

Deloitte.

The Energy, Resources & Industrials AI Dossier

A selection of high-impact use cases



About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect all the different dimensions of the robust, highly dynamic, and rapidly evolving Artificial Intelligence ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the “Age of With.”

The Deloitte AI Institute aims to promote the dialogue and development of AI, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, the future of work and talent, and applied AI use cases. Combined with Deloitte’s deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you are in: whether you are a board member or a C-Suite leader driving strategy for your organization—or a hands-on data scientist bringing an AI strategy to life—the Deloitte AI Institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet-ups and live events. Let’s explore the future of AI together.

www.deloitte.com/us/AllInstitute



Foreword

Artificial intelligence (AI) continues to advance by leaps and bounds, delivering breathtaking capabilities once thought to be far off in the future. With a remarkable capacity to understand complex inputs and generate valuable outputs—and the rapidly emerging ability to execute real-world actions through intelligent agents and physical AI—AI is opening the door to innovations and new ways of working that were almost unthinkable just a few years ago.

As the AI landscape evolves, so does this compendium. Our latest edition features 130 of the most compelling use cases for AI across six major industries:



Consumer



Energy, Resources & Industrials



Financial Services



Government & Public Services



Life Sciences & Health Care



Technology, Media & Telecommunications

For each of these industries, we explore innovative uses for AI that can address enterprise challenges in new ways, expand and improve capabilities in every business function, and deliver advantages in efficiency, speed, scale, and capacity. To further provide context and clarity, each case specifies the primary business function it supports and whether agentic and/or physical AI is used. These labels are presented for informational purposes, helping you quickly grasp the intention and scope of each case.

Of course, every powerful tool presents potential risks, and AI is no exception. To help you better understand and manage the risks associated with AI, we use Deloitte’s Trustworthy AI™ framework throughout this compendium to illuminate factors that contribute to trust and ethics in AI deployments, and to offer practical steps for strengthening governance and risk mitigation. The specific objective of our Trustworthy AI™ framework is to help organizations create AI systems that are (1)

fair and impartial, (2) robust and reliable, (3) transparent and explainable, (4) safe and secure, (5) responsible and accountable, and (6) private.

Given AI’s rapidly expanding scope and reach, this compendium offers just a glimpse of what the technology can do. Our goal is to convey what AI is currently capable of, and even more important, to inspire the next wave of AI-driven innovation. As AI technology continues to improve and organizations increasingly embrace it, we anticipate even more impressive and compelling use cases in the future—including those that have yet to be imagined.

We hope the use cases highlighted here will spark new ideas, provide a foundation for successful deployments, and set organizations on a path to harness the maximum value from this powerful new technology.



Nitin Mittal
Global Head of Generative AI
Deloitte Consulting LLP



Beena Ammanath
Global Deloitte AI Institute Leader
Deloitte Consulting LLP



Jim Rowan
US Head of AI
Deloitte Consulting LLP

The Energy, Resources & Industrials AI Dossier

AI is emerging as a critical enabler of transformation across the energy, resources, and industrial sectors. In industries defined by asset intensity, operational complexity, and margin pressure, AI offers the ability to sense, predict, and act with speed and precision. From optimizing energy production and forecasting demand to predicting equipment failures and supporting engineers in the field, AI is helping organizations operate more efficiently, safely, and sustainably.

Recent advances—particularly in autonomous control systems, industrial-grade machine learning, and physical AI—are expanding the frontiers of what’s possible. AI agents are managing assets and overseeing field operations with minimal human input. Generative models are accelerating engineering design and fine-tuning operational planning. And physical AI systems are synthesizing sensor data from thousands of assets simultaneously, operating in environments too hazardous for people, and executing physical tasks with greater precision, strength, and consistency than human workers.

These capabilities are arriving at a critical moment. Energy, resources, and industrials companies face mounting pressure to decarbonize, digitize, and build resilience in the face of global volatility. AI can support these goals by unlocking more value from existing assets, enabling predictive and

prescriptive analytics, and surfacing hidden efficiencies in large-scale operations. Integrating real-time data from sensors, satellites, and enterprise systems, AI can offer decision-makers a more holistic, adaptive view of their operations.

However, realizing this potential requires more than model development. It demands robust infrastructure, workforce upskilling, data governance, and alignment with regulatory frameworks—particularly as AI applications begin to intersect with safety-critical systems. Companies that take a disciplined, value-focused approach to AI adoption are already seeing meaningful returns, from reduced downtime and energy waste to faster innovation cycles.

As capabilities mature, AI is set to become not just a tool for operational excellence in the energy, resources, and industrials sectors, but a cornerstone of competitiveness in an increasingly automated and resource-constrained world.

In industries defined by asset intensity, operational complexity, and margin pressure, AI offers the ability to sense, predict, and act with speed and precision.

Note: The tags below each use case indicate its primary business function and whether Agentic or Physical AI is used.

Tags

Primary business function

Agentic AI

Physical AI

AI-driven predictive maintenance

Avoiding downtime through autonomous, multi-agent diagnosis and intervention

Agentic AI systems can monitor industrial equipment health, anticipate failures, diagnose root causes, and proactively schedule maintenance—keeping equipment running smoothly and reducing maintenance costs.

ISSUE/OPPORTUNITY

Industrial equipment failure can trigger substantial costs; yet, traditional maintenance is often reactive or rigidly scheduled, leading to unexpected breakdowns or wasteful over-maintenance. Critical factors such as

labor shortages, weather events, and the move toward electrification are all increasing the need for AI-powered predictive maintenance of industrial equipment and more dynamic, data-driven asset maintenance planning.

Tags

Manufacturing & Quality

Agentic AI

HOW AI CAN HELP

Continuous sensor monitoring and anomaly detection

AI agents can analyze vibration, temperature, pressure, and other IoT sensor data in real time, leveraging predictive analytics to flag deviations from baseline performance and predicting impending failures.

Root cause analysis and diagnosis

When anomalies are detected, specialized agents can assess historical failure logs, maintenance records, and environmental conditions to pinpoint likely failure modes.

Automated work order generation and scheduling

Other agents can generate detailed work orders and schedule tasks based on production cycles, resource availability, and cost constraints.

Simulation and reinforcement learning

Multi-agent reinforcement learning systems can simulate inspection intervals and failure scenarios to reduce maintenance expenses and downtime.

Human-centric integration and continuous improvement

AI agents can collaborate with human maintenance teams: presenting findings in clear, natural language, validating outcomes, helping to prioritize alerts and recommend next steps, and refining models over time based on new data and outcomes.

AI-driven predictive maintenance

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Predictive accuracy should be validated across varied asset types and environmental conditions. Also, AI agents should be tested against historical failure cases and simulated breakdown scenarios, with fall back mechanisms for human review in uncertain situations.



Transparent and explainable

Explainable AI outputs improve adoption and help build trust. Agents need to provide transparent reasoning for their recommendations and actions (e.g., “Vibration on bearing exceeds historical threshold during peak load,” or “Leaf spring failure consistent with past incidents”), supported by traceable data sources.



Safe and secure

Industrial systems are vulnerable to cyber threats, and the consequences of a breach can be severe. Agent platforms should include intrusion detection, secure communication with edge devices, and resilience against malicious sensor tampering or spoofing.



Responsible and accountable

Although AI agents can provide valuable decision support, ultimately human technicians and maintenance managers are responsible for critical decisions and actions. As such, clear escalation protocols need to be established for ambiguous or high-risk alerts.

POTENTIAL BENEFITS

Less unplanned downtime

Real-time, automated detection can trigger repairs early when needed, turning potential disruptions into planned maintenance and reducing productivity losses.

Lower maintenance costs

Focusing on condition-based needs can reduce unnecessary maintenance, minimizing spare-part inventory and technician labor.

Extended asset lifespan and operational efficiency

Continuously monitoring the condition of industrial equipment enables more precise upkeep and longer service life. Data-driven insights improve scheduling and reduce waste, boosting overall productivity and sustainability.

Autonomous drone-based infrastructure inspection

Conducting unmanned, AI-guided inspections of physical assets

Autonomous drones, guided or enhanced by agentic AI, can inspect physical infrastructure such as power lines, pipelines, and transmission towers—capturing and analyzing quality imagery at lower cost and risk.

ISSUE/OPPORTUNITY

Traditional infrastructure inspections in energy, mining, utilities, and industrial environments typically rely on human teams using scaffolding, helicopters, or ropes. In addition to being costly, dangerous, and slow, these manual inspection methods are often disruptive to ongoing operations. What's more, they can miss subtle problems such as hidden defects and weather-related damage.

Autonomous drone systems, especially docked or “drone in a box” variants, enable frequent, hands free inspections that reduce costs, minimize risk to human personnel, and improve overall quality and speed.

Tags

Field Services

Agentic AI

HOW AI CAN HELP

Automated flight and mission control

Drones launch autonomously from preprogrammed docks, fly designated routes, capture high-resolution imagery, and then return for charging and data offload—all without human pilots.

AI-driven defect detection

Onboard analytics can automatically process visual, thermal, or LiDAR data to highlight problems such as cracks, corrosion, vegetation encroachment, and structural anomalies.

Agentic inspection orchestration

An orchestration agent can oversee fleets of drones, schedule inspection missions dynamically, monitor inspection results in near-real time, and then automatically trigger follow-up maintenance workflows when anomalies are detected. This greatly reduces the need for manual coordination.

Extended-range and precision navigation

Advanced AI systems now support beyond-visual-line-of-sight (BVLOS) operations over long distances using perception-aware controllers to maintain accurate positioning and avoid obstacles.

Autonomous drone-based infrastructure inspection

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Inaccurate or inadequate inspections can lead to safety risks and failures. Drone systems need to be validated against varied situations and adverse environments, not just routine test missions and historical defect cases. Also, escalation paths should be established for manual review of ambiguous findings.



Transparent and explainable

AI should provide tangible evidence (such as high-resolution images, heat maps, or annotated photos) accompanied by clear explanations that enable maintenance teams to confidently verify issues.



Private

Some drone inspections may inadvertently capture sensitive data. Agents need to enforce data minimization and restricted-access policies to protect private or proprietary information during data collection and transfer.



Safe and secure

Since drones operate near power lines and potentially dangerous industrial equipment, systems must include breach-resistant communication, secure command chains, and rigorous safeguards against signal spoofing or unauthorized control.



POTENTIAL BENEFITS

Safety improvements

Human workers no longer need to climb towers, pilot helicopters, or enter hazardous zones for routine inspections, reducing the risk of injuries.

Cost and operational efficiency

Autonomous drones reduce inspection time and logistics overhead, significantly lowering costs and enabling more frequent monitoring without disrupting operations.

Proactive maintenance and uptime

Higher-frequency missions and more consistent data allow earlier detection of issues. This enables maintenance to be done before failure occurs, reducing unplanned downtime and enhancing asset reliability.



