Deloitte.

Together makes progress

The Physical AI Dossier

A selection of high-impact use cases across seven major industries

GLOBAL AI & EMERGING MARKETS

Note: *Artificial intelligence regulation is continuing to evolve across the world. The use cases presented in this deck are illustrative and are intended to support discussion of potential applications of physical artificial intelligence. Before deploying any artificial intelligence system, organizations should assess the regulatory requirements that may apply to their specific use case and implementation context and consider any corresponding safeguards that may need to be implemented.*

Foreword

Artificial intelligence is expanding from the screen into the physical world.

A new generation of AI systems can now perceive physical environments, reason about them, and take action within them. Physical AI is not a distant prospect; it is actively being deployed in factories, warehouses, utility networks, hospitals, farms, city streets, and homes. And the pace of adoption is accelerating.

This dossier features use cases across six major industries—Consumer; Energy, Resources & Industrials; Financial Services; Government & Public Services; Life Sciences & Health Care; and Technology, Media & Telecommunications—as well as a chapter of use cases that apply broadly across many industries.

For each industry, how Physical AI is being used, or may soon be used, to address operational challenges, improve safety and reliability, and create new sources of value is explored. The use cases span the range of Physical AI applications and form factors, including autonomous mobile robots (AMRs), drones, humanoid robots, autonomous vehicles, quadrupeds, and task-specific machines.

At the frontier of this evolution are dark factories—highly autonomous operating environments where Physical AI enables systems to run continuously with minimal human presence under governed oversight—demonstrating that Physical AI is not a distant prospect, but an emerging operational reality.

Deploying Physical AI at scale is not simply a technology challenge. It requires reimagining how work gets done, how humans and machines collaborate, and how accountability is defined when autonomous systems act on an organization's behalf. The following pages address where Physical AI stands today, where it's headed, and the governance principles that should guide its responsible deployment.

The goal is to help business and government leaders assess where Physical AI is most relevant to their organizations, understand what successful deployment actually requires, and build the strategic perspective to act with both ambition and discipline.

Nitin Mittal
Global AI Leader
Deloitte LLP

Jim Rowan
Deloitte US, Head of AI
Deloitte Consulting LLP

Sam Roddick
Chief Strategy Officer, Global Consulting Services
Deloitte Global

Chris Lewin
AP Strategic Growth Offering Leader
Deloitte Singapore T&T Pte Ltd.

Cross Industry Physical AI Dossier

Summary: Physical AI Use Cases For Simulation

Simulation is the control plane for Physical AI

Simulation and digital twins are foundational enablers of Physical AI. They provide a safe, scalable environment where intelligent machines can be designed, trained, validated, and governed before acting in the real world[1].

Simulation creates high-fidelity virtual environments—combining physics-based modeling, synthetic data, and AI-driven scenarios—to test how robots, autonomous systems, and physical infrastructure behave under real operating conditions, including rare and hazardous edge cases. Digital twins extend this capability into live operations by continuously mirroring physical assets, environments, and workflows, to help enable closed-loop learning between virtual and physical systems.

Together, simulation and digital twins reduce deployment risk, accelerate development cycles, and lower the cost of experimentation—while improving safety, reliability, and regulatory confidence. This simulation-first approach applies across industries, from warehouses and factories to utilities, networks, healthcare facilities, and financial infrastructure.

As Physical AI scales, simulation is no longer a design-phase tool. It becomes shared, enterprise-wide infrastructure that turns experimental automation into production-ready systems—to help enable organizations move faster with confidence, not caution.

Solution: Digital Twin + Simulation

Simulation-first development and digital twins (1/2)

Validating physical systems virtually before real-world deployment

DESCRIPTION

Simulation environments, synthetic data, and digital twins underpin the design, training, testing, validation, and certification support of Physical AI systems, including machines, robots, and vehicles, prior to real-world deployment. Simulation platforms help enable iterative, consequence-free learning that helps organizations identify risks, accelerate prototyping, significantly reduce development costs and timelines, and improve the operational readiness of robotic and autonomous systems before deployment.

ISSUE/OPPORTUNITY

Physical AI systems operate in environments where failures cause physical harm, making real-world learning, testing, and iteration prohibitively risky without simulation-based validation.

Testing autonomous robots, vehicles, or industrial equipment in live environments exposes workers and equipment to risk, limits the range of scenarios that can be safely evaluated, and makes it difficult to reproduce rare or hazardous edge cases consistently. Inconsistent risk analysis methods across development teams further complicate regulatory approval, with safety assessments varying in depth and documentation quality.

The opportunity is to shift most development, training and validation work into simulation, where AI systems can be tested against thousands of scenarios—including edge cases that would be impractical or unsafe to recreate physically, before hardware is deployed.

HOW PHYSICAL AI CAN HELP

Scenario generation at scale

AI creates high-variance test conditions covering normal operations, edge cases, equipment failures, and hazardous scenarios that would be difficult, costly, or dangerous to replicate in physical environments.

Evidence generation for compliance

Simulation outputs—scenario logs, performance metrics, failure mode analyses—provide structured documentation that supports regulatory submissions and certification processes across industries and jurisdictions.

Regression testing for updates

Repeatable simulation-based test suites verify that software updates or model changes do not degrade existing capabilities, to help enable continuous development without requiring full physical re-validation.

Standardized risk analysis workflows

AI tools structure hazard identification and risk assessment consistently across development teams, reducing variability in safety documentation and supporting more predictable regulatory review processes.

Design validation before build

Digital twins help enable testing of system designs and software changes against virtual representations of physical equipment and environments, identifying issues before physical trials and reducing the need for hardware modifications.

Reduced real-world trial burden

Fewer physical experiments are required when simulation has already validated system behavior across a broad scenario space, lowering development costs and reducing exposure to test-related incidents.

POTENTIAL BENEFITS

Faster development cycles

Shorter iteration loops because design flaws and performance gaps are identified in simulation rather than through physical testing.

Reduced deployment risk

Systems reach physical deployment with broader validation coverage and fewer untested failure modes.

Smoother approvals

Better-structured documentation reduces back-and-forth with regulators and accelerates certification timelines.

Lower redesign costs

Earlier detection of issues reduces expensive hardware modifications and late-stage rework.

Simulation-first development and digital twins (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Simulation environments used to validate Physical AI must be sufficiently faithful to real-world conditions to provide meaningful safety assurance. A digital twin that fails to represent environmental variability, sensor noise, or rare edge cases generates validation evidence that overstates readiness—potentially advancing systems to deployment with untested failure modes.

Responsible and accountable

Regulatory authorities can be expected to require accountability for the adequacy of simulation coverage—including how edge cases were selected, how real-world fidelity was validated, and what simulation limitations were acknowledged. These governance records should be maintained throughout the system lifecycle and sufficient to support regulatory certifications and examinations.

Transparent and explainable

Engineers, safety assessors, and regulators reviewing simulation-based validation evidence need to understand what scenarios were tested, how they were generated, what failure modes were covered, and where simulation limitations may affect conclusions. Opaque simulation methodologies can undermine the credibility of certification evidence and complicate regulatory review of Physical AI deployments.

