Press Release
For Immediate Release

HITACHI

Hitachi develops guardrail technology for the safe operation of Physical AI in social infrastructure

Seamlessly supports safety analysis during design and the runtime verification of AI-generated control commands against safety constraints

Tokyo, September 14, 2026 - Hitachi, Ltd. (TSE:6501, "Hitachi") has developed guardrail technology for the safe operation of physical AI in social infrastructure, including the industrial, energy, and mobility sectors. During design, the technology uses AI and a knowledge graph[*1] to support safety experts in conducting preliminary hazard analysis and identifying candidate safety constraints. During operation, a control guardrail verifies and constrains AI-generated control commands in real time on the basis of these safety constraints, thereby providing seamless support for the continued safe operation of physical AI. Hitachi conducted a preliminary evaluation using publicly available design specifications[*2] and a simulation of an autonomous mobile system[*3] operating in an environment where automated machines and people coexist. The simulation results confirmed that autonomous control using the technology satisfied the configured safety constraints while reducing unnecessary stops by 89%. This technology is intended to help enable the application of physical AI at mission-critical sites.

Going forward, Hitachi will verify the technology using use cases and operating conditions that reflect actual operations, positioning it as an OT safety technology supporting the Integrated World Infrastructure Model (IWIM)[*4], an intelligent foundation model for social infrastructure. Hitachi will also enhance the technology to support the evolution of HMAX by Hitachi (HMAX), a portfolio of next-generation solutions that transforms social infrastructure through AI, and aims to deploy it in mission-critical social infrastructure across the industrial, energy, and mobility sectors.

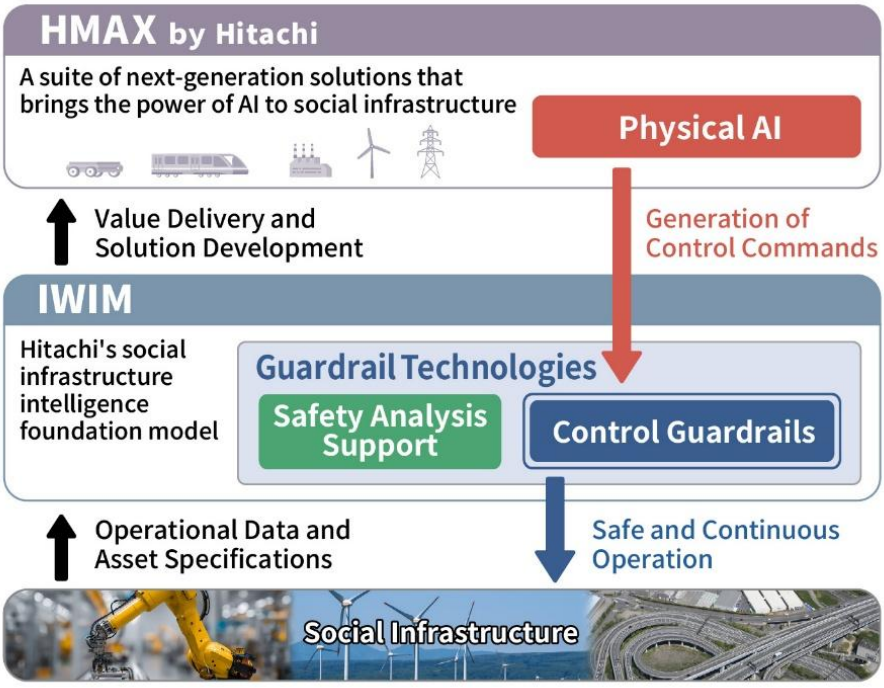


Figure 1. Positioning and overview of guardrail technology for applying Physical AI to social infrastructure

*1 Knowledge graph: A data structure that organizes and represents, as a network, information on the equipment, functions, and controls that constitute the target system, together with the relationships among them.
*2 Publicly available design specifications: Publicly available control-system specifications and design information used for evaluation. In this preliminary evaluation, Hitachi evaluated the potential of the technology to support safety analysis using public documents containing information required for control-system safety analysis.
*3 Autonomous mobile system: Vehicles, transport equipment, robots, and other systems that use sensors and similar devices to perceive their surroundings and move autonomously under the control of a control system.
*4 An intelligent foundation model announced by Hitachi in November 2025. Developed on the basis of Hitachi's multi-layered operational technology (OT) knowledge accumulated over many years and its deep understanding of physical phenomena, the model is designed to implement physical AI safely and reliably in mission-critical social infrastructure. It provides integrated support for the design, construction, and operation of equipment and systems through appropriate decision-making and control, even under complex conditions, thereby enabling their advancement.