Autonomous field operations management

Improving safety and efficiency by using AI agents to support field operations

Agentic AI systems can improve field operations by managing task coordination and automating frontline decision-making, enabling field workers to focus on complex, high-value activities rather than routine logistics.

ISSUE/OPPORTUNITY

Field operations such as utility maintenance, site monitoring, asset inspection, and emergency response are typically labor-intensive and highly fragmented. Workers in remote locations often face incomplete data, manual task assignments, shifting priorities, and

limited coordination. These challenges can slow response times, increase safety risks, and make it harder to scale up operations. AI agents can help address such issues by autonomously handling coordination, data collection, and routine decision-making.

Tags

Field Services

Agentic AI

HOW AI CAN HELP

Task management and response automation

AI agents can identify operational issues such as inspection gaps, maintenance alerts, and compliance breaches, while also providing real-time support to field teams, bridging the gap between manual and autonomous operations. They can generate task lists, assign them to the right field crews, and follow up on execution, reducing coordination delays and helping to avoid missed maintenance cycles.

Context-aware adaptation

Agents can respond in real time to changes such as weather shifts, unexpected hazards, and shifting regulatory requirements. They can also reschedule tasks, redirect resources, and escalate safety alerts as needed.

Specialist collaboration

Multi-agent AI teams can mimic expert field units. For example, a maintenance agent can diagnose faults; a scheduling agent can reroute technicians; and a compliance agent can confirm actions comply with regulatory requirements. By working together, AI agents can provide specialized yet cohesive action.

Continuous learning and feedback

Agents can learn from experiences such as safety incidents, unexpected site constraints, and real-world repair times. This feedback loop improves their task allocation and decision logic, boosting long-term effectiveness and responsiveness.

Autonomous field operations management

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Field operations take place in unpredictable environments. Agents should be tested against variable and adverse conditions, such as network outages and natural disasters. Fallback paths should include human intervention for ambiguous cases.



Transparent and explainable

Technicians, supervisors, and regulators need to understand AI decisions. Agents should provide clear rationales such as “rescheduled inspection due to stage-1 weather alert,” backed by task logs and field data.



Safe and secure

Safety in field operations is of paramount importance. As such, agentic systems must include cybersecurity safeguards such as encrypted communications, device authentication, anomaly detection, and protection against unauthorized commands.



Responsible and accountable

Governance structures should define escalation paths, sign-off mechanisms, and oversight protocols for AI-generated plans. Although agents can provide useful support, ultimate responsibility for critical decisions and actions should remain with human field supervisors and operations managers.



POTENTIAL BENEFITS

Faster field response

AI agents can close the gap between detection and action by assigning tasks and dispatching crews within minutes of an alert, eliminating unnecessary delays.

Higher operational uptime

Proactive task assignment and adaptive scheduling can help ensure that critical infrastructure remains up and running, minimizing downtime and safety disruptions.

Scalable field capacity with fewer resources

AI agents can help smaller teams effectively manage larger service areas, expanding operational reach without requiring proportional staffing increases.



Intelligent commercial operations

Driving smarter bidding, pricing, and customer engagement with AI agents

AI can streamline and accelerate commercial operations by using agents to automate contract pricing, bid preparation, demand forecasting, and customer engagement workflows.

ISSUE/OPPORTUNITY

Commercial teams in the energy, resources and industrial sector face complex pricing challenges across multiple product lines, geographic regions, and regulatory regimes. Traditional methods are labor-intensive and revolve around spreadsheet models and fragmented data sources. This highly manual approach slows down responses to competitive pressure and limits the ability to grow.

Businesses in the sector need systems that can accelerate commercial workflows, integrate real-time signals, and personalize offers to customers at scale. Multi-agent AI offers a potential solution, providing a path to more flexible bidding and pricing through automated analysis and continuous adaptation.

HOW AI CAN HELP

Pricing and bid preparation agents

A pricing agent analyzes input costs, market trends, competitor rates, and regulatory constraints. A bid preparation agent then drafts customized proposals based on customer histories, relevant contract terms, and compliance guidelines.

Demand forecasting and scenario modeling

A demand agent pulls in real-time signals—such as weather forecasts, energy prices, and consumption patterns—then simulates demand under different conditions to inform pricing and bid decisions.

Customer engagement coordination

An orchestration agent coordinates the bid, pricing, and demand agents to ensure cohesive action. It sequences tasks, manages version control, and forwards AI output to human commercial teams for review and execution.

Contract rollout and monitoring

After the sale, specialized agents monitor contract performance and market deviations, flagging margin erosion that needs to be addressed and triggering renewals when favorable conditions arise.

Intelligent commercial operations

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Pricing errors can reduce margins and trigger regulatory issues. Agents must be tested on historical scenarios and stress-tested for market shocks.



Fair and impartial

Agents must avoid bias toward certain customers or regions. Pricing models should be audited to ensure fairness, and human commercial teams should have the ability to override automated suggestions.



Safe and secure

Given the sensitivity of pricing data and customer value, agentic systems must include secure authentication, intrusion detection, and validation of external inputs to guard against manipulation or cyber interference.



Responsible and accountable

Escalation protocols should be in place for high-value or sensitive proposals, with humans retaining final responsibility for commercial offers and contract decisions.



POTENTIAL BENEFITS

Faster bidding cycles at lower cost

AI automation can streamline and accelerate bid preparation, reducing costs and improving responsiveness and competitiveness.

Higher margin capture

Dynamic pricing based on real-time data and scenario modeling helps maximize margins.

Scalability

AI agents help commercial teams efficiently manage larger customer portfolios, enabling faster business growth without proportional staffing increases.

Better strategic alignment

By freeing commercial teams from routine administrative tasks, AI agents allow people to focus more attention on strategic market expansion, negotiations, and relationship building.



Expediting experiments and design

Materials design

AI empowers materials designers to explore a wider design space, optimize material properties, and expedite the discovery of new materials.

ISSUE/OPPORTUNITY

Developing new materials is challenging, costly, and time-consuming. One reason is that the chemical space is vast and complex while the number of chemically feasible molecules is unknown. Also, the

materials discovery, development, and optimization process present different complexities at each stage, increasing the time required to reach a final design.

Tags

R&D/Product Development

HOW AI CAN HELP

Streamline experimental process

Using AI to determine the most efficient experimental procedures for probing or optimizing materials can streamline the experimental stages of development by removing redundant experiments and undertaking those that are cost- and time-optimized.

High-entropy alloy (HEA) engineering

Traditional techniques for developing HEAs with excellent physical, chemical, and mechanical properties are time-consuming and costly, making AI modelling a promising alternative development pathway.

Expediting experiments and design

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Intellectual property or a similar competitive advantage could be compromised by using AI in materials design, as models trained on proprietary or sensitive data could potentially reveal valuable insights or design strategies to competitors.



Responsible and accountable

Companies should be mindful to identify and mitigate unintended negative ramifications of materials designed with the support of AI, such as long-term environmental impacts from materials that cannot be manufactured in responsible and sustainable ways.

POTENTIAL BENEFITS

Fueling innovation

AI applications have the capability to rapidly generate and prioritize a wide range of virtual materials with diverse compositions and structures. This virtual screening process allows researchers to identify potential candidates for specific applications or material properties much more quickly than traditional experimental methods.

Bringing down costs

Through efficiency savings and the rationalization and/or elimination of experiment consumables, the organization can reduce development costs.

Enabling discovery

AI maximizes the likelihood of discovering materials with superior properties by leveraging its ability to efficiently explore and navigate a vast design space of potential materials.

Understanding the ore

Minerals processing optimization

AI can make the process of chemical separation of minerals from ore more cost- and time-efficient, safer, and more environmentally sustainable.

ISSUE/OPPORTUNITY

In mineral processing, chemical additives must be matched to the exact contents of the ore to separate as much material as possible from waste minerals without destroying it. The process is complicated due to the fact that modelling and testing each compound is time-and

effort-intensive, complex mineralogy and interrelationships between minerals can hinder recovery, and environmentally hazardous chemicals are often necessary to process certain compounds.

Tags

Manufacturing & Quality

HOW AI CAN HELP

Ore characterization and mapping

AI models can be trained on large datasets of mineral samples to generate synthetic samples that mimic the characteristics of real-world ores. Comprehensive databases can be built for mineral identification, classification, and prediction of ore properties, permitting insights into the behavior and composition of different ores without testing on known processing assays.

Process optimization

Models that simulate the physical and chemical processes involved in mineral processing can help optimize factors like grinding parameters, flotation conditions, and separation techniques. This can improve efficiency, reduce energy consumption, and enhance mineral recovery rates.

Understanding the ore

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

AI models may struggle to generalize mineral samples and processing scenarios that are significantly different from the training data. The model might not capture the full range of variations and unique characteristics of novel ores, which could lead to suboptimal processing recommendations. Also, if AI models cannot interpret complicated physical and chemical qualities like particle size distribution, mineral composition, and processing conditions, the model may generate suboptimal strategies or overlook critical factors.

POTENTIAL BENEFITS

Accelerated exploration

The cost and time needed to characterize ore and develop a processing workflow can be significantly reduced, and cost and efficiency trade-offs can be optimized to maximize mineral recovery while minimizing operational costs.

Eco-friendly operations

Keener insights into mineralogy using AI can help reduce the amount of environmentally damaging additives and resources needed for processing without sacrificing production volume or efficiency.

Occupational health

Optimized processing can help reduce human exposure to toxic chemical additives and fine particle dust, which contributes to a safer work environment.

Optimize the design

Site design generation

AI can support the development of site plans by automating aspects of the design process, providing designers with new possibilities and reducing the associated time and cost.

ISSUE/OPPORTUNITY

Site planning is a multi-stage, iterative process to optimize cost, efficiency, and safety, but it is also an expensive and time-consuming exercise involving numerous stakeholders and third-party specialists. Site planning can require surveys in remote, sometimes hostile

locations. Forecasting near- and long-term impacts involves assessing a multitude of factors, and site-specific activities such as topological and geological surveying can be labor intensive and expensive.

Tags

R&D/Product Development

HOW AI CAN HELP

Automated layout generation

Designers can use AI to analyze site constraints, design requirements, and input from engineers to quickly generate layout options for site plans that consider factors such as zoning regulations, operational use, and user preferences.

Design optimization

AI can help optimize site plans by analyzing parameters like solar orientation, traffic flow, and accessibility to suggest optimal infrastructure placements. This can help improve energy efficiency, support better space utilization, and enhance the user experience.

Efficient documentation and annotation

By analyzing design elements and structures in the generated plans, AI can automatically annotate the plans with relevant information, such as dimensions, materials, and specifications. This automation could save designers considerable time and effort, allowing them to focus on higher level design tasks.

Optimize the design

MANAGING RISK AND PROMOTING TRUST



Responsible and accountable

AI for design optimization may focus primarily on efficiencies, such as cost reduction or time savings, while potentially neglecting other important considerations, such as environmental sustainability, community impact, or long-term adaptability. The model should be configured to balance multiple objectives and prioritize trade-offs to achieve better overall outcomes. Using AI for site planning also raises legal considerations around intellectual property, ownership of AI-generated designs, liability for design flaws, and privacy restrictions for sensitive or proprietary data.