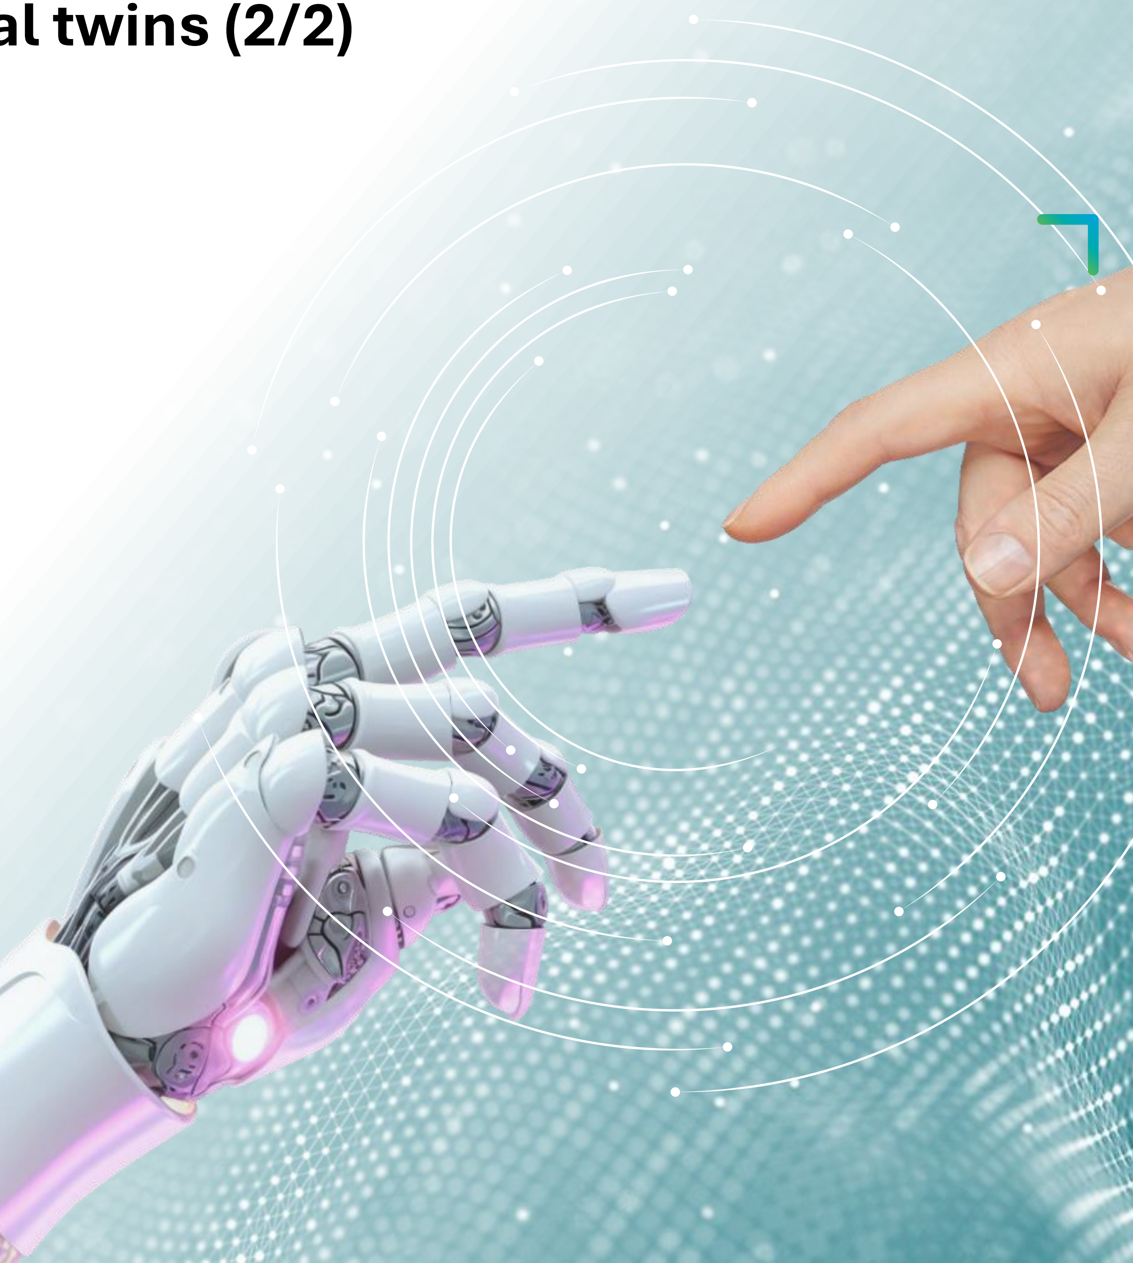

Solution: Digital Twin + Simulation

Simulation-trained, human-supervised closed-loop remediation (1/2)

Safe, simulation-trained intervention for physical systems

DESCRIPTION

Physical AI systems monitor and reason over live physical infrastructure—networks, grids, plants, and facilities—using simulation-trained models to recommend and execute corrective actions. Human-in-the-loop controls validate changes before physical interventions occur, enabling safe, adaptive remediation across safety conscious environments.

ISSUE/OPPORTUNITY

Physical AI systems—robots, drones, autonomous vehicles, and intelligent infrastructure—must operate reliably in complex, safety-conscious environments. However, testing and training directly in live environments is costly, disruptive, and risky, while real-world edge cases are difficult to anticipate. Organizations across industries face growing pressure to scale Physical AI quickly without compromising safety, compliance, or operational continuity. The opportunity lies in using simulation and digital twins to shift experimentation, learning, and validation into virtual environments—allowing Physical AI systems to mature faster while keeping humans accountable for final decisions.

HOW PHYSICAL AI CAN HELP

Simulation-trained intervention policies

AI models are trained in digital twins to learn safe remediation actions before being allowed to recommend changes in live physical systems.

Edge-based physical signal interpretation

Physical AI fuses telemetry from sensors, cameras, meters, and equipment controllers to detect anomalies that software-only monitoring would miss.

Closed-loop learning with digital twins

Performance data from live environments feeds back into digital twins, continuously improving models, policies, and predictions over time.

Approval-based execution

Humans approve changes before implementation, maintaining accountability and enabling operators to reject recommendations when local knowledge suggests alternative actions.

Action impact prediction before execution

Proposed remediation steps are stress-tested in simulation to predict downstream physical effects (safety, stability, service impact) before approval.

POTENTIAL BENEFITS

Faster resolution

Routine fixes occur sooner as AI presents validated solutions immediately rather than requiring operators to research procedures and manually configure changes through multiple system interfaces.

Lower operator load

Manual effort declines as operators shift from diagnosis and solution development to review and approval, enabling smaller control room teams to manage larger grid footprints.

Stronger governance and trust

Human-in-the-loop validation supports regulatory compliance and builds organizational confidence in Physical AI systems.

Trust preservation

Risk remains controlled through mandatory human approval, addressing regulatory requirements and organizational concerns while still capturing efficiency benefits from AI assistance.

