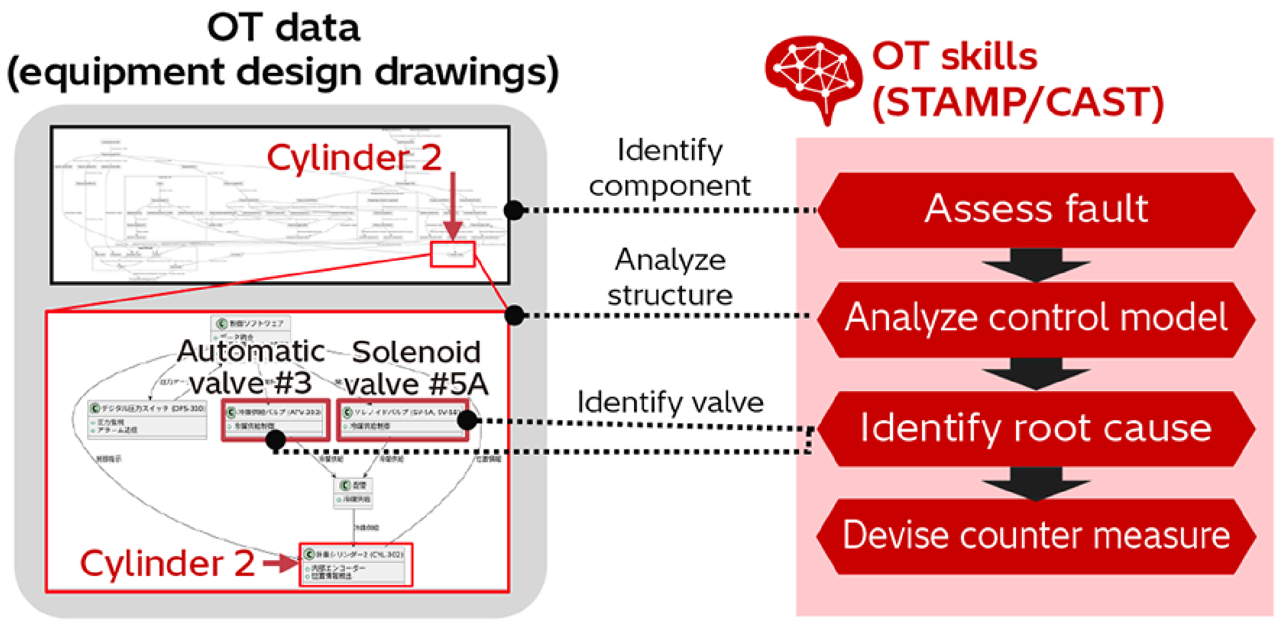

[9] How OT Knowledge is Combined with AI Agents

10. UDL for Rail Lifecycle Digital Thread

Unified Data Layer (UDL) is an AI-ready data foundation that supports the evolution of Lumada 3.0 by adding semantic understanding on top of existing siloed enterprise data, without requiring data movement or system changes. At Hitachi Rail's newly built railcar assembly factory in Hagerstown, Maryland, UDL logically integrated siloed data across multiple systems, enabling faster decision-making and shorter ramp-up times for applications.

Building on this success, efforts are underway to expand UDL step by step from the assembly factory to maintenance depots, enabling information flow that was previously fragmented. As manufacturing and maintenance are connected, new challenges arise, not only handling differences in data definitions across systems, but also maintaining consistent context across the product lifecycle. UDL enables a digital thread linking manufacturing and maintenance, opening new opportunities to apply AI-driven insights across the rail lifecycle. For example, visual inspection results and maintenance history collected at depots can be fed back into design and manufacturing to accelerate root-cause analysis and improve maintainability.

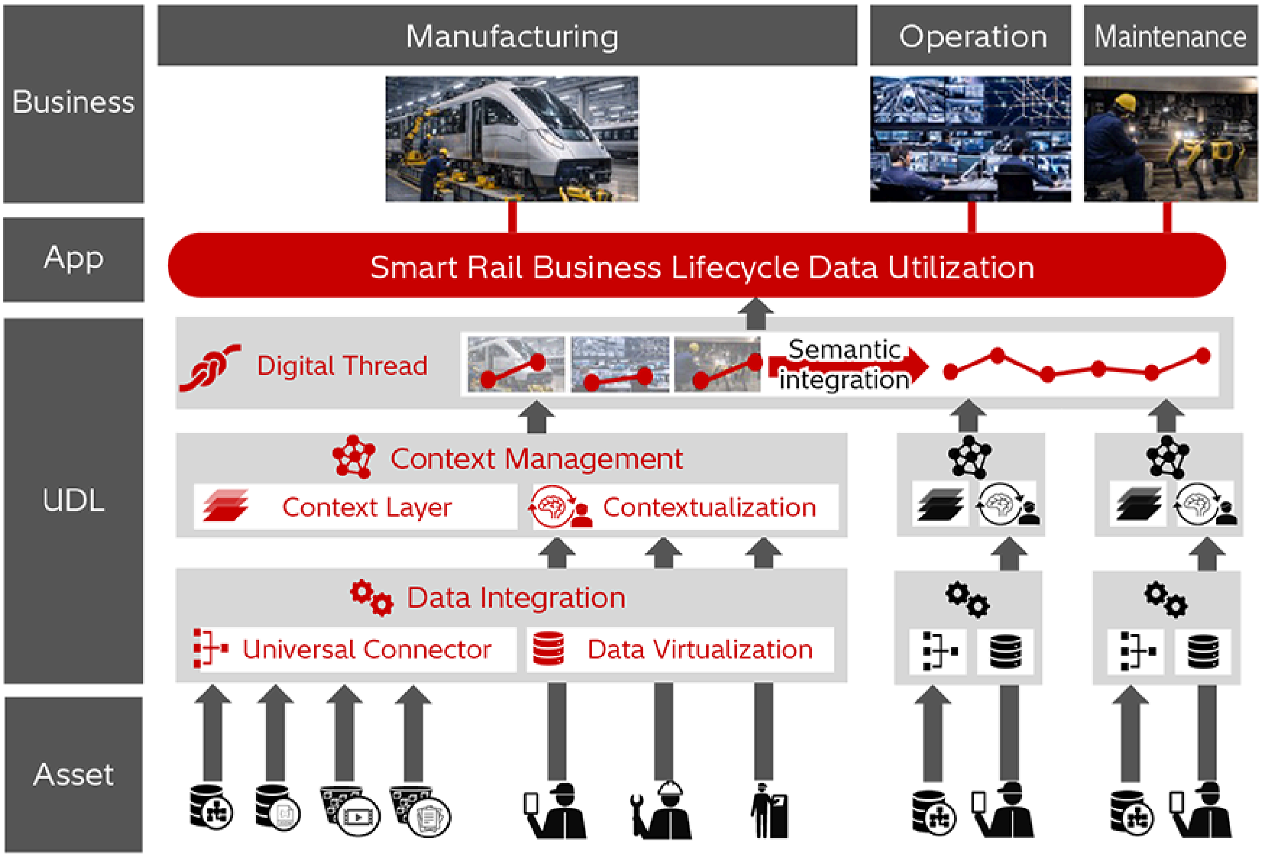

[10] Digital Thread Connecting Multiple Rail Businesses via UDL