POTENTIAL BENEFITS

Acceleration with automation

Using AI for site planning can accelerate the completion of time-consuming processes.

Discovering new solutions

With AI quickly creating a variety of site designs, the planning process can include a greater diversity of designs and promote innovative planning solutions.

Reducing risk

AI can simulate and analyze potential hazards and safety risks in site plans. AI-generated planning would consider factors such as weather events, traffic patterns, and emergency response routes. It could propose alternative design options to proactively minimize risks to safety and reduce potential property damage in case of unforeseen events.

Enhancing employee safety

Personalized OHS training

AI can be used to develop personalized and immersive occupational health and safety (OHS) training materials that allow trainees to be safely exposed to realistic scenarios and thereby reduce or better respond to real OHS incidents.

ISSUE/OPPORTUNITY

Traditional OHS training may only cover some potential scenarios, and it lacks practical opportunities to apply new skills and knowledge. Workers need to be prepared for emergency scenarios but cannot practice managing these scenarios in a real-world setting due to the cost and risk involved.

Tags

Learning & Development

HOW AI CAN HELP

Virtual reality (VR) training

Combined with VR, AI can be used to develop virtual training environments that replicate operational conditions. With realistic scenarios that simulate OHS incidents, trainees can navigate hazardous situations, identify risks, and improve their OHS awareness and response capabilities in a safe setting.

Customized training content

AI can be used to customize training materials based on specific job roles, site conditions, or regulatory requirements. This technology can analyze large volumes of data, such as incident reports, OHS guidelines, or compliance standards and generate tailored content, including videos, interactive modules, or quizzes.

Enhancing employee safety

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Real-life emergencies can be highly stressful and traumatic, and replicating these scenarios virtually could imperil the psychological safety of trainees. The final design of simulations should be reviewed by human trainers to remove potentially harmful visualizations.



Responsible and accountable

AI-generated training materials should be continuously monitored to identify any potential issues, inaccuracies, or outdated information. Regular updates to the training content should be made to reflect the latest safety guidelines, regulations, and best practices.



Fair and impartial

AI-generated training materials should be designed to be accessible to all types of learners, including individuals with disabilities. Organizations should consider providing closed captions for videos, adjustable training scenarios to accommodate different skill levels, and alternative formats for content.



POTENTIAL BENEFITS

Safety through preparedness

Increased training engagement and readiness for emergencies supports workforce safety and fewer OHS incidents.

Customized training

A personalized approach to OHS training helps address the specific needs of workers, ensuring they receive relevant and targeted instruction.

Dynamic compliance

Changes in legislation, regulation, and policies can be quickly reflected in training materials by using AI to make updates.



Peering below the surface

Hydrocarbon reservoir exploration

AI can be used to optimize exploration success rates, reduce costs, and mitigate risks associated with hydrocarbon reservoir location and characterization.

ISSUE/OPPORTUNITY

Oil and gas exploration involves a high degree of uncertainty and risk. Advanced technologies and extensive data analysis are needed to navigate the subsurface and accurately locate and characterize reservoirs. Extracting oil and gas from underground reservoirs requires advanced drilling techniques and technologies,

and harsh environmental conditions, deep water, and complex logistics make offshore exploration difficult. This makes exploration a capital-intensive and time-consuming process involving multiple stages of seismic surveys, analysis, drilling, and testing.

Tags

R&D/Product Development

HOW AI CAN HELP

Seismic data analysis

To overcome incomplete, low volume, or poor-quality seismic data, AI can support enhanced data analysis and interpretation. AI can be used to generate new data samples that resemble the patterns and characteristics of the existing seismic data, address missing or incomplete seismic data, improve data quality through denoising or resolution enhancement, and more effectively interpret complex data patterns.

Reservoir characterization

By analyzing data sources such as well logs, core samples, and production data, AI can create models that simulate the more complete behaviors of hydrocarbon reservoirs. This enables a better understanding of the reservoir dynamics, which helps optimize production strategies and improve recovery rates.

Peering below the surface

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

False positives or misinterpretations may result in costly and time-consuming drilling operations that do not yield productive reservoirs, making human expertise crucial to validating insights and decision-making. Also, AI models could overlook critical factors or geological nuances that human geoscientists would recognize and could fail to contextualize the data when generating outputs. Without a contextual understanding, AI models and interpretations may lack accuracy or fail to capture the full complexity of reservoirs.

POTENTIAL BENEFITS

Informed investments and decisions

A deeper, more complete understanding of the characteristics of hydrocarbon reservoirs reduces uncertainty and enables better investment decisions.

Amplifying exploration

Improved data quality supports more accurate subsurface modeling, imaging, and structure characterization, which leads to an increased ability to accurately locate hydrocarbon reservoirs.

Smarter strategy

With an earlier and more complete understanding of reservoir characteristics, less time is needed to optimize production strategies.

A smart eye in the sky

Smart summaries for drone surveying

AI can assist in summarizing large volumes of drone footage and enable querying to enhance productivity and efficiency.

ISSUE/OPPORTUNITY

In the mining industry, drones are increasingly used for tasks such as mapping, management of dam tailings, safety management, blast assessment, environmental monitoring, and haul road optimization. In the case of Optical Gas Imaging (OGI) to detect gases and volatile organic compounds leaking from vessels (e.g., pipelines), unmanned drones mounted with OGI cameras have proven useful for surveying a variety of equipment

over vast areas. Using drones in this way permits frequent scans and reduced costs associated with fugitive gases. Yet, while advanced AI solutions (e.g., volumetric monitoring) have been developed for applications using drone footage, manual inspection of drone footage is still required for environmental monitoring, security review, safety assessment, and retrospective analysis.

HOW AI CAN HELP

Smart summaries

Combined with computer vision solutions, AI can create smart assistive summaries in natural language from thousands of hours of drone footage. Assistive smart summaries can be based on a pre-determined template requested by the user, where observations are generated about elevations, topology, lighting, vegetation, and other factors. Summaries can also be queried in natural language so questions can be asked without the assessor manually reviewing all footage.

Querying the footage

When using OGI to detect leaks, there may be instances where a leak is irreparable but still must be managed. With AI, specific sites can be efficiently reviewed and monitored using simple natural language queries.

A smart eye in the sky

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

AI models may struggle to interpret environmental indicators, assess ecological impacts, or consider local conditions and regulations. Training data availability and quality in particular can impact the AI model's ability to generalize and handle diverse environmental scenarios. Inadequate or biased training data may result in limited or skewed analysis and summaries.



Private

Drone footage may contain sensitive information, including personally identifiable information, facial images, or confidential business information, and the footage may also be captured on private properties or areas with restricted access. In using AI to analyze and summarize the footage, unsecure data handling and access can raise privacy concerns as well as legal and regulatory implications.

POTENTIAL BENEFITS

Supplementing human expertise

Querying smart assistive summaries helps ensure critical observations are not missed due to human error or cost and time constraints.

Faster time to insight

Replacing manual drone footage inspection with assistive summaries saves significant time and effort.

Resilient logistics and planning

Supply chain optimization

AI can support supply chain optimization by leveraging its ability to simulate, model, and generate data-driven insights.

ISSUE/OPPORTUNITY

Global supply chains are highly interconnected with many dependencies and multiple stakeholders. The inherent complexity creates challenges to efficiency, resilience, and cost avoidance, making supply chain intelligence a

critical component of supply chain management. What is needed is a way to rapidly analyze data from internal and external sources to identify patterns and areas for improvement.

Tags

Procurement/Sourcing & Supply Chain

HOW AI CAN HELP

Supply chain intelligence

AI can help identify and simulate potential disruptions or risks in the supply chain. By assessing port congestion, shipment routes, and tier-n supplier mapping, AI can be used to predict risks and their corresponding impact on operations, then recommend actions to mitigate those risks. This allows supply chain managers to proactively implement mitigation strategies, develop contingency plans, and improve overall resilience.

Scenario analysis and optimization

Supply chain managers can use AI to run what-if scenarios in a digital twin environment that reflects the real-world supply chain. By simulating the impact of changes in demand patterns, production capacity, inventory strategies or supplier reliability, supply chain managers can improve risk assessments and proactive decision-making based on real-time conditions.

Supply chain planning

AI enables supply chain professionals to use natural language to interact with advanced planning solutions. Questions concerning all supply chain areas, such as planning, inventory, supply assurance, order management, and global logistics, can be easily asked, helping even less experienced users navigate complex topics and data.

Supplier assessment

AI can assist in supplier evaluation and relationship management by analyzing financial reports, performance metrics, customer feedback, and other data and then generate insights and predictions around supplier performance, risk factors, and opportunities for collaboration. This helps supply chain professionals make informed decisions when selecting, negotiating with, and managing suppliers.

Resilient logistics and planning

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Supply chain management involves complex trade-offs, strategic considerations, and tacit knowledge that AI models may not fully capture. AI outputs may also fail to balance ethical considerations or long-term strategic goals. As such, human judgment and validation is central to the interpretation and augmentation of AI outputs.



Fair and impartial

When using AI for supplier evaluation, negotiating, and contracting, bias in the data or model could lead to unfair recommendations and discriminatory practices. By taking into account factors such as fair contract terms, social responsibility, and ethical sourcing practices, organizations can promote decision-making processes that are fair and transparent.

POTENTIAL BENEFITS

Resilient supply chains

Enhancing supply chain resilience allows the organization to respond quickly to changing market dynamics and permits greater agility to take advantage of emerging opportunities based on real-time insights and recommendations.

Enhanced performance

By prioritizing alerts that require human intervention and differentiating between noise and disruption, the organization can drive greater efficiency in the supply chain.

Optimizing efficiency

Making optimized decisions across the supply chain, from supplier selection to fulfilment optimization, helps reduce costs, minimize waste, and improve overall operational efficiency.

Enabling a better grid

Grid and energy efficiency optimization

AI can be used to better understand the state of the grid and factors that could support more efficient energy consumption, minimizing losses and improving overall grid efficiency.

ISSUE/OPPORTUNITY

Energy grids are massive and intricate systems with interconnected components operating in a dynamic and uncertain environment. Maintaining a balance between energy supply and demand is crucial for grid stability, but it is challenged by the difficulty in predicting and managing fluctuations in energy demand. The integration of intermittent

renewable energy sources (e.g., solar) further complicates the supply-demand balancing act as these depend on weather conditions. Regulatory frameworks, policies, and market structures also constrain the ability to balance technical optimization.

Tags

Operations

HOW AI CAN HELP

Promote informed customer behavior

Energy companies can incentivize consumers to adjust their energy consumption based on their specific energy use patterns using conversational chatbots powered by generative AI. AI models can analyze historical data and customer preferences to recommend personalized strategies to reduce energy usage. When there is an immediate need to reduce peak loads to improve grid stability, AI applications can be used to alert customers about what they can specifically do to help. What is more, conversational chatbots can be used as an educational tool for consumers to understand and optimize their energy usage.