Transferable remediation patterns across industries

Once trained, remediation logic can be adapted across utilities, telecom, manufacturing, and infrastructure with minimal rework.

Simulation-trained, human-supervised closed-loop remediation (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Simulation-trained remediation models must perform reliably when applied to live physical infrastructure; however, the real world inevitably presents conditions not fully captured in digital twins. A model that recommends incorrect remediation actions in live safety-critical environments because its training did not represent actual operating conditions creates the very failures it was designed to prevent.

Responsible and accountable

Human approval is required for all changes before implementation, preserving accountability for safety-critical interventions. However, this governance design principle must be enforced in practice. Effective accountability requires clear and consistent documentation of what the AI recommended, what simulation stress-testing showed, who approved the action, and what outcome occurred.

Transparent and explainable

Operators approving AI-recommended remediations need to understand what anomaly was detected, what action was proposed, and what downstream effects were predicted by simulation and stress-testing. True human oversight is only possible if AI's reasoning and outputs are transparent and explainable.

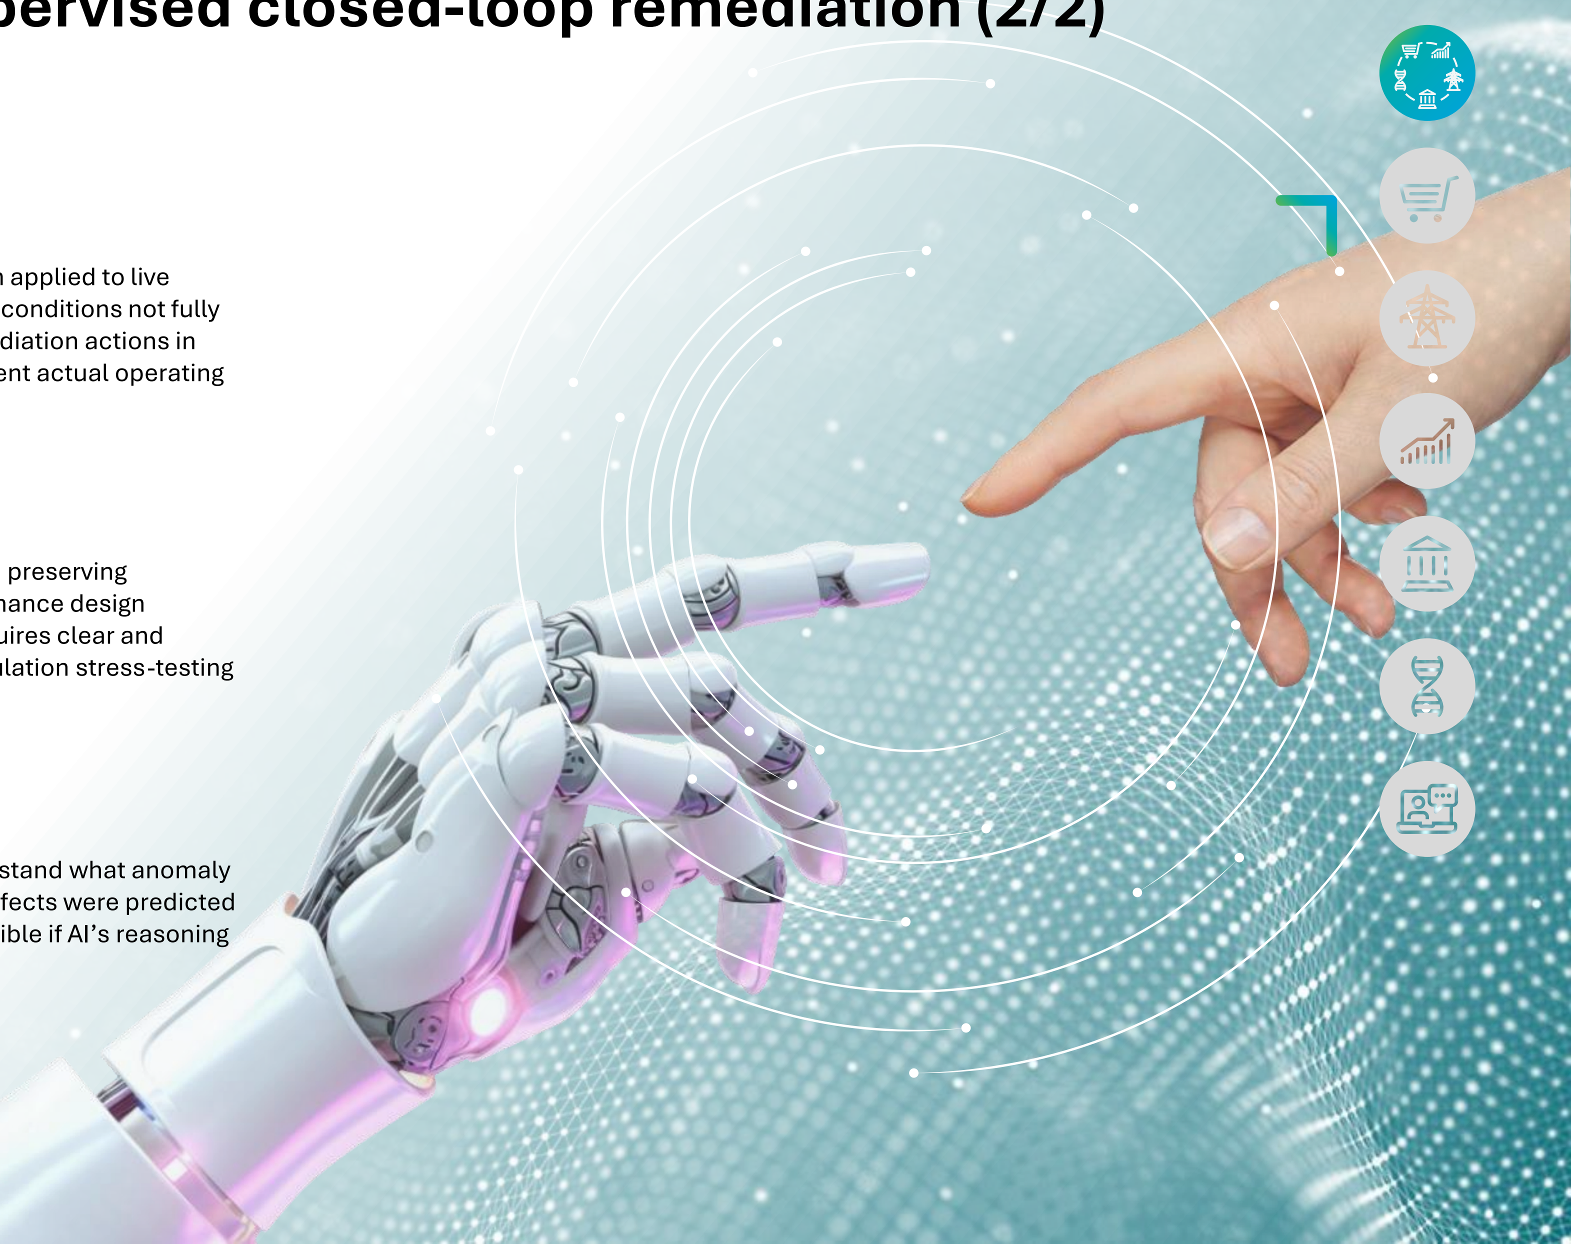

Simulation-driven remote operations and training (1/2)

Scaling expertise through digital twins and immersive simulation

DESCRIPTION

AR/VR digital twins create high-fidelity virtual replicas of offshore platforms and production facilities, to help enable specialists to perform remote troubleshooting, technical oversight, and operator training from onshore locations without traveling to hazardous sites.