Background and issues

In the social infrastructure sector, expectations are growing for physical AI that analyzes and makes decisions based on field data and translates those decisions into actions, such as controlling robots and automatically operating machinery and equipment. This demand is being driven by labor shortages, changes in field environments, and the need for more efficient operations. At mission-critical sites in the industrial, energy, and mobility sectors, however, it is not sufficient only to improve AI decision-making performance. AI-generated control commands must also be verified at runtime, and operations must be continued while avoiding unnecessary stops as far as possible. Conventional control architectures incorporating AI can reduce operational continuity because of unnecessary stops and the recovery work they require. In addition, preliminary hazard analyses and safety constraints must be organized from specifications and structural information during the early design stage. Because the relevant documents are diverse and extensive, this work places a heavy burden on safety experts.

Features of the technology developed to solve these issues

Hitachi has therefore combined its long-standing OT[*5] expertise with AI to develop guardrail technology for the safe operation of physical AI as a highly reliable foundational technology supporting IWIM. The technology seamlessly supports processes from safety analysis during design to the runtime verification of AI-generated control commands against safety constraints. Its key features are as follows.
*5 OT: Operational Technology. Technologies relating to equipment and systems used to control field equipment, including factories and power facilities.

1. AI technology that supports safety analysis in the early design stage

The technology organizes relationships among equipment, functions, and controls from extensive system specifications and structural information into a knowledge graph that AI can readily interpret. Based on these relationships, it supports safety experts in conducting preliminary hazard analysis and identifying candidate safety constraints. With expert review as a prerequisite, this supports the initial safety-analysis work required during the early design stage. It also supports the examination of candidate safety-constraint definitions by presenting responsible safety experts with analysis results whose source references can be verified.

2. Control guardrail technology that verifies and constrains Physical AI control commands at runtime

The technology defines the safety constraints identified during design as evaluation conditions based on stopping margin[*6], speed and acceleration, distance to obstacles, and other factors. At runtime, it evaluates in real time whether Physical AI-generated control commands satisfy these conditions, using a Control Barrier Function (CBF)[*7]. If a command does not satisfy the conditions, the technology enforces the safety constraints with the minimum intervention necessary, while preserving the AI's intended action as far as possible. This helps avoid both

approaching a hazardous state and reduced operational efficiency caused by unnecessary stops, thereby supporting continued safe operation.

*6 Stopping margin: An indicator of whether the current state of a vehicle, machine, or item of equipment remains within a range in which the safety constraints can be satisfied when the required stopping action is performed.

*7 Control Barrier Function (CBF): A method that mathematically represents the range within which a system can operate safely and evaluates and adjusts control commands so that the system state remains within that range.

3. Multi-layer protection architecture balancing safety and operational continuity

The technology employs Hitachi's proprietary multi-layer protection architecture*8, comprising an AI planner that generates control commands, a control guardrail*9 that verifies and constrains the commands at runtime using a CBF, and an independent safety system that performs the final stop separately from the AI. During normal operation, the system uses the AI's flexible commands. If a violation of the safety constraints is predicted, the control guardrail enforces the safety constraints with the minimum intervention necessary. If such intervention is not feasible, control switches to the independent safety system, which performs an emergency stop. This architecture suppresses unnecessary stops while maintaining the final safety function, thereby supporting both the safety and operational continuity required for social infrastructure.

*8 Patent pending by Hitachi. The patent-application status stated here is current as of the date of this announcement. Please note that the status of patents and related rights may subsequently change as a result of third-party invalidation proceedings, prosecution, and other circumstances.

*9 Control guardrail: A protective function that keeps AI-generated control commands within a range in which they can be executed safely.