Document and map digitization

AI can be used to digitize documentation, infrastructure maps, and records of energy use, as well as for image-to-image translation or image restoration (e.g., by removing noise, adjusting brightness,

and enhancing contrast). This improves the quality of the documents and yields searchable documents that can be used to train existing AI classification and forecasting tools.

Grid layout and expansion

AI can assist in designing optimal configuration and expansion plans for the energy grid. AI models can generate optimized grid designs that minimize transmission losses and maximize efficiency by considering factors such as population density, existing infrastructure, and energy demand projections.

Energy trading and market analysis

AI models can simulate the behavior of electricity markets under different scenarios, such as regulation changes or the introduction of new technologies. This can help energy companies optimize their trading strategies and make more informed investment decisions.

Enabling a better grid

MANAGING RISK AND PROMOTING TRUST



Private

Using AI in customer behavior analysis and chatbot interaction involves handling sensitive customer data. Risks include data breaches and unauthorized access to customer information and chat logs, and risk mitigation requires robust security measures, customer data protection, and adherence to privacy regulations.



Safe and secure

AI models are vulnerable to adversarial attacks, where malicious actors manipulate inputs to deceive or exploit the system, for example, to influence energy trading decisions or disrupt grid operations. Robust security measures and regular testing are necessary to mitigate such risks.

POTENTIAL BENEFITS

Diversifying energy sources

AI supports the integration of variable renewable energy sources while maintaining stability and reliability.

Dynamic demand response

Using AI for improved visibility of the grid's current state allows companies to better respond to fluctuations in demand.

Ongoing optimization

As more trends, data and documents are digitized and analyzed over time, AI enables continuous improvement in efficiency optimization and managing demand.

Simulation-first development & digital twins

Validating physical systems virtually before real-world deployment

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 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 the majority of 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 any hardware is deployed.

Tags

R&D/Product Development

Physical AI

HOW 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.

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.

Regression testing for updates

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

Design validation before build

Digital twins 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.

Simulation-first development & digital twins

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.



POTENTIAL BENEFITS

Faster development cycles

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

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.

Reduced deployment risk

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



Simulation-trained, human-supervised closed-loop remediation

Safe, simulation-trained intervention for physical systems

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 critical environments.

ISSUE/OPPORTUNITY

Physical AI systems—robots, drones, autonomous vehicles, and intelligent infrastructure—must operate reliably in complex, safety critical 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.

Tags

Operations

Physical AI

HOW 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.

Closed loop learning with digital twins

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

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.

Action impact prediction before execution

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

Approval-based execution

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

Simulation-trained, human-supervised closed-loop remediation

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.

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.

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.

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.

Transferable remediation patterns across industries

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

Simulation-driven remote operations and training

Scaling expertise through digital twins and immersive simulation

AR/VR digital twins create high-fidelity virtual replicas of offshore platforms and production facilities, enabling experts 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 expertise, faster decision making, and reduced physical exposure without compromising control or accountability.

Tags

Learning & Development

Physical AI

HOW 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 experts to train, rehearse, and validate actions in a risk free virtual environment before interacting with live systems.

Remote decision support with contextual awareness

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

Expert 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.

Governed human in the loop operations

AI supports recommendations and scenario evaluation, while humans retain approval authority for all 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.

Simulation-driven remote operations and training

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.

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.

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 allows 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.

Major cost efficiency gains

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

Scalable expertise 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 critical, remote, or complex environments—including ER&I, logistics, and healthcare facilities, where simulation first validation improves safety, reliability, and operational confidence.

Autonomous self-calibrating quality and process control

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

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.

Tags

Manufacturing & Quality

Physical AI

HOW 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.

Self monitoring and drift detection

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

Integrated inspection across stages

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

Predictive quality modeling

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

Automated calibration and correction

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

Autonomous self-calibrating quality and process control

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.



POTENTIAL BENEFITS

Reduced scrap and rework

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

Minimized downtime and production disruption

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

Continuous quality improvement

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

Sustained accuracy and reliability

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

Stronger auditability and validation

Comprehensive quality and process data supports regulatory audits and customer certifications.

Lower maintenance overhead

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



Voice-controlled physical AI assistants for industrial operations

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

Voice controlled physical AI systems designed specifically for industrial environments 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.

Tags

Operations

Physical AI

HOW 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 enables 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 critical tasks.

Voice-controlled physical AI assistants for industrial operations

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.



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.



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.

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.

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.

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.

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.

Predictive monitoring for environment health & safety

Vision-based safety enforcement at industrial scale

Physical AI systems embedded in vision enabled drones, robots, and fixed infrastructure continuously perceive safety conditions across ER&I environments, reason about evolving risk in context, and perform predictive monitoring (e.g., early detection of equipment anomalies or evolving hazards), triggering governed physical interventions or escalations to prevent incidents before harm occurs.

ISSUE/OPPORTUNITY

Environment health & safety (EHS) monitoring is often inconsistent and reactive when it relies on manual observation and reporting. Unsafe behaviors may go unnoticed until incidents occur, and enforcement can vary across shifts and sites. Supervisors cannot continuously observe all work areas, and workers may cut corners when oversight is absent, creating gaps in safety compliance that accumulate until an accident reveals the problem.

Manual incident reporting depends on workers recognizing and documenting near-misses or violations, but reporting rates are

low when workers fear repercussions or when the urgency of production deadlines overshadows safety protocols. Enforcement varies across shifts as different supervisors apply safety rules with different levels of rigor, and compliance tends to be highest when leadership is present and lowest during off-shifts when oversight is minimal.

Continuous, automated monitoring improves compliance while minimizing operational disruption, enabling consistent enforcement regardless of shift, location, or supervisor availability.

Tags

Compliance & Risk

Physical AI

HOW AI CAN HELP

Continuous video analysis

AI reviews camera/drone feeds in near real time to detect safety-relevant behaviors and conditions across all monitored areas simultaneously, providing coverage that manual observation cannot sustain.

Scalable multi-site deployment

A consistent detection framework can be rolled out across facilities, ensuring uniform safety standards and comparable metrics regardless of location, facility size, or local supervision practices.

Contextual filtering

AI can reduce noise by distinguishing true violations from benign activity in complex environments where authorized workers may temporarily enter restricted zones for legitimate maintenance or where PPE requirements vary by task.

Event-driven alerting

Detected risks trigger alerts and follow-up workflows for supervisors or EHS teams, enabling immediate intervention when unsafe conditions are identified rather than waiting for scheduled safety audits.

Non-intrusive integration

Systems leverage existing camera infrastructure and operate alongside current procedures, avoiding the cost and disruption of installing new sensor networks or redesigning facility layouts.

Rule-based safety interpretation

Models apply predefined rules (zones, PPE requirements) combined with sensor-driven context from robotic/drones to consistently detect violations without variation across shifts or individual judgment differences.

Predictive monitoring for environment health & safety

MANAGING RISK AND PROMOTING TRUST



Private

Continuous video monitoring of workers across an entire facility—capturing behaviors, locations, and movements throughout every shift—is one of the most expansive forms of workplace surveillance an employer can deploy. Organizations should apply data minimization principles, retain footage only as long as operationally necessary, anonymize where feasible, and be explicit with workers about what is monitored, how long data is retained, and who can access it and under what circumstances.



Responsible and accountable

AI-assisted safety monitoring creates accountability exposure in both directions: the system can be implicated for acting on a false positive that disrupts operations or triggers an unwarranted disciplinary action, and for failing to detect a genuine hazard that preceded a workplace injury. Governance must clearly establish that AI outputs are merely inputs to human judgment and that accountability for safety outcomes remains with qualified EHS personnel.



Fair and impartial

Biased detection in a safety enforcement context carries direct legal consequences, not just ethical ones. Models trained on historical violation data that over-represents certain demographic groups might disproportionately flag those workers for enforcement action, creating discrimination exposure under employment law.



POTENTIAL BENEFITS

Improved compliance

More consistent enforcement of safety rules and restricted zones as automated monitoring maintains the same vigilance regardless of shift, production pressure, or supervisor presence.

Standardized practices

Comparable monitoring and metrics across sites for governance enables enterprise-wide safety benchmarking, identification of high-risk locations, and evidence-based allocation of safety resources.

Reduced incident risk

Earlier detection of unsafe behavior through real-time physical environment (including drone imagery) monitoring enables intervention before violations escalate into injuries, equipment damage, or environmental releases.

Lower monitoring burden

Less manual observation and reporting overhead allows EHS personnel and supervisors to focus on investigating root causes, implementing corrective actions, and improving safety training rather than routine surveillance.



Autonomous haulage systems for safe & intelligent mining operations

Autonomous trucks optimizing safety, uptime, and haulage efficiency

Autonomous mining trucks and AMRs use LiDAR, radar, and AI-managed vision to handle hauling, inspections, and maintenance with safe sensor-based hand-offs. This same technology surveys electric grids and transformers, autonomously detecting infrastructure defects and safety hazards.

ISSUE/OPPORTUNITY

Mining operations involve large, expensive vehicles operating in hazardous conditions. Human error and limited visibility increase accident risk and equipment damage. Manual operations also constrain productivity.

Haul truck operators work long shifts in dusty, noisy environments with limited sightlines, navigating massive vehicles weighing hundreds of tons along narrow roads carved into pit walls. Fatigue, distraction, or misjudgment can result in catastrophic accidents including collisions with other equipment, running off roadways, or striking workers on foot. Driver recruitment and retention are challenging given

the remote locations, harsh working conditions, and monotonous nature of repetitive hauling cycles. Equipment damage from operator error—such as overloading, improper dumping, or collisions—creates costly downtime and repair expenses that significantly impact mining economics.

The opportunity is to deploy physical AI systems that enable autonomous operation while maintaining extremely high safety standards through redundancy and continuous monitoring, removing workers from hazardous roles while improving operational consistency.

Tags

Operations

Physical AI

HOW AI CAN HELP

Sensor fusion for perception

AI combines truck LiDAR, radar, and cameras with complementary drone RGB/thermal imaging and AMR/quad 3D mapping to build an accurate, multi altitude site model.

Safety-zone enforcement

AI maintains dynamic safety buffers around vehicles that adjust based on speed, load weight, road conditions, and proximity to other equipment or personnel.

Redundant safety validation

Multiple sensors cross-check detections to avoid false alarms, ensuring that critical safety determinations are confirmed by independent sensor systems before actions are taken.

Continuous telemetry monitoring

Vehicle health and operation are tracked centrally, enabling fleet managers to monitor performance, predict maintenance needs, and intervene remotely if needed.

Controlled-environment deployment

Autonomy is restricted to defined mining zones with known characteristics, avoiding the complexity of public roads and unpredictable environments outside the controlled mining area.

Fail-safe autonomy design

Any uncertainty triggers immediate stops, prioritizing safety over productivity when sensor disagreement, unexpected obstacles, or system anomalies are detected.