ISSUE/OPPORTUNITY

Operating and maintaining complex, hazardous, or geographically remote physical assets traditionally requires expert personnel to be physically present on site. This dependence increases safety risk, travel cost, downtime, and delays—particularly when specialist availability, weather conditions, or access constraints limit rapid response. Training new operators is similarly constrained, as hands-on learning in live environments is costly, slow to scale, and exposes people and assets to operational risk. Traditional remote monitoring tools lack the fidelity and interactivity needed to support effective troubleshooting, skill transfer, and decision-making for complex physical systems. The opportunity is to shift oversight, training, and validation into high-fidelity simulation and digital twins—enabling safe, scalable knowledge, faster decision-making, and reduced physical exposure without compromising control or accountability.

HOW PHYSICAL AI CAN HELP

Simulation-first skill transfer and validation

AI uses high-fidelity digital twins to model assets, procedures, and failure scenarios, enabling operators and specialists to train, rehearse, and validate actions in a risk-free virtual environment before interacting with live systems.

Knowledge capture and replay

Simulation environments encode expert decision paths, diagnostic logic, and safe operating envelopes, allowing scarce expertise to be reused consistently across locations and shifts without requiring physical presence.

Remote decision support with contextual awareness

AI continuously synchronizes simulation models with live asset data, enabling remote specialists to reason about current conditions, test interventions virtually, and provide guidance grounded in predicted physical outcomes.

Governed human-in-the-loop operations

AI supports recommendations and scenario evaluation, while humans retain approval authority for physical actions—preserving accountability, safety, and regulatory alignment.

Simulation-trained action policies

Physical AI systems learn safe operating envelopes in simulation before executing actions on real assets, reducing reliance on trial-and-error in hazardous environments.

POTENTIAL BENEFITS

Transformational safety improvement

Keeping personnel out of hazardous offshore environments for routine technical service and troubleshooting represents the primary value driver, reducing exposure to the safety risks inherent in offshore operations.

Major cost efficiency gains

Eliminating the majority of specialist travel to remote sites for technical service and troubleshooting removes significant aviation and logistics costs that accumulate across large offshore portfolios.

Faster workforce development

Operator training and onboarding conducted in onshore back offices rather than on offshore platforms accelerates skills development, reduces training logistics costs, and helps enable more flexible scheduling of new hire cohorts.

Proven deployment

Digital twin remote operations technology has already been fully deployed and validated at scale in offshore operations, representing a mature capability with established return on investment rather than an emerging proof of concept.

Scalable knowledge transfer

Simulation-trained AI encapsulates expert knowledge, allowing consistent execution across assets without depending on scarce specialists.

Cross-industry applicability

Applicable wherever Physical AI systems operate in safety conscious, remote, or complex environments—including ER&I, logistics, and healthcare facilities, where simulation-first validation improves safety, reliability, and operational confidence.

Simulation-driven remote operations and training (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Digital twins used for remote troubleshooting must faithfully represent current physical conditions. A simulation that diverges from the live system state due to synchronization failures can lead remote experts to recommend interventions based on inaccurate virtual representations. Reliable synchronization between the digital twin and live physical data is a prerequisite for safe remote operations.

Responsible and accountable

In this use case, humans retain approval authority for all physical actions. AI supports recommendations but does not act autonomously. To maintain accountability, organizations need to document what the simulation predicted, what the remote expert recommended, and what action was approved and taken.

Transparent and explainable

Remote experts and operators using digital twins for troubleshooting need to understand how the simulation represents current physical conditions, what assumptions underlie predicted outcomes, and where simulation limitations may affect recommendation reliability.

Summary: Physical AI Use Cases Across Industries

Some of the most powerful Physical AI capabilities are horizontal—defined not by industry but by the operational problem they solve

Physical AI is difficult to define within sector boundaries. Many of the most impactful capabilities in this dossier are defined not by the industry they serve but by the operational challenge they address, and those challenges recur across industries in forms similar enough that the same underlying approach can be deployed broadly.

The cross-industry use cases collected here share a common characteristic: the problem they solve is sufficiently universal, and the solution sufficiently transferable, that limiting them to a single industry chapter would understate their relevance. For example, quality assurance, logistics automation, human-machine interaction in physical environments, and operational simulation are challenges that companies in many different industries face in similar forms. A capability developed to solve such a problem in one industry frequently transfers to others with meaningful adaptation but without reinvention.

This horizontal applicability changes the economics of Physical AI investment. Organizations that recognize the phenomenon can leverage implementations across multiple business units, geographies, and functions—compounding the return on their investments and building institutional capabilities that extend well beyond single deployment. Horizontal applicability also suggests a different framing for how leaders should evaluate Physical AI opportunities: not just "what is my industry doing?" but "what operational problem am I trying to solve, and where has it already been solved in a way I can learn from and adapt?"

As Physical AI matures, the boundary between industry-specific and cross-industry capabilities may continue to blur. The use cases in this chapter represent the leading edge of that convergence.

Solution: Edge + Robotics | Task-specific

Autonomous self-calibrating quality and process control (1/2)

From defect detection to self-maintaining, defect-preventing production systems

DESCRIPTION

Physical AI systems combine advanced vision, sensing, and closed-loop process control to continuously monitor production quality and their own operational performance. These systems detect defects, process drift, sensor misalignment, and environmental changes in real time, and autonomously recalibrate sensors, retrain models, or adjust equipment parameters to prevent defect propagation—maintaining accuracy and stability without relying on periodic human intervention.

ISSUE/OPPORTUNITY

Traditional quality systems detect defects after production, resulting in scrap, rework, batch losses, and delayed root-cause analysis. At the same time, Physical AI deployments themselves degrade over time as cameras shift, sensors drift, lighting changes, and thermal conditions evolve—requiring frequent manual recalibration and maintenance. These gaps create production inefficiencies, quality risk, and ongoing operational overhead. The opportunity is to move beyond reactive inspection and manual upkeep toward Physical AI systems that both prevent defects and self-maintain performance, sustaining quality and reliability continuously in dynamic production environments.

HOW PHYSICAL AI CAN HELP

Real-time defect analysis and closed-loop control

Vision systems identify emerging quality issues and automatically adjust process parameters (e.g., temperature, pressure, speed, positioning) before defects spread.

Integrated inspection across stages

Quality data is shared across process steps, enabling upstream adjustments based on downstream signals and optimizing the full production line.

Self-monitoring and drift detection

AI continuously evaluates its own accuracy, false-positive rates, and environmental inputs to detect performance degradation early.