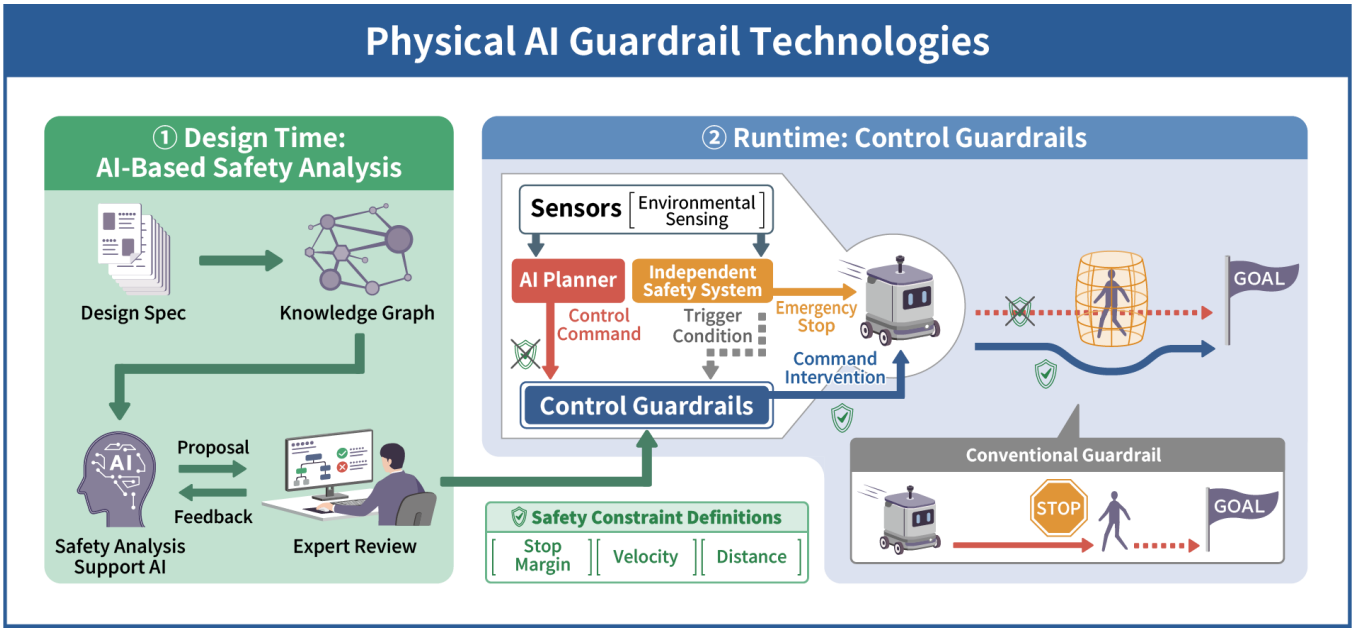


Figure 2. Overview of guardrail technology for the safe operation of Physical AI

Confirmed results

In the preliminary evaluation using publicly available design specifications, the technology identified 20 of 22 major hazards*10, indicating its potential to support initial safety analysis. In the simulation of an autonomous mobile system operating in an environment where automated machines and people coexist, Hitachi also confirmed an 89% reduction in unnecessary stops compared with a conventional emergency-stop control method, using an architecture in which the final stop is performed by a safety system independent of the AI. The results indicate that autonomous control using the technology can avoid unnecessary stops and improve operational continuity while satisfying the configured safety constraints*11 in simulation.

*10 Major hazards: Hazardous factors that could lead to serious consequences, such as accidents involving the controlled system, equipment damage, or service interruption. The evaluation examined whether the technology could support initial safety analysis, subject to final expert review.

*11 Safety constraints: In this verification, the constraints were physical constraints configured in the simulation, including stopping margin, speed and acceleration, and distance to obstacles. The results do not demonstrate compliance with standards or fulfillment of all safety requirements for an actual system.

Looking ahead

To support the application of physical AI to social infrastructure, Hitachi will position this technology as an OT safety technology constituting IWIM and will advance verification using use cases and operating conditions that reflect actual operations. Hitachi will also enhance the technology to support the evolution of HMAX as an OS*12 for autonomously controlling and optimizing social infrastructure as a whole, aiming to deploy it in mission-critical social infrastructure across the industrial, energy, and mobility sectors and contribute to the use of physical AI that balances system safety with operational continuity.
Part of this work is scheduled to be presented at the IEEE International Conference on Intelligent Transportation Systems (IEEE ITSC 2026), to be held in Naples, Italy, from September 15 to 18, 2026. The technology is also scheduled to be exhibited at the International Trade Fair for Transport Technology (InnoTrans 2026), to be held in Berlin, Germany, from September 22 to 25, 2026.

*12 OS: An abbreviation for Operating System. Foundational software that manages and controls the operation of an entire computer or system and enables various applications and devices to work together.

Trademark Notice: All trademarks and product names are the property of their respective owners.

About Hitachi, Ltd.
Through its Social Innovation Business (SIB) that brings together IT, OT(Operational Technology) and products, Hitachi aims to be a global leader in continuously transforming social infrastructure through digital, contributing to a harmonized society where the environment, wellbeing, and economic growth are in balance.
Hitachi operates worldwide across four sectors – Digital Systems & Services, Energy, Mobility, and Connective Industries – as well as a Strategic SIB Business Unit focused on new growth areas. With Lumada at its core, Hitachi creates value by combining data, technology and domain knowledge to solve customer and social challenges. Revenues for FY2025 (ended March 31, 2026) totaled 10,586.7 billion yen, with 606 consolidated subsidiaries and approximately 290,000 employees worldwide. Visit us at www.hitachi.com.

Contact
Hitachi, Ltd.
Research & Development Group
https://www8.hitachi.co.jp/inquiry/hitachi-ltd/hqrd/news/en/form.jsp