Autonomous haulage systems for safe & intelligent mining operations

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

The safety case for autonomous haulage relies on the system outperforming human drivers; however, autonomous trucks can introduce their own failure modes in the conditions that define mining environments: wet roads that degrade traction sensing, dust that compromises vision, GPS-limited areas that undermine positioning, and sudden obstacle detection that causes emergency stops and lane breaches. Validation must account for these conditions explicitly, not treat them as edge cases.



Safe and secure

GPS spoofing, signal jamming, and unauthorized command injection are potential attack vectors against autonomous haulage systems—and on an active mine site where massive vehicles operate at speed, a successful attack could redirect vehicles, suppress safety alerts, or lock entire fleets with potentially fatal consequences. Cybersecurity must be governed with the same rigor as physical safety, with layered defenses and regular adversarial testing as standard operational requirements.



Responsible and accountable

Regulators may impose enforceable incident reporting and investigation obligations that create formal governance requirements for autonomous systems. Operational logs capturing vehicle perception, decision-making, and control states must be maintained with sufficient detail and accuracy to satisfy regulators and support liability determination.

POTENTIAL BENEFITS

Less accident risk

Accident risk is significantly reduced in hazardous environments as autonomous systems eliminate operator fatigue, distraction, and judgment errors while maintaining consistent adherence to safety protocols.

Reduced equipment losses

Expensive equipment losses are avoided through precise vehicle control that prevents damage from operator error such as collisions, overloading, or driving off roadways, extending the life of multi-million-dollar haul trucks.

Greater efficiency

Operational efficiency improves through continuous, predictable vehicle operation that maintains consistent cycle times without breaks or shift changes, maximizing equipment utilization around the clock.

Higher overall productivity

Mining productivity increases without compromising safety as autonomous fleets can be scaled more easily than human workforces.

Improved worker safety

Workforce exposure to danger declines as human operators are removed from hazardous roles involving heavy equipment navigation in confined spaces and extreme conditions, allowing workers to be redeployed to safer supervisory and maintenance positions.

Autonomous inspection and intervention in safety-critical physical environments

Drone-first inspection with robotic intervention

Physical AI robots operate in hazardous environments such as oil wells, energy facilities, and industrial plants. Small hybrid fleets such as drones for aerial inspection and a robotic hand equipped ground unit for simple actuation are used to keep people out of risk zones.

ISSUE/OPPORTUNITY

Human inspection in hazardous environments exposes workers to significant risk and limits inspection frequency. Inspectors must enter confined spaces with toxic gases, high temperatures, or radiation exposure to check equipment conditions and identify developing problems. The need for extensive safety protocols, personal protective equipment, and standby rescue teams makes each inspection expensive and time-consuming.

Inspection frequency is constrained by the hazards involved—facilities may only inspect critical equipment quarterly or annually

when continuous monitoring would be preferable. However, if failures are detected late—or occur between inspections—the consequences can be catastrophic, including explosions, fires, toxic releases, or environmental contamination.

The opportunity is to deploy physical AI robots that continuously inspect and intervene where safe, enabling constant vigilance in areas too dangerous for human presence.

HOW AI CAN HELP

Selective physical intervention

Robots perform simple corrective actions such as closing valves, tightening fittings, or applying temporary sealants to contain minor issues before they require major repairs

Remote human supervision

Operators oversee robot behavior from control rooms, reviewing sensor data, directing inspection priorities, and authorizing interventions without entering hazardous areas.

Early fault identification

Issues are detected before escalation as continuous robotic patrols identify subtle changes in equipment condition that would be missed during infrequent manual inspections.

Multisensor anomaly detection

AI interprets gas, temperature, and visual signals to identify developing problems such as gas concentration changes indicating leaks, thermal anomalies suggesting equipment stress, or visual evidence of corrosion and structural degradation.

Autonomous navigation in hazards

Drones handle overhead access and unstable surfaces remotely while robotic hand platforms are ruggedized to approach and work in short range hazardous pockets where wheeled AMRs cannot.

Tags

Compliance & Risk

Physical AI

Autonomous inspection and intervention in safety-critical physical environments

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

These systems are deployed specifically because human inspection is too dangerous or infrequent, which means a missed gas leak, thermal anomaly, or structural defect does not just represent underperformance—it produces false assurance that can be more dangerous than no monitoring at all. AI detection models must be validated continuously against the specific degradation modes that define these environments (e.g., dust, heat, humidity, and gas interference), not just against controlled laboratory benchmarks.



Safe and secure

Autonomous systems with physical intervention capabilities operating inside critical energy infrastructure are some of the most tempting targets for cyberattack in the physical AI landscape. Intervention workflows must require human confirmation that cannot be bypassed by automated processes or externally injected commands under any circumstances.



Transparent and explainable

Remote operators supervising autonomous inspection systems from control rooms must be able to understand why the AI has flagged a particular anomaly, how confident the detection is, and what intervention is being recommended. Alert and recommendation outputs should include the sensor evidence, detection logic, and confidence level that drove them, giving operators the information they need to make fully informed authorization decisions.



POTENTIAL BENEFITS

Safety improvement

Human exposure is reduced, keeping workers out of toxic, explosive, or structurally compromised environments while maintaining thorough inspection coverage.

Operational continuity

Facilities remain online during inspections since robotic systems can work without requiring equipment shutdowns, safety lockouts, or evacuation zones that halt production.

Cost reduction

Emergency repairs decline as proactive identification prevents expensive unplanned shutdowns, rush repairs with premium labor costs, and production losses from extended outages.

Failure avoidance

Early detection prevents incidents by identifying minor equipment degradation before it escalates into catastrophic failures that threaten lives and operations.



Precision-critical high-value manufacturing

AI-guided perception and actuation in safety-critical assembly environments

Physical AI systems combine sensing, reasoning, and robotic actuation to execute precision-critical assembly tasks in environments where tolerance, safety, and reliability are non negotiable. These systems continuously perceive alignment, force, and surface conditions; reason about acceptable operating envelopes; and execute physical actions under human supervision—enabling consistent outcomes at a scale and precision difficult to achieve manually.

ISSUE/OPPORTUNITY

Heavy industrial manufacturing of complex assets such as aircraft, ships, and spacecraft involves precision-critical assembly operations that must meet extremely tight tolerances and rigorous safety standards. Workers perform repetitive tasks like drilling holes, installing fasteners, applying sealants, and conducting dimensional inspections across massive structures where even minor deviations can compromise structural integrity or safety certification. Manual execution introduces variability as worker fatigue, attention levels, and individual technique affect quality consistency. For example, an aircraft fuselage

may require tens of thousands of fasteners installed at precise torque specifications, while spacecraft components demand micron-level precision that challenges human capability.

Quality defects discovered late in production trigger expensive rework that can delay delivery schedules, while undetected errors create catastrophic safety risks. The opportunity is to deploy robotic systems for repetitive precision tasks while maintaining human oversight for critical decision points, combining automation consistency with experienced technician judgment.

Tags

Manufacturing & Quality

Physical AI

HOW AI CAN HELP

Precision perception

AI detects fine tolerances through advanced sensing that measures dimensions, alignment, and surface conditions at levels exceeding human visual capability, ensuring that components meet exacting specifications.

Human-supervised operation

Critical steps are overseen by experienced technicians who monitor robot performance, verify quality checkpoints, and intervene when situations require judgment or fall outside normal parameters.

Repeatable execution

Robots perform consistent actions with identical technique, torque application, and positioning across thousands of repetitions, eliminating the variability introduced by human fatigue or attention lapses.

Integration with tooling

Robots align with existing processes and production workflows, working alongside human teams without requiring complete facility redesigns or replacement of proven manufacturing methods.

Precision-critical high-value manufacturing

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

The entire value proposition of this system—that AI-guided robotic execution is more consistent and precise than human performance—must be demonstrable to regulators and certification bodies before it can be deployed in industries such as aerospace and defense. Validation must cover the full range of production conditions, not just the configurations present during initial commissioning.



Responsible and accountable

Full traceability is a regulatory requirement in aerospace and defense manufacturing. Every AI-guided measurement, execution step, and human verification sign-off must be logged with enough detail and accuracy to reconstruct exactly what happened and who was responsible.



Transparent and explainable

Although human-supervised operation is a core design requirement, such oversight only works if technicians can find out what the AI measured, what tolerance deviation was detected, and why a specific action was recommended. Technicians who cannot meaningfully challenge or verify AI outputs cannot provide genuine oversight.



POTENTIAL BENEFITS

Throughput gains

Predictable throughput under safety and quality constraints. Production accelerates as robots work continuously without breaks, maintain consistent cycle times, and eliminate delays.

Labor efficiency

Manual repetition is reduced, allowing human expertise focused on judgment intensive tasks, quality verification, and problem-solving that requires human expertise rather than physically demanding repetitive operations.

Cost control

Rework is minimized as consistent robot execution reduces defects that would otherwise require expensive repairs, schedule delays, and in severe cases, scrapping of high-value components or assemblies.

Quality consistency

Errors decline through repeatable execution that eliminates common defect sources—such as improper torque, misalignment, or inconsistent material application—that occur with manual operations.



Autonomous agriculture and precision farming

AI-driven agricultural field operations

Physical AI systems orchestrate drones and autonomous ground robots to continuously sense field conditions, reason over crop and environmental signals, and execute targeted physical interventions under human supervision.

ISSUE/OPPORTUNITY

Agricultural operations face mounting challenges from labor shortages, rising input costs, and the need to increase productivity on limited arable land. Farmers struggle to find seasonal workers for labor-intensive tasks like weeding, thinning, and harvesting, while blanket application of seeds, fertilizers, pesticides, and water across entire fields wastes resources and increases environmental impact.

Manual field monitoring is time-consuming and imprecise, as farmers walk fields to visually assess crop health, often missing early signs of problems until they've spread. Traditional precision agriculture relies on periodic data collection and manual decision making, which limits responsiveness to rapidly changing field conditions. The opportunity is to deploy physical AI systems that continuously interpret environmental signals and coordinate physical actions in real time—scaling expert judgment across vast agricultural operations without requiring constant human presence.

HOW AI CAN HELP

Environmental sensing

Drones run scheduled multispectral/thermal passes to map crop stress and pest hotspots while AMRs carry close range sensors to validate and collect plant level samples.

Selective intervention

Actions target specific areas or individual plants based on actual need rather than blanket treatment, applying herbicides only where weeds are detected or adjusting seeding density based on soil quality.

Closed loop physical task execution

AMRs and drones navigate fields independently, performing seeding, weeding, spraying, or harvesting—while continuously validating outcomes against expected results.

Simulation and model driven intervention policies

Intervention strategies are validated against historical field data and seasonal simulations before deployment, reducing risk and improving predictability in large scale agricultural operations.

Governed human in the loop control

Operators approve intervention policies, review exceptions, and retain authority over high impact decisions, ensuring safety, compliance, and environmental accountability.