Automated calibration and correction

Systems autonomously recalibrate cameras, sensors, and models when misalignment or drift is detected, without stopping production.

Predictive quality modeling

Machine-learning models forecast defect risk based on process variables, enabling proactive intervention rather than reactive correction.

POTENTIAL BENEFITS

Reduced scrap and rework

Preventing defects at the source lowers material waste, batch losses, and costly rework cycles.

Sustained accuracy and reliability

Self-calibrating systems maintain consistent performance over time, avoiding degradation between maintenance cycles.

Minimized downtime and production disruption

Automated correction prevents failures that would otherwise require systems to be taken offline.

Lower maintenance overhead

Reduced reliance on manual calibration and inspection frees specialized staff for higher-value work.

Continuous quality improvement

Closed-loop learning progressively tightens process control beyond human-achievable consistency.

Stronger auditability and validation

Quality and process data supports regulatory audits and customer certifications.

Autonomous self-calibrating quality and process control (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Systems autonomously recalibrating sensors, retraining models, and adjusting process parameters must do so reliably. An incorrect self-calibration that degrades detection accuracy can allow defects to propagate undetected. The self-monitoring capability must itself be monitored and validated to ensure it correctly identifies genuine drift rather than triggering inappropriate recalibrations that disrupt production.

Responsible and accountable

Autonomous process adjustments without human review can adversely affect output quality and equipment behavior. Organizations must define which adjustments fall within pre-validated safe operating envelopes and which require engineer validation before execution—preventing unreviewed AI decisions from affecting production quality or equipment integrity at scale.

Transparent and explainable

To validate that corrections are appropriate and detect when self-calibration logic responded to a spurious signal, process engineers need to understand what triggered the self-calibration, what adjustment was made, and what outcome resulted. Systems that self-correct opaquely are difficult to audit, certify, and troubleshoot when quality issues arise in production.

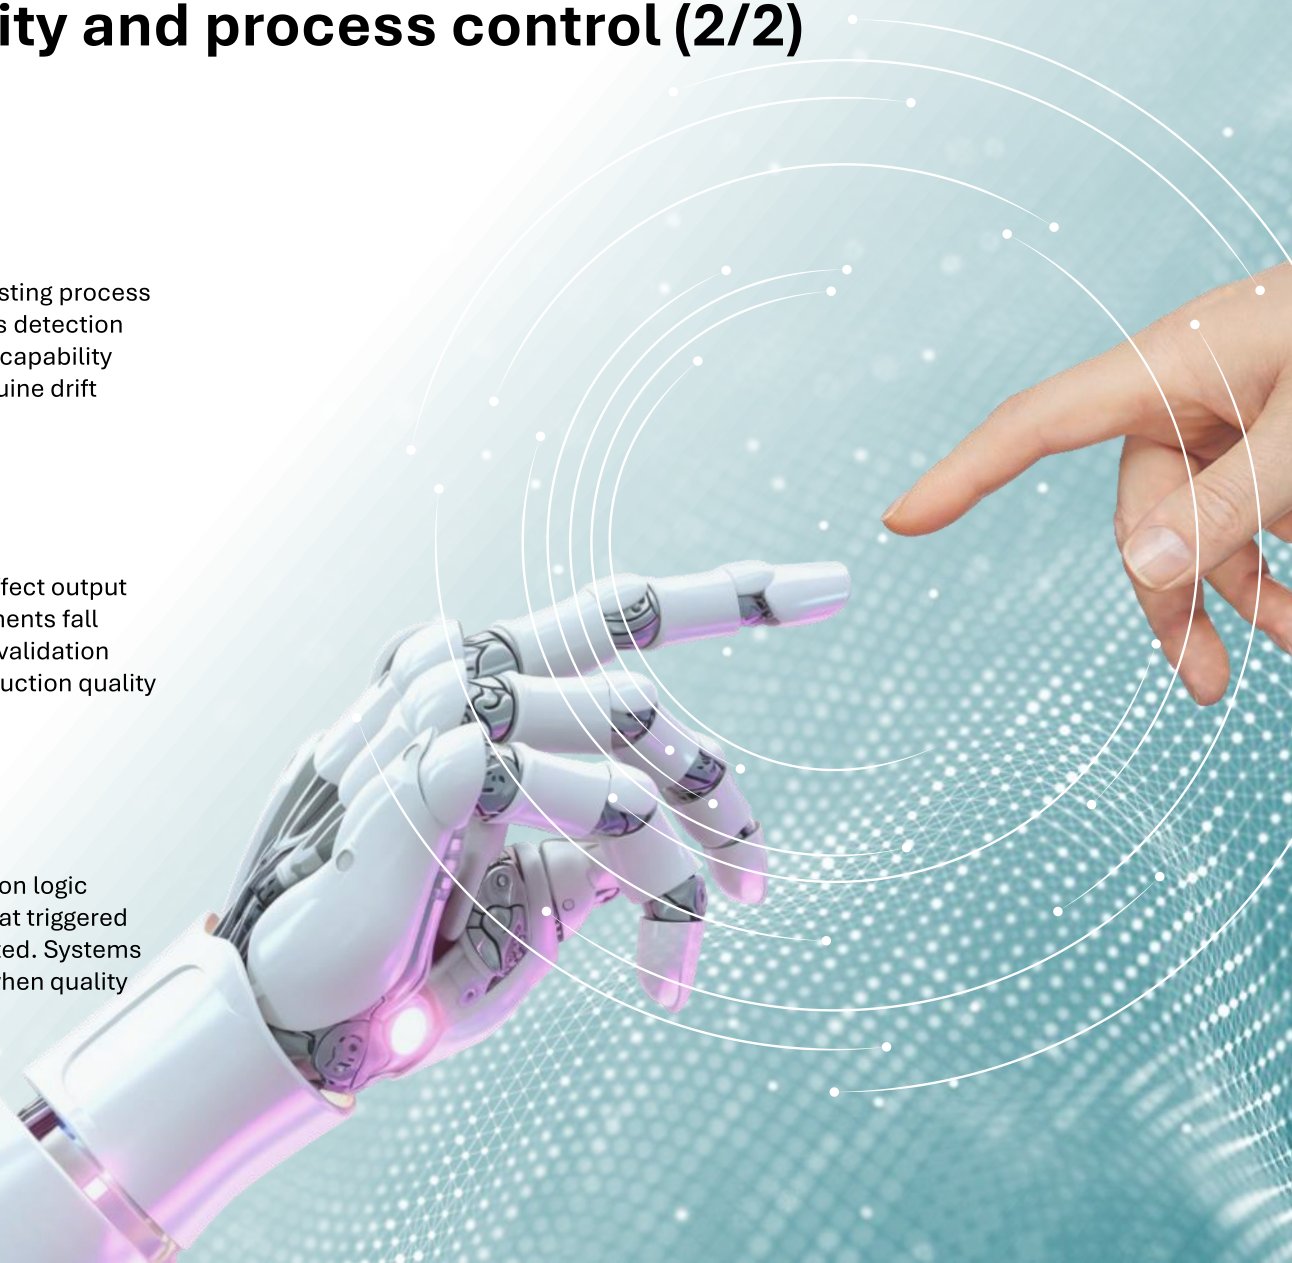

Solution: Edge + Robotics | Autonomous mobile robots (AMRs)

Semi-autonomous warehouse loading and unloading robotics (1/2)

Human-supervised Physical AI for dynamic material handling

DESCRIPTION

Mobile robots assist with unloading containers, rearranging goods, and optimizing pallet placement within warehouse environments under supervised operation, increasing throughput while maintaining human fallback capability.

ISSUE/OPPORTUNITY

Warehouse loading and unloading are physically demanding, repetitive, and throughput-sensitive processes. Manual handling introduces variability and ergonomic risk while limiting scalability during peak demand. Workers manually unload shipping containers, lift heavy boxes onto pallets, and stack goods in warehouse locations, performing physically taxing work that causes injuries, fatigue, and high turnover.

During peak seasons or promotional periods, warehouses struggle to find sufficient labor to process increased container volumes, creating bottlenecks that delay inventory availability and frustrate customers expecting rapid delivery. Pallet stacking quality varies by worker skill and fatigue level, leading to inefficient space utilization when goods are stacked loosely or unstably.

Traditional fixed automation often requires costly infrastructure redesign. Fixed conveyor systems and automated storage require extensive facility modifications and work well only for standardized products, not the diverse container contents typical of modern distribution centers.

The opportunity is to deploy Physical AI systems that can perceive, reason, and act in the physical world, while keeping humans in the loop for judgment, safety, and accountability. This enables automation of physically demanding tasks without sacrificing operational resilience or control.

HOW PHYSICAL AI CAN HELP

Perception of unstructured physical environments

Vision and sensor fusion help enable systems to identify objects, assess orientation, detect obstacles, and understand spatial constraints in real time.

Supervised autonomy controls

Robots execute predefined tasks while escalating exceptions to human supervisors when they encounter situations outside normal parameters or require judgment calls.

Physical reasoning and adaptive manipulation

AI models infer stable grasp points, load balance, and movement paths, adjusting actions dynamically as conditions change.

Environmental sensing integration

Robots detect obstacles and adjust paths in real time, navigating around workers, equipment, and temporary obstructions common in busy warehouse environments.

Fallback continuity mechanisms

Manual override capability helps to ensure uninterrupted operations during technical interruptions, help enable workers to complete tasks manually if robots experience downtime.

POTENTIAL BENEFITS

Reduced unloading time

Robotic assistance increases processing speed as machines work continuously without fatigue, enabling faster container turnover during peak periods.

Improved space utilization

Optimized stacking enhances warehouse capacity by consistently applying space-efficient pallet configurations that maximize vertical storage and minimize wasted space.

Lower ergonomic strain

Automation reduces physically intensive tasks, decreasing worker injuries from heavy lifting and repetitive motions.

Operational resilience through fallback mechanisms

Human fallback maintains continuity during outages as workers can manually perform robot tasks if equipment fails, preventing complete operational stoppage.

Cross-industry applicability

Applicability across logistics, manufacturing, healthcare facilities, construction sites, and industrial plants—wherever physical materials should be handled safely in dynamic environments.

Semi-autonomous warehouse loading and unloading robotics (2/2)

MANAGING RISK AND PROMOTING TRUST

Safe and secure

Mobile robots in shared warehouse environments must safely detect and respond to human presence—particularly during high-pressure peak periods when workers are fatigued and moving unpredictably. Validation of safety boundaries must cover actual operating conditions that include the congested, fast-paced peak periods when robot assistance is most needed and human-robot proximity is greatest.

Responsible and accountable

Supervised autonomy requires robots to escalate exceptions to human supervisors when situations fall outside normal parameters. Clear governance must define what triggers escalation, who is accountable when robot actions cause product damage or worker injury, and what logs must be maintained to support investigation and liability determination when incidents occur.

Transparent and explainable

Human supervisors handling escalations must understand what the system perceived and why it is requesting an exception so they can make informed decisions when authorizing requests. Opaque robot escalations that supervisors are unable to interpret undermine the governance model this use case depends on.