Autonomous agriculture and precision farming

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

AI models trained on data from specific geographies, crop varieties, or seasons might struggle to handle different situations. Before autonomous systems can be trusted to operate at scale, they must be validated against the full range of climates, soil types, and crop conditions present in the targeted deployment environment.



Private

Farm data generated by continuous drone and sensor monitoring is commercially sensitive information that could reveal the economic position of individual farm operators. In an industry where large agricultural technology vendors actively seek to aggregate and monetize field data, governance policies must clearly define what is collected, who owns it, and whether vendors can use farm data to train shared models without explicit consent.



Fair and impartial

AI systems developed and validated primarily on large commercial operations may systematically underserve smaller, more diverse, or mixed-crop farms—potentially recommending strategies optimized for commodity monocultures that are poorly suited to other contexts.



POTENTIAL BENEFITS

Yield optimization

Crop outcomes improve through precise intervention that addresses plant needs at the right time and location, reducing stress from over or under-treatment and maximizing productive plant growth.

Input efficiency

Resources are used more precisely as targeted application reduces waste from blanket treatment, lowering costs for seeds, fertilizers, pesticides, and water while minimizing environmental runoff.

Scalability

Large areas are managed efficiently as autonomous systems can cover extensive acreage with consistent quality, enabling farmers to increase operational scale without proportional increases in labor overhead.

Labor reduction

Manual work declines as autonomous systems handle repetitive field tasks, reducing dependence on seasonal labor that is increasingly difficult to recruit and retain.



Workforce scheduling and dispatch

Predictive coordination of physical repair crews

Physical AI models forecast likely network faults using historical failures, environmental signals, and asset data, continuously updating field-force scheduling to dispatch the right engineer by skill, proximity, and urgency. Next-gen inspection wearables with lightweight industrial glasses gives technicians hands-free, real-time guidance, equipment data, and AI diagnostics via heads-up overlays, replacing tablets and manuals and enabling fully hands-free maintenance.

ISSUE/OPPORTUNITY

Field service operations are costly and hard to optimize because failures are unpredictable, assets are geographically dispersed, and skill needs vary, so static schedules and manual dispatch drive inefficiencies, longer repairs, and avoidable travel. This is especially acute in energy, industrial, and resource operations spread across vast areas, where dispatchers juggle specialized crews, competing emergencies, and long distances that can consume much of a technician's day, making rapid response critical to avoid safety incidents and production losses.

In the field, technicians still depend on tablets, phones, or paper manuals for procedures and diagnostics, forcing constant context-switching that slows work and increases safety risk; earlier AR alternatives were too bulky and unreliable for industrial use.

The opportunity is to use physical AI to anticipate where physical interventions will be needed and dynamically coordinate field resources to respond faster and at lower cost.

Tags

Field Services

Physical AI

HOW AI CAN HELP

Fault likelihood prediction

AI models forecast where physical equipment issues are most likely to occur based on asset age, operating conditions, maintenance history, and environmental stressors such as weather or load patterns.

Skill-based assignment

Engineers are matched to jobs based on capability, ensuring specialized equipment such as high-voltage transformers or pressure vessels receives attention from appropriately certified technicians.

Schedule optimization

Predictions feed directly into workforce scheduling systems, enabling proactive assignment of technicians to areas where failures are anticipated rather than waiting for emergency calls.

Manager oversight

Supervisors retain control over final decisions, reviewing AI recommendations and adjusting assignments based on factors the system cannot assess such as customer relationships or operational priorities.

Real-time visual guidance overlay

AR glasses display step-by-step repair instructions, measurements, and equipment specifications directly overlaid on the physical equipment technicians are working on, eliminating the need to reference separate screens.

Continuous schedule updates

Assignments are refreshed frequently as conditions change, rerouting technicians when higher-priority emergencies arise or when predicted failures materialize sooner than expected.

Remote expert collaboration

Technicians can share their AR field of view with remote experts who can provide guidance, annotate the technician's view, and collaborate on complex repairs without traveling to site.

Workforce scheduling and dispatch

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

AI predictions that consistently misdirect field crews—deploying technicians to the wrong locations or failing to anticipate failures that subsequently cause outages—undermine the operational case for the system and erode dispatcher trust. Model accuracy must be validated continuously against actual failure outcomes across the entire asset portfolio, accounting for seasonal variation and aging infrastructure that changes risk profiles over time.



Transparent and explainable

Human supervisor oversight is a core design requirement, but true oversight is only possible if dispatchers and technicians can understand why the system made a particular recommendation. When scheduling logic is opaque, workers cannot meaningfully challenge assignments, supervisors cannot make genuinely informed adjustments, and human-in-the-loop design becomes a talking point rather than a genuine safeguard.



Fair and impartial

Scheduling algorithms optimized primarily for operational efficiency can distribute workloads, travel burdens, and undesirable assignments unequally across the workforce if equitable treatment is not built into the optimization criteria. Historical dispatch patterns used to train or calibrate the system may reflect prior management decisions that were themselves inequitable, embedding those patterns into automated recommendations that appear objective but replicate unfair outcomes at scale.

POTENTIAL BENEFITS

Productivity gains

Engineer utilization improves as optimized routing reduces idle travel time and ensures technicians are dispatched to locations where their specific skills are needed most.

Cost reduction

Travel inefficiencies decline as AI minimizes unnecessary mileage, consolidates nearby service calls into single trips, and reduces emergency overtime by enabling proactive maintenance.

Service reliability

Issues are resolved sooner through better resource allocation, preventing minor problems from escalating into major outages that affect production, safety, or customer service.

Faster physical repairs

Response times decrease through predictive positioning of crews near anticipated failure locations and intelligent routing that accounts for real-time traffic and site access constraints.

Increased technician productivity

Hands-free access to information and guidance enables technicians to work faster and more efficiently, eliminating time spent searching manuals, measuring equipment, or switching between tools and tablets.

Inspection of network and utility infrastructure

Vision-based inspection of physical assets

Vision-driven physical AI systems use satellite imagery, LiDAR, drones, and environmental data. These systems monitor poles, towers, substations, cables, and surrounding environments to detect physical degradation, vegetation encroachment, and environmental risks earlier and more consistently than manual inspections.

ISSUE/OPPORTUNITY

Manual inspection of distributed infrastructure is slow, costly, and infrequent, especially in remote areas. Failures often originate from gradual physical degradation. Utility companies maintain thousands of miles of transmission lines, distribution poles, and towers spread across diverse terrain including forests, mountains, and farmland where access requires specialized vehicles and significant travel time. Inspectors visually examine poles for rot, cracks, or corrosion, check cables for fraying, assess vegetation encroachment, and identify structural issues with towers.

Manual inspection cycles may occur only every few years due to cost, allowing problems to develop undetected between visits.

Critical infrastructure in hard-to-reach locations receives even less frequent attention, increasing the risk of unexpected failures. Weather events, wildlife damage, and natural aging cause continuous degradation that accumulates between inspection cycles.

HOW AI CAN HELP

Vision based inspection at scale

Drones and mobile inspection platforms capture high resolution visual, thermal, and LiDAR data across vast and hard to reach infrastructure, enabling frequent inspection without physical site access.

Large-area coverage

Assets across wide geographies are inspected systematically through aerial surveys that can cover hundreds of miles of infrastructure in days rather than months required for ground-based inspection.

AI driven defect and risk detection

Computer vision models identify early signs of asset degradation—such as corrosion, structural fatigue, conductor wear, or vegetation proximity—before issues escalate into outages or safety incidents.

Human validation

Inspectors review results before authorizing repairs, verifying AI findings and applying expertise to distinguish urgent problems from conditions that can be monitored.

Risk based prioritization of interventions

AI ranks detected issues based on severity, asset criticality, environmental conditions, and potential customer impact, enabling field teams to focus resources on the highest risk assets first.

Tags

Field Services

Physical AI

Inspection of network and utility infrastructure

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

The system's value relies on AI inspection being more consistent and comprehensive than infrequent manual visits—which means a model that underperforms doesn't just reduce efficiency; it produces false assurance in the locations where the inspection gap is most dangerous. Validation must explicitly cover the full range of environmental and asset conditions across the deployment territory, not just those well-represented in training data.



Responsible and accountable

When AI-driven risk prioritization determines which assets receive attention first, and a deprioritized asset subsequently fails, the basis for that triage decision must be auditable and defensible—to regulators, affected customers, and the public. The human validation step must constitute genuine review, with clear documentation of what inspectors assessed and on what basis, rather than nominal sign-off on outputs they cannot meaningfully interrogate.



Fair and impartial

AI prioritization optimized for asset criticality and customer impact risks systematically directing maintenance resources toward urban or commercially significant infrastructure at the expense of rural or lower-density areas where assets may be older, more vulnerable, and failures more isolating for the customers affected. Utility operators have public service obligations across their entire territory, and AI-driven triage should be regularly audited to ensure those obligations are being met equitably.



POTENTIAL BENEFITS

Lower inspection cost

Fewer manual visits required as aerial inspection covers large areas quickly, reducing vehicle expenses, labor hours, and specialized equipment needed for accessing remote terrain.

Improved asset reliability

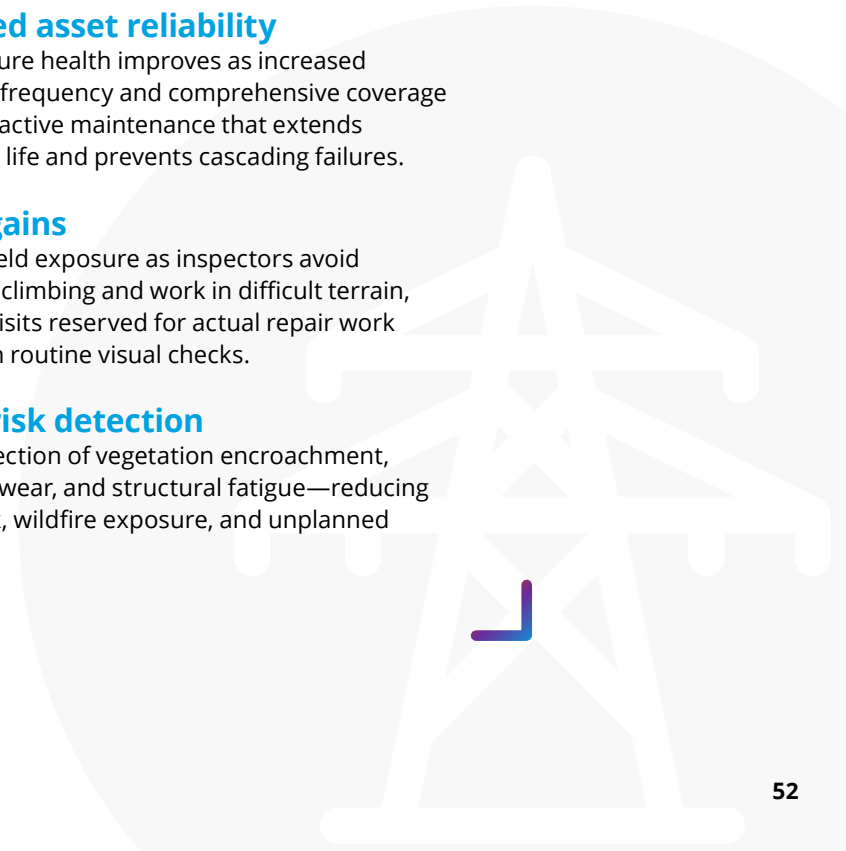
Infrastructure health improves as increased inspection frequency and comprehensive coverage enable proactive maintenance that extends equipment life and prevents cascading failures.