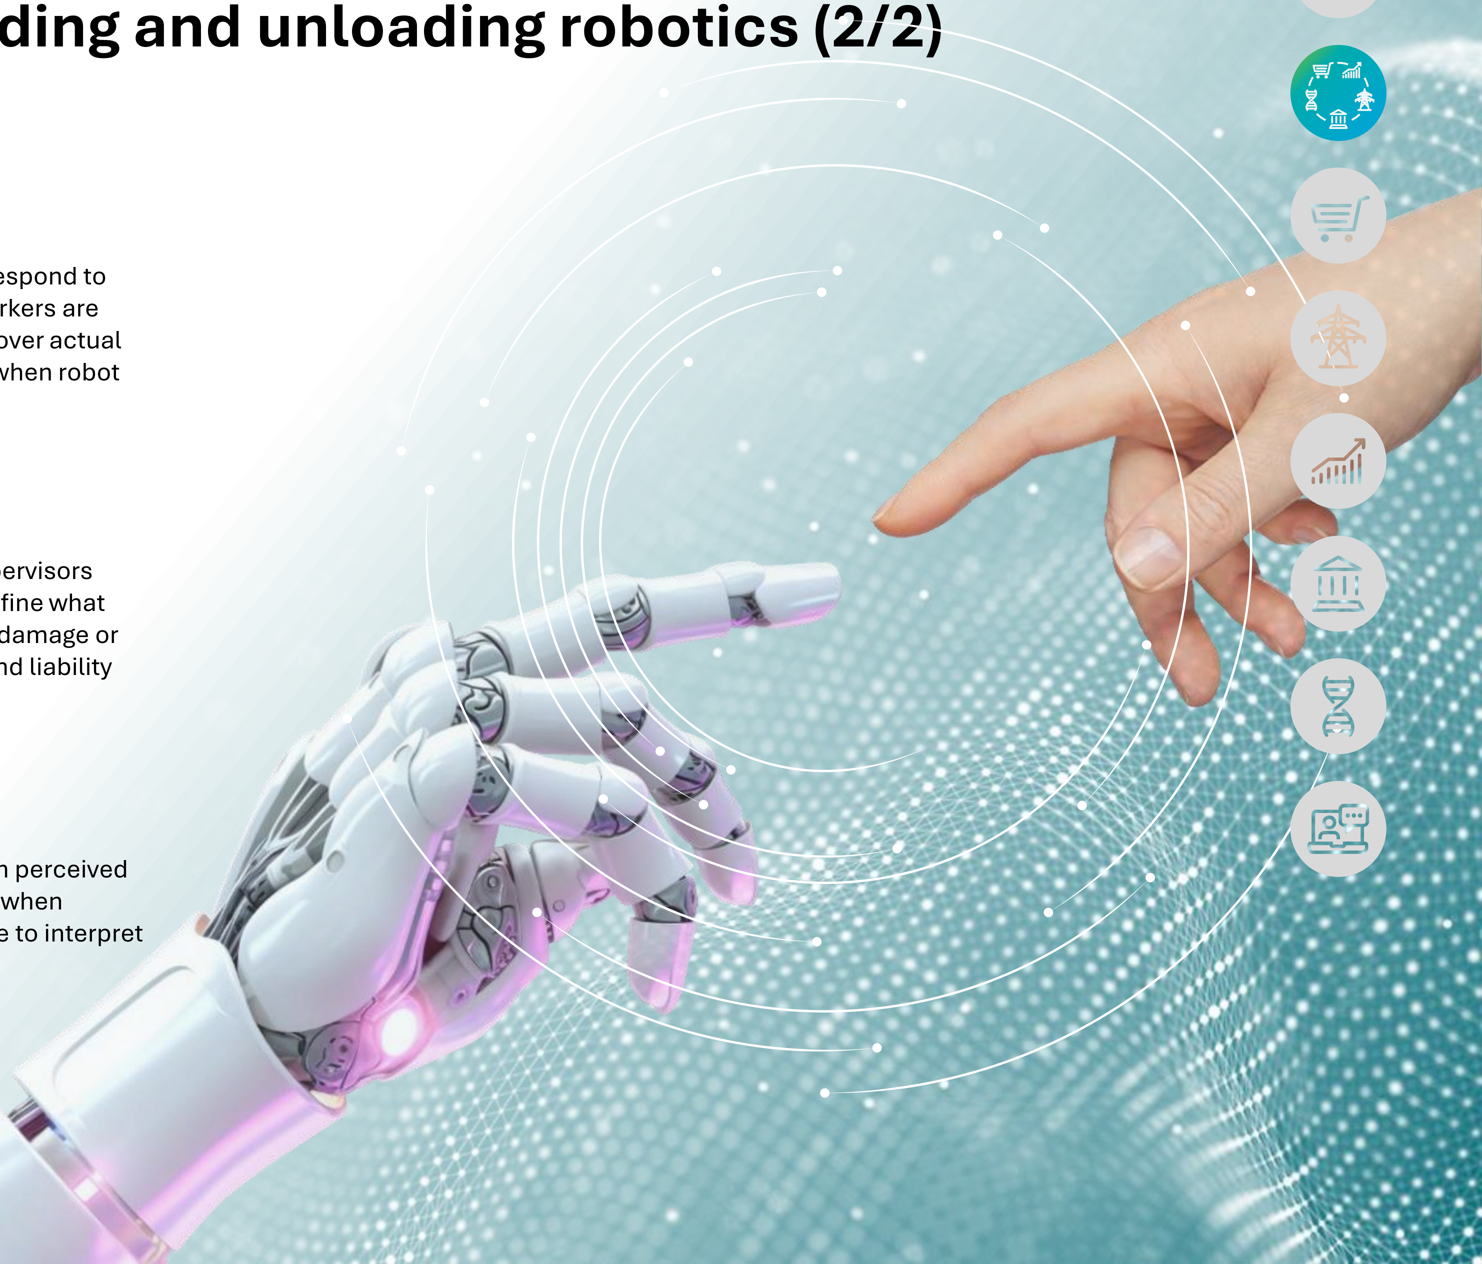

Solution: Edge + Robotics | Task-specific

Voice-controlled Physical AI assistants for industrial operations (1/2)

Hands-free human–machine collaboration in safety-critical environments

DESCRIPTION

Voice-controlled Physical AI systems designed specifically for industrial environments help enable workers to interact naturally with physical machines and robotic systems, request information, trigger actions, and receive alerts through voice commands, even in high-noise factory and field settings. These systems combine voice recognition optimized for industrial acoustics with multi-modal interfaces (voice, touch, visual) to support human-AI collaboration in hands-occupied or hazardous work environments.

ISSUE/OPPORTUNITY

Industrial workers operating machinery, conducting inspections, or performing maintenance often need to interact with AI systems and access digital information while their hands and visual attention are occupied with physical tasks. Traditional interfaces requiring screens and keyboards or touchscreens force workers to stop physical work to interact with systems, interrupting workflow, reducing efficiency, and creating safety risks when workers must divert attention from potentially hazardous operations.

Consumer voice assistants fail in industrial environments due to noise, task complexity, and lack of operational context. The opportunity is to enable safe, hands-free interaction tailored to industrial realities.

HOW PHYSICAL AI CAN HELP

Noise-robust voice recognition

Advanced speech recognition models trained specifically for industrial environments filter machine noise, mechanical sounds, and background conversations to accurately recognize worker commands in realistic production settings.

Contextual AI understanding

Natural language processing helps enable workers to ask questions and give commands conversationally rather than memorizing specific phrases, with AI understanding context from current tasks and equipment states to interpret intent correctly.

Multi-modal interaction design

Systems combine voice input with touch, gesture, and visual confirmations, allowing workers to choose the most appropriate interaction method based on immediate conditions and task requirements rather than forcing single-mode interaction.

Human-in-the-loop operation

Assistants support decision-making and execution without autonomous control, preserving human judgment and accountability.

Hands-free, safety-aware interaction

Workers access information, trigger actions, and receive alerts without stopping work or shifting attention away from safety conscious tasks.

POTENTIAL BENEFITS

Hands-free operation in critical environments

Workers can access information and control systems—and receive AI insights—without interrupting physical tasks or diverting visual attention from safety-critical work, improving both efficiency and safety in hands-occupied operations.

Reduced training and adoption barriers

Natural voice interaction lowers the learning curve for AI system adoption, enabling workers to leverage AI capabilities without extensive technical training on complex interfaces or memorizing command sequences.

Faster response to AI alerts and recommendations

Voice-based notifications and alerts reach workers immediately without requiring them to check screens, enabling faster response to quality issues, safety warnings, or process anomalies detected by AI monitoring systems.

Cross-industry applicability

Applicable wherever workers interact with physical systems under safety, time, or mobility constraints—including manufacturing, energy, logistics, healthcare facilities, construction sites, and field service operations.

Voice-controlled Physical AI assistants for industrial operations (2/2)

MANAGING RISK AND PROMOTING TRUST

Robust and reliable

Voice recognition must perform reliably in real industrial environments, not just controlled testing conditions. Background noise, acoustic variability, and worker speech patterns all affect accuracy. A system that misinterprets commands during safety-critical or time-sensitive operations creates the workflow interruptions and safety risks it is deployed to prevent.

Fair and impartial

Voice recognition trained predominantly on specific languages or accents may perform less accurately for workers whose speech patterns or dialects differ from training data—creating unequal access to AI-assisted capabilities across a diverse workforce. Organizations should validate recognition accuracy across the full linguistic and demographic diversity of workers who will use the system.

Transparent and explainable

Workers using voice AI for safety-critical guidance and equipment alerts need to understand when they are receiving AI-generated recommendations versus factual data, how confident the system is, and how to override outputs. In hands-occupied environments where cross-checking is difficult, transparency about system limitations is directly relevant to safe operational decision-making.