Safety gains

Reduced field exposure as inspectors avoid hazardous climbing and work in difficult terrain, with field visits reserved for actual repair work rather than routine visual checks.

Earlier risk detection

Earlier detection of vegetation encroachment, conductor wear, and structural fatigue—reducing outage risk, wildfire exposure, and unplanned shutdown.



Defect detection for industrial machinery

Vision-supported inspection with human validation

AI-enabled vision systems support the end-to-end repair lifecycle of large industrial components (e.g., turbine parts and blades) by automating intake identification, augmenting defect detection during inspection, and enabling robotic inspection and selective repair, while keeping human experts in the decision loop for all safety critical judgments. The solution combines computer vision, operational context, and robotics to improve consistency, coverage, safety, and cycle time across repair and maintenance operations.

ISSUE/OPPORTUNITY

Manual inspection of complex industrial components is time-consuming, requires specialized expertise, and is subject to variability across inspectors, shifts, and facilities. Defects may be missed during initial inspection or identified late in the repair cycle, increasing rework costs and extending equipment downtime. Yet, fully automated inspection systems are not reliable enough for safety-critical applications due to false positives from lighting variations, surface contamination, and reflection artifacts, as well as false negatives that miss subtle defects in challenging viewing conditions.

Augmenting human inspectors with AI-based vision systems can improve coverage and consistency while retaining human decision authority for final accept/reject determinations—combining the pattern recognition capabilities of computer vision with the judgment and accountability that expert inspectors provide.

Tags

Manufacturing & Quality

Physical AI

HOW AI CAN HELP

Defect pattern recognition

Vision models analyze component imagery to highlight cracks, wear patterns, erosion, and anomalous surface conditions that warrant closer inspector review, improving detection of subtle defects that might otherwise be missed.

Human-in-the-loop confirmation

All AI outputs are treated as recommendations requiring expert validation, preserving inspector judgment and accountability for final component disposition.

Correlation with operational history

AI connects observed defects to field operating conditions, sensor data, and maintenance history to provide contextual clues for root cause analysis, helping inspectors understand whether damage patterns align with expected wear or indicate abnormal operating conditions.

Inspection prioritization

AI directs inspector attention to high-risk regions based on component type, operational history, and observed patterns,

improving the efficiency of limited inspection resources by focusing expert review where it matters most.

Bounded decision support

AI is constrained to assistance and highlighting roles rather than automated accept/reject authority, maintaining human responsibility for safety-critical determinations.

Robustness to inspection artifacts

Models trained on diverse imagery learn to reduce sensitivity to lighting variations, glare, surface contamination, and reflection artifacts that create false positives, improving the signal-to-noise ratio for human reviewers.

Defect pattern recognition across inspection stages

Computer vision models analyze imagery from manual, robotic, or in situ inspections to highlight cracks, erosion, wear, delamination, and anomalous surface conditions—directing inspector attention to high-risk regions that warrant expert review.

Defect detection for industrial machinery

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

False positives waste inspector time and erode trust; false negatives allow genuine defects to pass undetected through safety-critical components. Both have serious consequences in critical contexts such as turbine and industrial component repair. Robustness must be validated across the full range of lighting conditions, surface states, and component types that the system will encounter, not just favorable training conditions.



Responsible and accountable

In this use case, humans retain final accept/reject authority. However, this only provides genuine protection if applied in practice. Inspection records must document which defects were AI-flagged, what human experts concluded, and who made the final decision—creating an audit trail sufficient for regulatory compliance and liability determination.



Transparent and explainable

The human-in-the-loop design only works if inspectors can understand what was flagged and why—enabling true judgment rather than default deference to opaque outputs. When AI correlates observed defects with operational history to suggest root causes, inspectors need sufficient visibility into that reasoning to assess whether it is contextually sound before acting on it.

POTENTIAL BENEFITS

Earlier defect identification

Potential issues are surfaced sooner in the repair cycle, enabling proactive repair planning and reducing the risk that defects progress undetected through multiple inspection stages.

Quality consistency

Lower variance in detection rates across inspectors, shifts, and facilities improves repeatability of inspection outcomes and supports more predictable repair planning and component lifecycle management.

Higher inspection coverage

AI-assisted systems reduce the variability introduced by inspector workload, fatigue, and experience differences, leading to more consistent and thorough reviews across components, sites, and time periods.

Engineering efficiency

Reduced time spent on low-risk scanning and documentation allows inspectors to focus expertise on complex assessments, root cause analysis, and judgment-intensive decisions.

Reduced cycle time and downtime

Automated intake and prioritized inspection accelerate engineering triage and repair planning, shortening asset turnaround and turbine shutdown durations.

Extended asset life and lower cost

Earlier detection and targeted repairs prevent damage progression, reducing premature component replacement and overall maintenance cost.

Chemical manufacturing workflow automation

Autonomous handoff control for chemical lines

Embedded AI (including physical-AI systems on robots and equipment) manage process handoffs across the manufacturing lifecycle, autonomously coordinating materials, sequencing tasks, allocating resources, and resolving routine exceptions to eliminate bottlenecks and accelerate throughput. AMRs handle materials and inter stage movement while robotic hands perform valve turns, sampling, and minor adjustments at equipment handoffs without human entry.

ISSUE/OPPORTUNITY

Chemical manufacturing involves hundreds of process handoffs where information, materials, or decisions transfer between steps or personnel. When one production stage completes, operators must notify the next team, confirm material availability, verify equipment readiness, and coordinate timing—creating coordination bottlenecks that slow overall throughput. Each handoff creates opportunities for delays, errors, and inefficiencies as workers wait for approvals, materials sit idle between stages, or miscommunication causes incorrect sequencing.

Manual coordination of these workflows consumes significant labor hours and limits production speed as supervisors spend time orchestrating activities rather than addressing exceptions or improving processes. Traditional automation only addresses repetitive individual tasks such as valve control or temperature

regulation, not the complex decision-making and coordination required across end-to-end processes that involve judgment about priorities, resource allocation, and exception handling.

Fragmented, machine-specific automation solutions further increase complexity. Different OEM equipment often uses isolated control logic, making integration across lines and sites difficult, slowing replication of improvements, and increasing engineering effort to scale operations.

Deploying a physical-AI-first workflow automation and AI platform—embedded on robots, AMRs, quadrupeds, autonomous vehicles, and process equipment—creates an opportunity to unlock latent capacity, materially increase productivity, and scale output without proportional labor growth.

HOW AI CAN HELP

Physical task sequencing

Physical AI coordinates AMR movement ensuring materials, samples, and intermediate products move seamlessly without manual scheduling or waiting periods.

Human-agent hybrid processes

All safety critical actions and non standard conditions require human approval. Physical AI operates within defined safety envelopes, preserving accountability and regulatory compliance.

Unified physical control across heterogeneous equipment

Physical AI coordinates across OEM machines, robots, and material handling systems, enabling system level optimization rather than isolated, machine by machine automation.

Real-time process-aware sequencing

AI models continuously re-sequence physical tasks based on throughput, safety constraints, and equipment status—reducing idle time between process steps and preventing downstream bottlenecks.

Scalable, repeatable operations

Standardized physical AI control logic enables consistent execution across lines and sites, supporting scale-up without re-engineering workflows or adding coordination layers.

Unified AI control layer across equipment

A standardized AI control architecture spans multiple machine types and OEMs, replacing isolated point solutions and enabling system-level optimization rather than machine-by-machine tuning.

Tags

Manufacturing & Quality

Physical AI

Chemical manufacturing workflow automation

MANAGING RISK AND PROMOTING TRUST



Robust and reliable

Autonomous coordination of process handoffs in chemical manufacturing is not a context where degraded performance is merely inconvenient. Incorrect sequencing of hazardous materials between process stages can create conditions that are difficult to reverse quickly. Validation must cover heterogeneous equipment combinations, edge cases, and the full range of exception scenarios the system will likely encounter in production.



Safe and secure

The physical actions this system controls—valve operations, material transfers, robotic movements—make the integrity of the AI control layer a safety issue, not just a cybersecurity issue. Unauthorized access or manipulation of process sequencing logic in a chemical facility could trigger dangerous physical actions. Security protections must be commensurate with the hazard level of the processes being coordinated.



Responsible and accountable

Human approval of safety-critical actions and non-standard conditions only provides true protection if the audit trail is accurate and complete. In a regulated chemical manufacturing environment, documentation of which decisions were executed autonomously, which were escalated, and who approved them is not optional governance; it is a regulatory requirement and the basis for liability determination.



POTENTIAL BENEFITS

50%+ productivity improvement

Eliminate most handoff delays and coordination overhead, dramatically accelerating production cycles by removing waiting periods between process stages.

System level optimization across equipment

Physical AI provides a unified control layer across heterogeneous OEM equipment, robots, and AMRs—optimizing workflows at the line and plant level rather than machine by machine.

Improved repeatability across sites

Standardized agent and control logic delivers consistent process outcomes across production lines and facilities.

Error reduction

Remove human errors at process transitions through consistent, automated handoffs and decision execution that apply the same logic every time.

Lower integration effort

Reduced bespoke engineering and interface work through unified AI orchestration and control layers.



Industrial facilities and space optimization

AI and sensor-driven building efficiency

Physical AI sensors and edge devices across buildings stream occupancy and environmental signals to edge AI controllers that optimize lighting, HVAC, and space use in near real time.

ISSUE/OPPORTUNITY

Office buildings are often inefficient due to static operating rules and limited visibility into real usage patterns. Lighting, heating, and cooling systems typically run on fixed schedules that ignore actual occupancy, wasting energy in empty spaces while failing to adjust for peak usage periods when demand exceeds capacity. Manual processes lead to wasted energy and underutilized space as facilities teams lack real-time visibility into how buildings are actually being used versus how scheduling systems suggest they should be used, preventing optimization based on observed patterns.

The opportunity is to operate buildings more responsively using real occupancy data while preserving safety and manual override capabilities that allow facilities managers to intervene when AI-driven adjustments conflict with operational needs or emergency situations.

Tags

Operations

Physical AI

HOW AI CAN HELP

Occupancy pattern analysis

AI combines sensor, access, and booking data to learn true space utilization and identify underused areas.

Meeting-room utilization optimization

AI aligns booking systems with actual presence to reduce friction between scheduled and actual usage.

Dynamic environment adjustment

Lighting and energy systems are adjusted based on real usage rather than schedules, reducing waste in unoccupied areas.

Localized execution

Decisions are executed close to building systems for responsiveness without requiring centralized coordination.

Override and safety controls

Facility teams retain manual override to pause AI adjustments when operational or safety needs require intervention.

Energy demand forecasting

Models predict consumption and identify opportunities to reduce waste through pattern analysis.

Industrial facilities and space optimization

MANAGING RISK AND PROMOTING TRUST



Private

Continuous occupancy monitoring generates detailed records of how individuals move through building spaces throughout the working day—more than most occupants would expect from a facilities optimization tool. Data should be used exclusively for building management purposes, with clear retention limits and explicit communication to occupants about what is monitored and how their data is used.



Responsible and accountable

Automated adjustments affect occupant comfort and well-being in ways that are immediately felt but not easily traced to their cause. Facilities teams must maintain meaningful oversight, with accessible escalation paths when adjustments cause problems. Accountability for building conditions must remain with qualified facilities managers—not be effectively delegated to an optimization system that occupants cannot see or challenge.



Fair and impartial

Optimization that prioritizes high-traffic areas may systematically deliver inferior lighting, temperature, or air quality to workers in lower-utilization spaces—based on where they sit, rather than their actual needs. This is a foreseeable and preventable outcome of efficiency-first design. Organizations should actively monitor whether AI-driven optimization produces equitable environmental conditions across the building population.



POTENTIAL BENEFITS

Energy savings

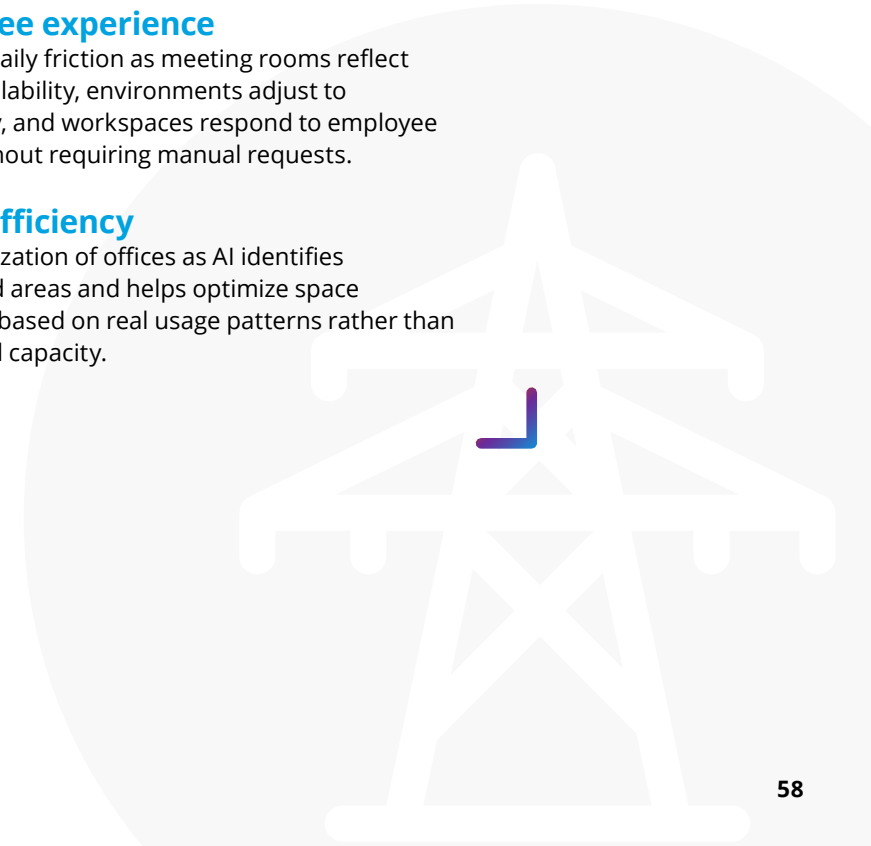
Lower operating costs through optimized heating, cooling, and lighting that respond to actual occupancy rather than fixed schedules or conservative assumptions.

Employee experience

Reduced daily friction as meeting rooms reflect actual availability, environments adjust to occupancy, and workspaces respond to employee needs without requiring manual requests.

Space efficiency

Better utilization of offices as AI identifies underused areas and helps optimize space allocation based on real usage patterns rather than theoretical capacity.



Autonomous and semi-autonomous robotics in construction

Physical AI-enabled execution across dynamic construction sites

Autonomous and semi autonomous machines are taking on essential construction duties—earthmoving, grading, rebar tying, bricklaying, and on site additive manufacturing. The rising “print to build” approach is tightening the link between physical systems and digital twins, streamlining design to field transfers for MEP (mechanical, electrical, plumbing) deliverables and cutting both handoff time and execution errors when AI driven equipment performs on site tasks.

ISSUE/OPPORTUNITY

Robotics in construction represents a pressing issue and a strategic opportunity: persistent labor shortages, rising project complexity, and the demand for greater efficiency threaten schedules, budgets, and safety if traditional methods continue. At the same time, integrating robotics with data capture and AI offers a clear path to fill

workforce gaps, boost on-site productivity and accuracy, reduce safety risks, and enable predictive, data-driven planning and management—shifting projects from manual, fragmented workflows to integrated, scalable delivery models.

Tags

Operations

Physical AI

HOW AI CAN HELP

Automated bricklaying systems

Robotic bricklayers can place up to roughly 3,000 bricks per day—well beyond typical manual output—taking on repetitive masonry so experienced masons can concentrate on complex or detail work.

3D printing and additive manufacturing

Large-format robotic 3D printers can build entire structures in days, dramatically shortening construction timelines. Demonstration projects show this approach can make housing delivery faster and more resource-efficient.

Demolition and renovation robots

Radio- or remote-controlled demolition robots handle hazardous, confined-space work with precision, lowering worker exposure to danger and improving control over selective demolition and refurbishment tasks. operating conditions.

Model-to-field execution

Digital twins and BIM models are directly translated into physical actions, reducing interpretation errors and rework.

Surveying and monitoring with drones

AI-enabled drones fitted with advanced optical sensors generate detailed 3D site maps within hours, accelerating site surveys, progress tracking, and safety inspections so teams can spot and address risks earlier.

Autonomous and semi-autonomous robotics in construction

MANAGING RISK AND PROMOTING TRUST



Safe and secure

Construction sites are among the most hazardous working environments, and autonomous systems that fail to reliably detect human proximity or enforce operational boundaries create exactly the injury risk the technology is meant to reduce. Safety boundaries must be validated under actual site conditions—not just controlled testing—with robust mechanisms for responding to unexpected human entry into autonomous zones.



Responsible and accountable

Construction is already a complex multi-party contracting environment before AI is introduced. When autonomous equipment causes structural errors, property damage, or worker injury, accountability across developer, manufacturer, integrator, and contractor is even harder to resolve. Operational logs capturing system decisions and human oversight actions must be constantly maintained—not assembled retrospectively when an incident demands explanation.



Fair and impartial

Workforce displacement in construction is immediate and concentrated; bricklaying, rebar tying, and demolition represent significant employment for lower-skilled workers with limited alternative pathways. Productivity gains from autonomous systems should not come at the expense of this workforce without meaningful investment in transition support and reskilling.



POTENTIAL BENEFITS

Increased efficiency

Automation refines workflows, lowers staffing expenses, and shortens project schedules by handling repetitive tasks more quickly and consistently.

Improved quality

Automated systems deliver repeatable precision and consistency, which reduces errors, defects, and costly rework.

Higher sustainability

Robotics can reduce material waste and energy use, supporting greener construction methods and lower environmental impact.

Enhanced safety

Robots take on dangerous activities and confined-space work, cutting worker exposure and reducing the incidence of injuries.

Long-term cost savings

While capital costs can be significant up front, robotics often deliver lifecycle savings through reduced labor needs, less waste, and faster completion times.



Contacts



Beena Ammanath
Global Deloitte AI Institute
Leader, Deloitte AI Institute
United States, Lead
Deloitte Consulting LLP
bammanath@deloitte.com



Tiago Durao
Deloitte AI Institute
Portugal, Lead
Deloitte Portugal
tdurao@deloitte.com



Geert Hallemeesch
Deloitte AI Institute
Belgium, Lead
Deloitte Belgium
ghallemeesch@deloitte.com



Chris Lewin
Deloitte AI Institute
Asia Pacific Lead
Deloitte Singapore
chrislewin@deloitte.com



Audrey Ancion
Deloitte AI Institute
Canada, Lead
Deloitte Canada
knuttall@deloitte.com



Richard Eudes
Deloitte AI Institute
France, Lead
Deloitte France
reudes@deloitte.fr



Jan Hejtmanek
Deloitte AI Institute
Central Europe, Lead
Deloitte Central Europe
jhejtmanek@deloitte.com



Wessel Oosthuizen
Deloitte AI Institute
Africa, Lead
Deloitte Africa
woosthuizen@deloitte.com



Naser Bakhshi
Deloitte AI Institute
Netherlands, Lead
Deloitte Netherlands
nbakhshi@deloitte.nl



Roman Fan
Deloitte AI Institute
China, Lead
Deloitte China
rfan@deloitte.com



Prashanth Kaddi
Deloitte AI Institute
India, Lead
Deloitte India
pkaddi@deloitte.com



Sulabh Soral
Deloitte AI Institute
United Kingdom, Lead
Deloitte United Kingdom
ssoral@deloitte.com



Dr. Bjoern Bringmann
Deloitte AI Institute
Germany, Lead
Deloitte Germany
bbringmann@deloitte.com



Tomas Meca Figueres
Deloitte AI Institute
Spain, Lead
Deloitte Spain
tomecafigueras@deloitte.es



Sultanbek Khunkaev
Deloitte AI Institute
Middle East, Lead
Deloitte United Arab Emirates
sukhunkaev@deloitte.com



Dr. Elea Wurth
Deloitte AI Institute
Australia, Lead
Deloitte Australia
ewurth@deloitte.com.au



Martin Cabrera
Deloitte AI Institute
Chile, Lead
Deloitte
mcabreraa@deloitte.com



Alfredo Maria Garibaldi
Deloitte AI Institute
Italy, Lead
Deloitte Italy
agaribaldi@deloitte.it



Tomotake Kozu
Deloitte AI Institute
Japan, Lead
Deloitte Japan
tomotake.kozu@tohatsu.co.jp



Jefferson Denti
Deloitte AI Institute
Brazil, Lead
Deloitte Brazil
jdenti@deloitte.com



Nicolas Griedlich
Deloitte AI Institute
Luxembourg, Lead
Deloitte Luxembourg
ngriedlich@deloitte.lu



Carlos Labanda
Deloitte AI Institute
South-Latin America, Lead
Deloitte Colombia
clabanda@deloitte.com



This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or their related entities (collectively, the "Deloitte Network") 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 entity in the Deloitte Network shall be responsible for any loss whatsoever sustained by any person who relies on this communication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.