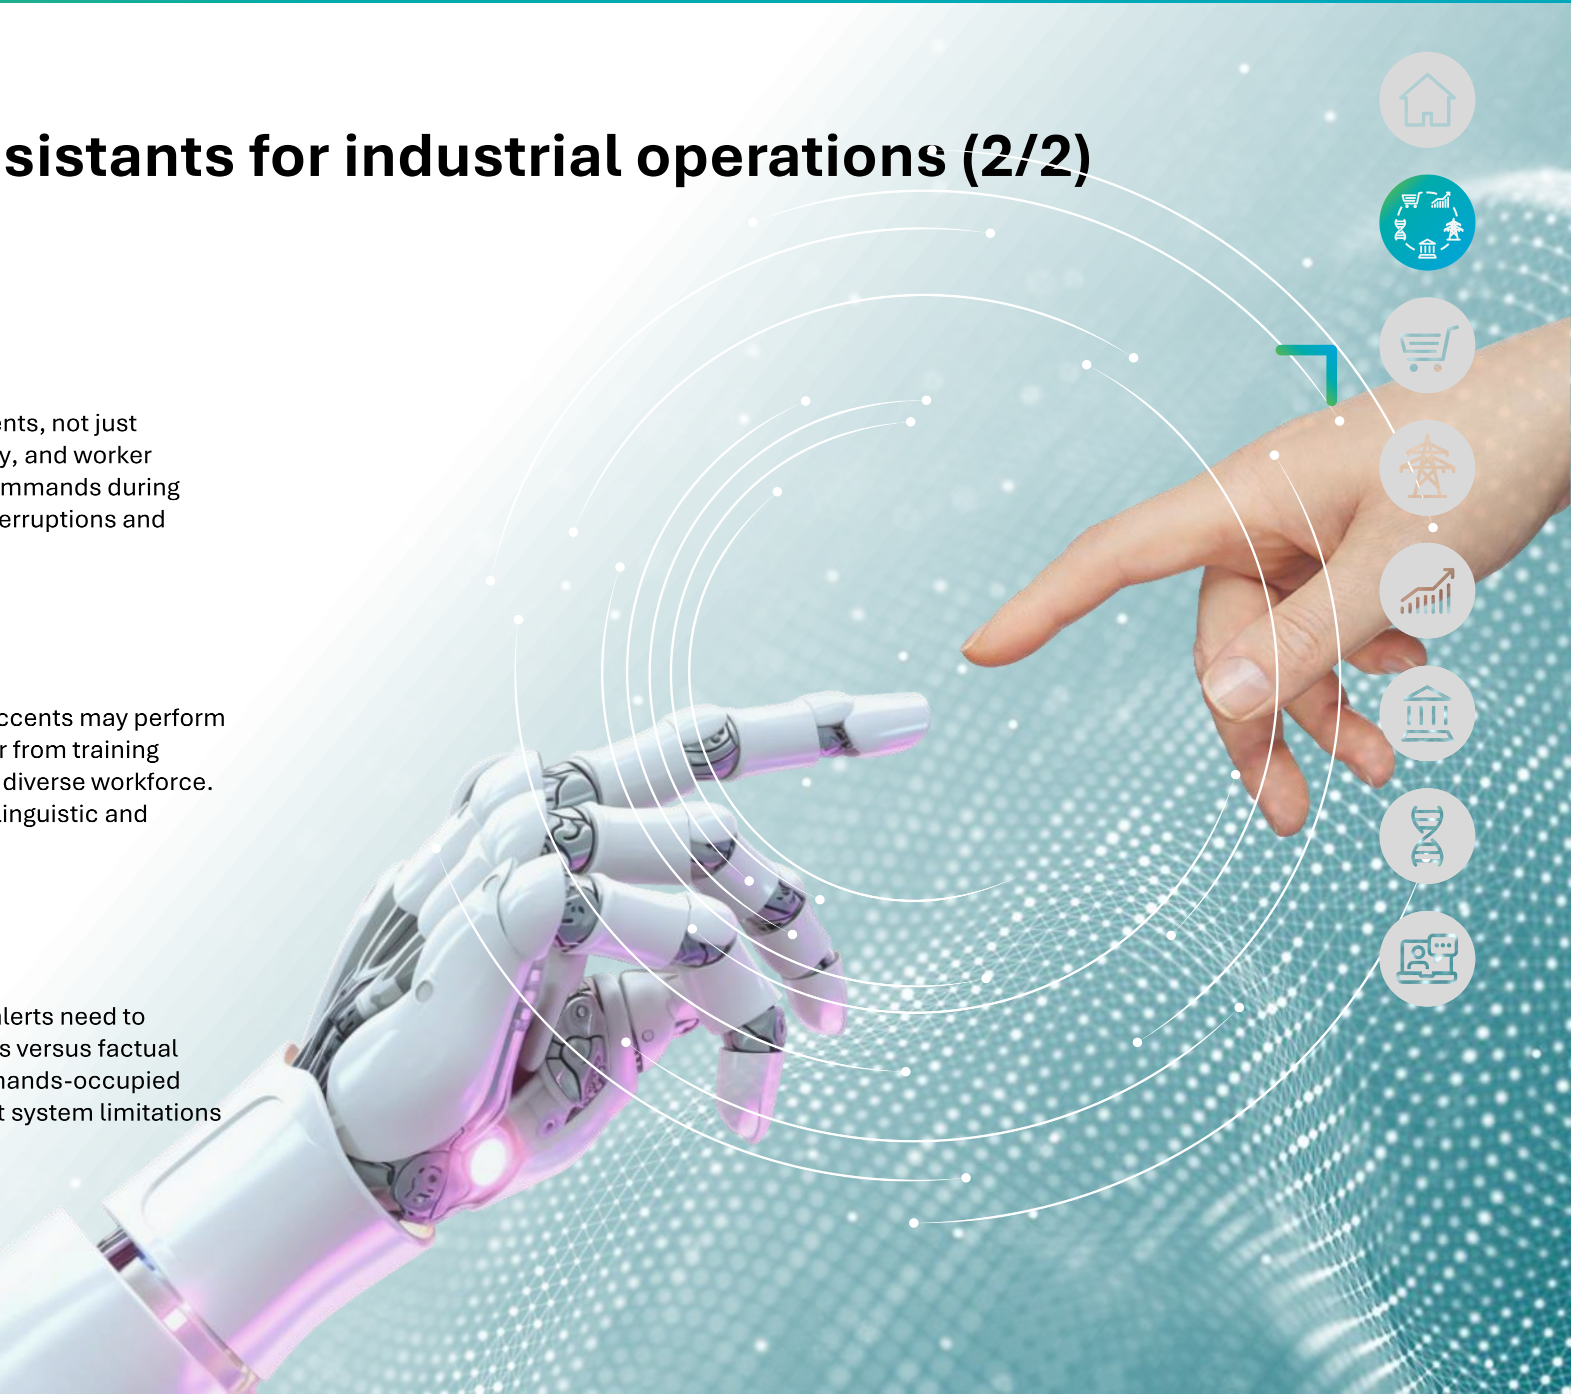

End Notes

1. [2506.06580] AI Simulation by Digital Twins: Systematic Survey, Reference Framework, and Mapping to a Standardized Architecture

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more. Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 460,000 people worldwide make an impact that matters at www.deloitte.com. This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.

© 2026. For information, contact Deloitte Global.